

Contract # AR2492**STATE OF UTAH COOPERATIVE CONTRACT**

1. CONTRACTING PARTIES: This contract is between the Division of Purchasing and the following Contractor:

Unisys Corporation
Name
801 Lakeview Drive #100
Address
Bluebell PA 19422
City State Zip

LEGAL STATUS OF CONTRACTOR

- ☐ Sole Proprietor
☐ Non-Profit Corporation
☒ For-Profit Corporation
☐ Partnership
☐ Government Agency

Contact Person Harsh Bajpai Phone #703-439-6200 Email Harsh.Bajpai@unisys.com
Vendor #34660A Commodity Code #920-05

2. GENERAL PURPOSE OF CONTRACT: Contractor is permitted to provide the Cloud Solutions identified in Attachment B to Participating States once a Participating Addendum has been signed
3. PROCUREMENT PROCESS: This contract is entered into as a result of the procurement process on Bid#CH16012.
4. CONTRACT PERIOD: Effective Date: 09/30/2016 Termination Date: 09/15/2026 unless terminated early or extended in accordance with the terms and conditions of this contract. Note: Pursuant to Solicitation #CH16012, Contract must re-certify its qualifications each year.
5. Administrative Fee, as described in the Solicitation and Attachment A: The Contractor shall pay to NASPO ValuePoint, or its assignee, a NASPO ValuePoint Administrative Fee of one-quarter of one percent (0.25% or 0.0025) no later than 60 days following the end of each calendar quarter. The NASPO ValuePoint Administrative Fee shall be submitted quarterly and is based on sales of the Services.
6. ATTACHMENT A: NASPO ValuePoint Master Terms and Conditions, including that attached Exhibits
ATTACHMENT B: Scope of Services Awarded to Contractor
ATTACHMENT C: Pricing Discounts and Pricing Schedule
ATTACHMENT D: Contractor's Response to Solicitation #CH16012
ATTACHMENT E: Microsoft Additional Use Rights and Restrictions, AWS Public Sector Access Policy, and Google Apps for Work (for Customers)
- Any conflicts between Attachment A and the other Attachments will be resolved in favor of Attachment A.**
8. DOCUMENTS INCORPORATED INTO THIS CONTRACT BY REFERENCE BUT NOT ATTACHED:
a. All other governmental laws, regulations, or actions applicable to the goods and/or services authorized by this contract.
b. Utah State Procurement Code and the Procurement Rules.
9. Each signatory below represents that he or she has the requisite authority to enter into this contract.

IN WITNESS WHEREOF, the parties sign and cause this contract to be executed.

CONTRACTOR

Harsh Bajpai Oct. 07, 2016
Contractor's signature Date


STATE

Christopher Hughes 10.6.16
Director, Division of Purchasing Date

Harsh Bajpai, Director. of Cloud & Security Sol. & Svcs
Type or Print Name and Title

<u>Christopher Hughes</u>	<u>801-538-3254</u>	<u>christopherhughes@utah.gov</u>
Division of Purchasing Contact Person	Telephone Number	Email

(Revision 16 June 2016)



Attachment A: NASPO ValuePoint Master Agreement Terms and Conditions

1. Master Agreement Order of Precedence

a. Any Order placed under this Master Agreement shall consist of the following documents:

- (1) A Participating Entity's Participating Addendum² ("PA");
- (2) NASPO ValuePoint Master Agreement Terms & Conditions, including the applicable Exhibits³ to the

Master Agreement and the cloud service providers service terms;
; and

- (3) A Service Level Agreement issued against the Participating Addendum.

b. These documents shall be read to be consistent and complementary. Any conflict among these documents shall be resolved by giving priority to these documents in the order listed above. Contractor terms and conditions that apply to this Master Agreement are only those that are expressly accepted by the Lead State and must be in writing and attached to this Master Agreement as an Exhibit or Attachment.

2. Definitions - Unless otherwise provided in this Master Agreement, capitalized terms will have the meanings given to those terms in this Section.

Confidential Information means any and all information of any form that is marked as confidential or would by its nature be deemed confidential disclosed by either Party to the other Party or its employees or agents in the performance of this Master Agreement, including, but not necessarily limited to (1) any Purchasing Entity's records, (2) personnel records, (3) information concerning individuals, human resources, financial costs and information, inventory, purchasing or merchandising plans, strategies or forecasts and (4) information relating to Contractor's techniques, software and tools used to provide the Services is confidential information of the disclosing Party.

Contractor means the person or entity providing solutions under the terms and conditions set forth in this Master Agreement. Contractor also includes its employees, subcontractors, agents and affiliates who are providing the services agreed to under the

Master Agreement. The cloud services providers providing public and hybrid cloud services are not Contractor's subcontractors or agents and the agreements between Contractor and these cloud services providers are not subcontracts.

Data means all information, (including text, sound, software or image files created by or in any way originating with a Participating Entity or Purchasing Entity, and all information that is uploaded into the Service and information that is uploaded into the Service and available and authorized by the service provider for export/download from the Service and information that is available and authorized by the service provider for Export/download from the Service that is the output of any computer processing or other electronic manipulation, of any information that was created by or in any way originating with a Participating Entity or Purchasing Entity, in the course of using and configuring the Services provided under this Agreement.

Data Breach means any actual or reasonably suspected non-authorized access to or acquisition of computerized Non-Public Data or Personal Data that compromises the security, confidentiality, or integrity of the Non-Public Data or Personal Data, or the ability of Purchasing Entity to access the Non-Public Data or Personal Data.

Data Categorization means the process of risk assessment of Data. See also "High Risk Data", "Moderate Risk Data" and "Low Risk Data".

Disabling Code means computer instructions or programs, subroutines, code, instructions, data or functions, (including but not limited to viruses, worms, date bombs or time bombs), including but not limited to other programs, data storage, computer libraries and programs that self-replicate without manual intervention, instructions programmed to activate at a predetermined time or upon a specified event, and/or programs purporting to do a meaningful function but designed for a different function, that alter, destroy, inhibit, damage, interrupt, interfere with or hinder the operation of the Services, the service provider's infrastructure, Purchasing Entity's software, applications and/or its end users processing environment, the system in which it resides, or any other software or data on such system or any other system with which it is capable of communicating.

Fulfillment Partner means a third-party contractor qualified and authorized by Contractor, and approved by the Participating State under a Participating Addendum, who may, to the extent authorized by Contractor, fulfill any of the requirements of this Master Agreement including but not limited to providing Services under this Master Agreement and billing Customers directly for such Services. Contractor may, upon written notice to the Participating State, add or delete authorized Fulfillment Partners as necessary at any time during the contract term. Fulfillment Partner has no authority to amend this Master Agreement or to bind Contractor to any additional terms and conditions.

High Risk Data is as defined in FIPS PUB 199, Standards for Security Categorization of Federal Information and Information Systems ("High Impact Data").

Infrastructure as a Service (IaaS) as used in this Master Agreement is defined the capability provided to the end user to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The end user does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Intellectual Property means any and all patents, copyrights, service marks, trademarks, trade secrets, trade names, patentable inventions, or other similar proprietary rights, in tangible or intangible form, and all rights, title, and interest therein.

Lead State means the State centrally administering the solicitation and any resulting Master Agreement(s).

Low Risk Data is as defined in FIPS PUB 199, Standards for Security Categorization of Federal Information and Information Systems ("Low Impact Data").

Master Agreement means this agreement executed by and between the Lead State, acting on behalf of NASPO ValuePoint, and the Contractor, as now or hereafter amended.

Moderate Risk Data is as defined in FIPS PUB 199, Standards for Security Categorization of Federal Information and Information Systems ("Moderate Impact Data").

NASPO ValuePoint is the NASPO ValuePoint Cooperative Purchasing Program, facilitated by the NASPO Cooperative Purchasing Organization LLC, a 501(c)(3) limited liability company (doing business as NASPO ValuePoint) is a subsidiary organization the National Association of State Procurement Officials (NASPO), the sole member of NASPO ValuePoint. The NASPO ValuePoint Cooperative Purchasing Organization facilitates administration of the cooperative group contracting consortium of state chief procurement officials for the benefit of state departments, institutions, agencies, and political subdivisions and other eligible entities (i.e., colleges, school districts, counties, cities, some nonprofit organizations, etc.) for all states and the District of Columbia. The NASPO ValuePoint Cooperative Development Team is identified in the Master Agreement as the recipient of reports and may be performing contract administration functions as assigned by the Lead State.

Non-Public Data means High Risk Data and Moderate Risk Data that is not subject to distribution to the public as public information. It is deemed to be sensitive and confidential by the Purchasing Entity because it contains information that is exempt by statute, ordinance or administrative rule from access by the general public as public information.

Participating Addendum means a bilateral agreement executed by a Contractor and a Participating Entity incorporating this Master Agreement and any other additional

Participating Entity specific language or other requirements, e.g. ordering procedures specific to the Participating Entity, other terms and conditions.

Participating Entity means a state, or other legal entity, properly authorized to enter into a Participating Addendum.

Participating State means a state, the District of Columbia, or one of the territories of the United States that is listed in the Request for Proposal as intending to participate. Upon execution of the Participating Addendum, a Participating State becomes a Participating Entity.

Personal Data means data alone or in combination that includes information relating to an individual that identifies the individual by name, identifying number, mark or description can be readily associated with a particular individual and which is not a public record. Personal Information may include the following personally identifiable information (PII): government-issued identification numbers (e.g., Social Security, driver's license, passport); financial account information, including account number, credit or debit card numbers; or Protected Health Information (PHI) relating to a person.

Platform as a Service (PaaS) as used in this Master Agreement is defined as the capability provided to the consumer to deploy onto the cloud infrastructure end user created or -acquired applications created using programming languages and tools supported by the provider. This capability does not necessarily preclude the use of compatible programming languages, libraries, services, and tools from other sources. The end user does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Product means any deliverable under this Master Agreement, including Services and software.

Protected Health Information (PHI) shall have the same meaning as stated in 45 CFR § 160.103 of HIPAA and limited to such protected health information received by service provider or transmitted, maintained, created or received by service provider on behalf of Participating/Purchasing Entity. PHI excludes education records covered by the Family Educational Rights and Privacy Act (FERPA), as amended, 20 U.S.C. 1232g, records described at 20 U.S.C. 1232g(a)(4)(B)(iv) and employment records held by a covered entity in its role as employer. PHI may also include information that is a subset of health information, including demographic information collected from an individual, and (1) is created or received by a health care provider, health plan, employer or health care clearinghouse; and (2) relates to the past, present or future physical or mental health or condition of an individual; the provision of health care to an individual; or the past, present or future payment for the provision of health care to an individual; and (a) that identifies the individual; or (b) with respect to which there is a reasonable basis to believe the information can be used to identify the individual.

Purchasing Entity means a state, city, county, district, other political subdivision of a State, and a nonprofit organization under the laws of some states if authorized by a Participating Addendum, who issues a Purchase Order against the Master Agreement and becomes financially committed to the purchase.

Services mean any of the specifications described in the Scope of Services that are supplied or created by the Contractor pursuant to this Master Agreement.

Security Incident means the actual unauthorized access to a Purchasing Entity's Non-Public Data and Personal Data the Contractor believes could reasonably result in the use, disclosure or theft of a Purchasing Entity's Non-Public Data within the possession or control of the Contractor. A Security Incident also includes a major security breach to the Contractor's system, regardless if Contractor is aware of unauthorized access to a Purchasing Entity's Non-Public Data. A Security Incident may or may not turn into a Data Breach.

Service Level Agreement (SLA) means the SLA obligations the service provider publishes in its service terms or a written agreement between both the Purchasing Entity and the Contractor that is subject to the terms and conditions in this Master Agreement and relevant Participating Addendum unless otherwise expressly agreed in writing between the Purchasing Entity and the Contractor. SLAs should include: (1) the technical service level performance promises, (i.e. metrics for performance and intervals for measure), (2) description of service quality, (3) identification of roles and responsibilities, (4) remedies, such as credits, and (5) an explanation of how remedies or credits are calculated and issued.

Software as a Service (SaaS) as used in this Master Agreement is defined as the capability provided to the consumer to use the service providers, Contractor's or third party applications available in the service provider's marketplace running on a Contractor's infrastructure (commonly referred to as 'cloud infrastructure'). The applications are accessible from various client devices through a thin client interface such as a Web browser (e.g., Web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Solicitation means the documents used by the State of Utah, as the Lead State, to obtain Contractor's Proposal.

Statement of Work means a written statement in a solicitation document or contract that describes the Purchasing Entity's service needs and expectations.

3. Term of the Master Agreement: The initial term of this Master Agreement is for ten (10) years with no renewal options.

4. Amendments: The terms of this Master Agreement shall not be waived, altered, modified, supplemented or amended in any manner whatsoever without prior written approval of the Lead State and Contractor.

5. Assignment/Subcontracts: Contractor shall not assign, sell, transfer, or sublet rights, or delegate responsibilities under this Master Agreement, in whole or in part, without the prior written approval of the Lead State.

The Lead State reserves the right to assign any rights or duties, including written assignment of contract administration duties to the NASPO Cooperative Purchasing Organization LLC, doing business as NASPO ValuePoint.

6. Discount Guarantee Period: All minimum pricing discounts are provided in Attachment C must be guaranteed for the entire term of the Master Agreement. Participating Entities and Purchasing Entities shall receive the immediate benefit of price or rate reduction of the services provided under this Master Agreement.

7. Termination: Unless otherwise stated, this Master Agreement may be terminated by either party upon 90-days written notice prior to the effective date of the termination. Further, any Participating Entity may terminate its participation upon 60 days written notice, unless otherwise limited or stated in the Participating Addendum. Any termination under this provision shall not affect the rights and obligations attending orders outstanding at the time of termination, including any right of any Purchasing Entity to indemnification by the Contractor, rights of payment for Services delivered and accepted, data ownership, Contractor obligations regarding Purchasing Entity Data, rights attending default in performance an applicable Service Level of Agreement in association with any Order, Contractor obligations under Termination and Suspension of Service, and any responsibilities arising out of a Security Incident or Data Breach. Termination of the Master Agreement due to Contractor default may be immediate if defaults cannot be reasonably cured as allowed per the Default and Remedies provision.

8. Confidentiality, Non-Disclosure, and Injunctive Relief

a. Confidentiality. Each party acknowledges that it and its employees or agents may, in the course of providing a Product under this Master Agreement, be exposed to or acquire information that is confidential to the other Party or Purchasing Entity's or Purchasing Entity's clients. Any reports or other documents or items (including software) that result from the use of the Confidential Information by the receiving Party shall be treated in the same manner as the Confidential Information. Confidential Information does not include information that (1) is or becomes (other than by disclosure by the receiving Party) publicly known; (2) is furnished by the disclosing Party to others without restrictions similar to those imposed by this Master Agreement; (3) is rightfully in the receiving Party's possession without the obligation of nondisclosure prior to the time of its disclosure under this Master Agreement; (4) is obtained from a source other than the disclosing Party without the obligation of confidentiality, (5) is disclosed with the written consent of the disclosing Party or; (6) is independently developed by employees,

agents or subcontractors of the receiving Party—who can be shown to have had no access to the Confidential Information.

b. Non-Disclosure. Each Party shall hold Confidential Information in confidence, using at least the industry standard of confidentiality, and shall not copy, reproduce, sell, assign, license, market, transfer or otherwise dispose of, give, or disclose Confidential Information to third parties or use Confidential Information for any purposes whatsoever other than what is necessary to the performance of Orders placed under this Master Agreement. Contractor shall advise each of its employees and agents of their obligations to keep Confidential Information confidential. In the provision of private and hybrid cloud services, Contractor shall use commercially reasonable efforts to assist Purchasing Entity in identifying and preventing any unauthorized use or disclosure of any Confidential Information. Without limiting the generality of the foregoing, Contractor shall advise Purchasing Entity, applicable Participating Entity, and the Lead State immediately if Contractor learns or has reason to believe that any Contractor employee who has had access to Confidential Information has violated or intends to violate the terms of this Master Agreement, and Contractor shall at its expense cooperate with Purchasing Entity in seeking injunctive or other equitable relief in the name of Purchasing Entity or Contractor against any such person. Except as required by applicable law or directed by Purchasing Entity, Contractor will not at any time during or after the term of this Master Agreement disclose, directly or indirectly, any Confidential Information to any person, except in accordance with this Master Agreement, and that upon termination of this Master Agreement or at Purchasing Entity's request, Contractor shall turn over to Purchasing Entity all documents, papers, and other matter in Contractor's possession that embody Confidential Information. Notwithstanding the foregoing, Contractor may keep one copy of such Confidential Information necessary for quality assurance, audits and evidence of the performance of this Master Agreement.

c. Injunctive Relief. Contractor acknowledges that breach of this section, including disclosure of any Confidential Information, will cause irreparable injury to Purchasing Entity that is inadequately compensable in damages. Accordingly, Purchasing Entity may seek and obtain injunctive relief against the breach or threatened breach of the foregoing undertakings, in addition to any other legal remedies that may be available. Contractor acknowledges and agrees that the covenants contained herein are necessary for the protection of the legitimate business interests of Purchasing Entity and are reasonable in scope and content.

9. Right to Publish: Throughout the duration of this Master Agreement, Contractor must secure prior approval from the Lead State or Participating Entity for the release of any information that pertains to the potential work or activities covered by the Master Agreement, including but not limited to reference to or use of the Lead State or a Participating Entity's name, Great Seal of the State, Coat of Arms, any Agency or other subunits of the State government, or any State official or employee, for commercial promotion which is strictly prohibited. News releases or release of broadcast e-mails pertaining to this Master Agreement or Participating Addendum shall not be made without prior written approval of the Lead State or a Participating Entity.

The Contractor shall not make any representations of NASPO ValuePoint's opinion or position as to the quality or effectiveness of the services that are the subject of this Master Agreement without prior written consent. Failure to adhere to this requirement may result in termination of the Master Agreement for cause.

10. Defaults and Remedies

a. The occurrence of any of the following events shall be an event of default under this Master Agreement:

- (1) Nonperformance of contractual requirements; or
- (2) A material breach of any term or condition of this Master Agreement; or
- (3) Any certification, representation or warranty by Contractor in response to the solicitation or in this Master Agreement that proves to be untrue or materially misleading; or

b. Upon the occurrence of an event of default, the party claiming default shall issue a written notice of default, identifying the nature of the default, and providing a period of at least 30 calendar days in which the party in default shall have an opportunity to cure the default. The Lead State shall not be required to provide advance written notice or a cure period and may immediately terminate for convenience this Master Agreement if the Lead State, in its sole discretion, determines that it is reasonably necessary to preserve public safety or prevent immediate public crisis.

Time allowed for cure shall not diminish or eliminate Contractor's liability for damages.

c. If the party in default is afforded an opportunity to cure and fails to cure the default within the period specified in the written notice of default, the party in default shall be in breach of its obligations under this Master Agreement and the party not in default shall have the right to exercise any or all of the following remedies:

- (1) Exercise any remedy provided by law; and
- (2) Terminate this Master Agreement; and
- (3) Suspend Contractor from being able to respond to future NASPO bid solicitations; and
- (4) Suspend Contractor's performance; and

(5) Withhold payment for the portion of the nonconforming Service at issue until the default is remedied.

d. Unless otherwise specified in the Participating Addendum, in the event of a default under a Participating Addendum, a Participating Entity shall provide a written notice of default as described in

this section and have all of the rights and remedies under this paragraph regarding its participation in the Master Agreement, in addition to those set forth in its Participating Addendum.

11. Changes in Contractor Representation: The Contractor must notify the Lead State of changes in the Contractor's key administrative personnel, in writing within 10 calendar days of the change. The Lead State reserves the right to approve changes in key personnel, as identified in the Contractor's proposal. The Contractor agrees to propose replacement key personnel having substantially equal or better education, training, and experience as was possessed by the key person proposed and evaluated in the Contractor's proposal.

12. Force Majeure: Neither party shall be in default by reason of any failure in performance of this Contract outside its reasonable control and without fault or negligence on their part. Such causes may include, but are not restricted to, acts of nature or the public enemy, acts of the government in either its sovereign or contractual capacity, terrorism, fires, floods, epidemics, quarantine restrictions, strikes, freight embargoes and unusually severe weather, but in every case the failure to perform such must be beyond the reasonable control and without the fault or negligence of the party.

13. Indemnification

a. The Contractor shall defend claims or causes or action asserted against NASPO, NASPO ValuePoint, the Lead State, Participating Entities, and Purchasing Entities, along with their officers, agents, and employees as well as any person or entity for which they may be liable (each an "Indemnified Party") and indemnify and hold harmless the indemnified party, from and against any resultant damages agreed to in settlement by Contractor or awarded by a court of final adjudication, reasonable attorneys' fees and related costs for any death, injury, or damage to tangible property to the extent caused by negligence whether act(s), error(s), or omission(s) of the Contractor, its employees or subcontractors or volunteers, at any tier during the performance of its obligations under the Master Agreement. Contractor's duties under this provision are dependent on the indemnified party giving Contractor (1) prompt written notice of such third party claim and (2) sole authority to defend or settle the claim.

b. Indemnification – Intellectual Property. The Contractor shall defend, indemnify and hold harmless NASPO, NASPO ValuePoint, the Lead State, Participating Entities, Purchasing Entities, along with their officers, agents, and employees as well as any person or entity for which they may be liable ("Indemnified Party"), from and against third party claims, or causes of action and pay any damages awarded by a court of final adjudication, reasonable attorneys' fees and related costs arising out of the claim that

the Service or its proper or reasonably expected or acceptable use, infringes that third party's patent, copyright, or trademark or makes unlawful use of its trade secret (an "Intellectual Property Claim" claim).

(1) The Contractor's obligations under this section shall not extend to any claims based on:

- (a) Contractor's compliance with Participating Entity/Purchasing Entity's designs, specifications or instructions or
- (b) Contractor's use of technical information or technology provided by the Participating Entity/Purchasing Entity; or
- (c) non-Contractor software, modifications a Participating Entity/Purchasing Entity makes to, or any specifications or materials a Participating Entity/Purchasing Entity provides or makes available for a Service; or
- (d) Participating Entity/Purchasing Entity's combination of the Service with a non- Contractor product, data or business process; or damages based on the use of a non- Contractor product, data or business process; or
- (e) Participating Entity/Purchasing Entity's use of either Contractor's trademark. (2)

(2) The Indemnified Party shall notify the Contractor promptly after receiving notice of an Intellectual Property Claim. Even if the Indemnified Party fails to provide reasonable notice, the Contractor shall not be relieved from its obligations unless the Contractor can demonstrate that it was prejudiced in defending the Intellectual Property Claim resulting in increased expenses or loss to the Contractor and then only to the extent of the prejudice or expenses. If the Contractor promptly and reasonably investigates and defends any Intellectual Property Claim, it shall have control over the defense and settlement of it. However, the Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Contractor's reasonable request and expense, information and assistance necessary for such defense. If Contractor reasonably believes that a Service may infringe or misappropriate a third party's intellectual property rights, Contractor will seek to (i) procure for Participating Entity/Purchasing Entity the rights to continue to use the Service or (ii) modify or replace it with a functional equivalent to make it non-infringing and notify Participating Entity/Purchasing Entity to discontinue use of the prior version, which Participating Entity/Purchasing Entity must do immediately. If the foregoing options are not commercially reasonable for Contractor, or if required by a valid judicial or government order, Contractor may terminate Participating Entity/Purchasing Entity's license or access rights in the Service. The foregoing constitutes

indemnified party's sole remedy and Contractor's sole and exclusive liability for all IP claims.

14. Independent Contractor: The Contractor shall be an independent contractor. Contractor shall have no authorization, express or implied, to bind the Lead State, Participating States, other Participating Entities, or Purchasing Entities to any agreements, settlements, liability or understanding whatsoever, and agrees not to hold itself out as agent except as expressly set forth herein or as expressly agreed in any Participating Addendum.

15. Individual Customers: Except to the extent modified by a Participating Addendum, each Purchasing Entity shall follow the terms and conditions of the Master Agreement and applicable Participating Addendum and will have the same rights and responsibilities for their purchases as the Lead State has in the Master Agreement, including but not limited to, any indemnity or right to recover any costs as such right is defined in the Master Agreement and applicable Participating Addendum for their purchases. Each Purchasing Entity will be responsible for its own charges, fees, and liabilities. The Contractor will apply the charges and invoice each Purchasing Entity individually.

16. Insurance

a. Unless otherwise agreed in a Participating Addendum, Contractor shall, during the term of this Master Agreement, maintain in full force and effect, the insurance described in this section. Contractor shall acquire such insurance from an insurance carrier or carriers licensed to conduct business in each Participating Entity's state and having a rating of A-, Class VII or better, in the most recently published edition of Best's Reports. Failure to buy and maintain the required insurance may result in this Master Agreement's termination or, at a Participating Entity's option, result in termination of its Participating Addendum.

b. Coverage shall be written on an occurrence basis. The minimum acceptable limits shall be as indicated below:

(1) Commercial General Liability covering premises operations, independent contractors, products and completed operations, blanket contractual liability, personal injury (including death), advertising liability, and property damage, with a limit of not less than \$1 million per occurrence/\$3 million general aggregate;

(2) CLOUD MINIMUM INSURANCE COVERAGE:

Level of Risk	Data Breach and Privacy/Cyber Liability including Technology Errors and Omissions
	Minimum Insurance Coverage
Low Risk Data	\$2,000,000
Moderate Risk Data	\$5,000,000
High Risk Data	\$10,000,000

(3) Contractor must comply with any applicable State Workers Compensation or Employers Liability Insurance requirements.

(4) Professional Liability. As applicable, Professional Liability Insurance Policy in the minimum amount of \$1,000,000 per occurrence and \$1,000,000 in the aggregate, written on a claims made form that provides coverage for its work undertaken pursuant to each Participating Addendum.

c. Contractor shall pay premiums on all insurance policies. Such policies shall also reference this Master Agreement and shall have a condition that they not be revoked by the insurer until thirty (30) calendar days after notice of intended revocation thereof shall have been given to Purchasing Entity and Participating Entity by the Contractor.

d. Prior to commencement of performance, Contractor shall provide to the Lead State a written endorsement to the Contractor's general liability insurance policy or other documentary evidence acceptable however such acceptance shall not be unreasonably withheld to the Lead State that (1) names the Participating States identified in the Request for Proposal as additional insureds to the extent of the liabilities assumed by Contractor as set forth in the indemnification provision of this Agreement, (2) provides that no material alteration, cancellation, non-renewal, or expiration of the coverage contained in such policy shall have effect unless the named Participating State has been given at least thirty (30) days prior written notice, and (3) provides that the Contractor's liability insurance policy shall be primary to the extent of the liabilities assumed by Contractor as set forth in the indemnification provision of this Agreement, with any liability insurance of any Participating State as secondary and noncontributory. Unless otherwise agreed in any Participating Addendum, the Participating Entity's rights and Contractor's obligations are the same as those specified in the first sentence of this subsection. Before performance of any Purchase Order issued after execution of a Participating Addendum authorizing it, the Contractor shall provide to a Purchasing Entity or Participating Entity who requests it the same information described in this subsection.

e. Contractor shall furnish to the Lead State, Participating Entity, and, on request, the Purchasing Entity copies of certificates of all required insurance within thirty (30) calendar days of the execution of this Master Agreement, the execution of a Participating Addendum, or the Purchase Order's effective date and prior to performing any work. The insurance certificate shall provide the following information: the name and address of the insured; name, address, telephone number and signature of the authorized agent; name of the insurance company (authorized to operate in all states); policy period, policy number, limits of liability; and an acknowledgment of the requirement for notice of cancellation by Contractor. Copies of renewal certificates of all required insurance shall be furnished within thirty (30) days after any renewal date. These certificates of insurance must expressly indicate compliance with each and every insurance requirement specified in this section. Failure to provide evidence of coverage may, at sole option of the Lead State, or any Participating Entity, result in this Master Agreement's termination or the termination of any Participating Addendum.

f. Coverage and limits shall not limit Contractor's liability and obligations under this Master Agreement, any Participating Addendum, or any Purchase Order.

17. Laws and Regulations: Contractor shall comply fully with all Federal and State laws and regulations applicable to the operation of its business in the provision of the Services.

18. No Waiver of Sovereign Immunity: In no event shall this Master Agreement, any Participating Addendum or any contract or any Purchase Order issued thereunder, or any act of a Lead State, a Participating Entity, or a Purchasing Entity be a waiver of any form of defense or immunity, whether sovereign immunity, governmental immunity, immunity based on the Eleventh Amendment to the Constitution of the United States or otherwise, from any claim or from the jurisdiction of any court.

This section applies to a claim brought against the Participating State only to the extent Congress has appropriately abrogated the Participating State's sovereign immunity and is not consent by the Participating State to be sued in federal court. This section is also not a waiver by the Participating State of any form of immunity, including but not limited to sovereign immunity and immunity based on the Eleventh Amendment to the Constitution of the United States.

19. Ordering

a. Master Agreement order and purchase order numbers shall be clearly shown on all acknowledgments, shipping labels, packing slips, invoices, and on all correspondence.

b. This Master Agreement permits Purchasing Entities to define project-specific requirements and informally compete the requirement among other firms having a Master Agreement on an "as needed" basis. This procedure may also be used when requirements are aggregated or other firm commitments may be made to achieve reductions in pricing. This procedure may be modified in Participating Addenda and adapted to Purchasing Entity rules and policies. The Purchasing Entity may in its sole discretion determine which firms should be solicited for a quote. The Purchasing Entity may select the quote that it considers most advantageous, cost and other factors considered.

c. Each Purchasing Entity will identify and utilize its own appropriate purchasing procedure and documentation. Contractor is expected to become familiar with the Purchasing Entities' rules, policies, and procedures regarding the ordering of supplies and/or services contemplated by this Master Agreement.

d. Contractor shall not begin providing Services without a valid Service Level Agreement or other appropriate commitment document compliant with the law of the Purchasing Entity.

e. Orders may be placed consistent with the terms of this Master Agreement during the term of the Master Agreement.

f. All Orders pursuant to this Master Agreement, at a minimum, shall include:

- (1) The services or supplies being delivered;
- (2) The place and requested time of delivery;
- (3) A billing address;
- (4) The name, phone number, and address of the Purchasing Entity representative;
- (5) The price per unit or other pricing elements consistent with this Master Agreement and the contractor's proposal;
- (6) A ceiling amount of the order for services being ordered; and
- (7) The Master Agreement identifier and the Participating State contract identifier.

g. All communications concerning administration of Orders placed shall be furnished solely to the authorized purchasing agent within the Purchasing Entity's purchasing office, or to such other individual identified in writing in the Order.

h. Orders must be placed pursuant to this Master Agreement prior to the termination date of this Master Agreement. Contractor is reminded that financial obligations of Purchasing Entities payable after the current applicable fiscal year are contingent upon agency funds for that purpose being appropriated, budgeted, and otherwise made available.

i. Notwithstanding the expiration or termination of this Master Agreement, Contractor agrees to perform in accordance with the terms of any Orders then outstanding at the time of such expiration or termination. Contractor shall not honor any Orders placed after the expiration or termination of this Master Agreement. Orders from any separate indefinite quantity, task orders, or other form of indefinite delivery order arrangement priced against this Master Agreement may not be placed after the expiration or termination of this Master Agreement, notwithstanding the term of any such indefinite delivery order agreement.

20. Participants and Scope

a. Contractor may not deliver Services under this Master Agreement until a Participating Addendum acceptable to the Participating Entity and Contractor is executed. The NASPO ValuePoint Master Agreement Terms and Conditions are applicable to any Order by a Participating Entity (and other Purchasing Entities covered by their Participating Addendum), except to the extent altered, modified, supplemented or amended by a Participating Addendum. By way of illustration and not limitation, this authority may apply to unique delivery and invoicing requirements, confidentiality requirements, defaults on Orders, governing law and venue relating to Orders by a Participating Entity, indemnification, and insurance requirements. Statutory or constitutional requirements relating to availability of funds may require specific language in some Participating Addenda in order to comply with applicable law. The expectation is that these alterations, modifications, supplements, or amendments will be addressed in the Participating Addendum or, with the consent of the Purchasing Entity and Contractor, may be included in the ordering document (e.g. purchase order or contract)

used by the Purchasing Entity to place the Order.

b. Subject to subsection 20c and a Participating Entity's Participating Addendum, the use of specific NASPO ValuePoint cooperative Master Agreements by state agencies, political subdivisions and other Participating Entities (including cooperatives) authorized by individual state's statutes to use state contracts is subject to the approval of the respective State Chief Procurement Official.

c. Unless otherwise stipulated in a Participating Entity's Participating Addendum, specific services accessed through the NASPO ValuePoint cooperative Master Agreements for Cloud Services by state executive branch agencies, as required by a Participating Entity's statutes, are subject to the authority and approval of the Participating Entity's Chief Information Officer's Office¹.

d. Obligations under this Master Agreement are limited to those Participating Entities who have signed a Participating Addendum and Purchasing Entities within the scope of those Participating Addenda. Financial obligations of Participating States are limited to the orders placed by the departments or other state agencies and institutions having available funds. Participating States incur no financial obligations on behalf of political subdivisions.

e. NASPO ValuePoint is not a party to the Master Agreement. It is a nonprofit cooperative purchasing organization assisting states in administering the NASPO ValuePoint cooperative purchasing program for state government departments, institutions, agencies and political subdivisions (e.g., colleges, school districts, counties, cities, etc.) for all 50 states, the District of Columbia and the territories of the United States.

f. Participating Addenda shall not be construed to amend the terms of this Master Agreement between the Lead State and Contractor.

g. Participating Entities who are not states may under some circumstances sign their own Participating Addendum, subject to the approval of participation by the Chief Procurement Official of the state where the Participating Entity is located. Coordinate requests for such participation through NASPO ValuePoint. Any permission to participate through execution of a Participating Addendum is not a determination that procurement authority exists in the Participating Entity; they must ensure that they have the requisite procurement authority to execute a Participating Addendum.

h. Resale. Subject to any explicit permission in a Participating Addendum, Purchasing Entities may not resell goods, software, or Services obtained under this Master Agreement. This limitation does not prohibit: payments by employees of a Purchasing Entity as explicitly permitted under this agreement; sales of goods to the general public as surplus property; and fees associated with inventory transactions with other

¹ Chief Information Officer means the individual designated by the Governor with Executive Branch, enterprise-wide responsibility for the leadership and management of information technology resources of a state.

governmental or nonprofit entities under cooperative agreements and consistent with a Purchasing Entity's laws and regulations. Any sale or transfer permitted by this subsection must be consistent with license rights granted for use of intellectual property.

21. Payment: Unless otherwise stipulated in the Participating Addendum, Payment is normally made within 30 days following the date of a correct invoice is received. Purchasing Entities reserve the right to withhold payment of a portion (including all if applicable) of disputed amount of an invoice. After 45 days the Contractor may assess overdue account charges up to a maximum rate of one percent per month on the outstanding balance. Payments will be remitted by mail. Payments may be made via a State or political subdivision "Purchasing Card" with no additional charge.

22. Data Access Controls: Unless otherwise provided for in the service provider terms, Contractor will provide access to Purchasing Entity's Data only to those Contractor employees, contractors and subcontractors ("Contractor Staff") who need to access the Data to fulfill Contractor's obligations under this Agreement. Contractor shall not access a Purchasing Entity's user accounts or Data, except on the course of data center operations, response to service or technical issues, as required by the express terms of this Master Agreement, or at a Purchasing Entity's written request.

Contractor may not share a Purchasing Entity's Data with its parent corporation, other affiliates, or any other third party without the Purchasing Entity's express written consent.

Contractor will ensure that, prior to being granted access to the Data, Contractor Staff who perform work under this Agreement have successfully completed annual instruction of a nature sufficient to enable them to effectively comply with all Data protection provisions of this Agreement; and possess all qualifications appropriate to the nature of the employees' duties and the sensitivity of the Data they will be handling.

23. Operations Management: Unless otherwise provided for in the service provider terms, the services

Provider shall maintain and follow appropriate technical and organization measures intended to protect customer Data against accidental, unauthorized or unlawful access, disclosure, alteration, loss or destruction.

24. Public Information: This Master Agreement and all related documents are subject to disclosure pursuant to the Purchasing Entity's public information laws.

25. Purchasing Entity Data: Purchasing Entity retains full right and title to Data provided by it. Contractor shall not collect, access, or use user-specific Purchasing Entity Data except as strictly necessary to provide Service to the Purchasing Entity. No information regarding Purchasing Entity's use of the Service may be disclosed, provided, rented or sold to any third party for any reason unless required by law or regulation or by an order of a court of competent jurisdiction. The obligation shall extend beyond the term of this Master Agreement in perpetuity.

Unless otherwise provided for in the service provider terms, Contractor shall not use any

information collected in connection with this Master Agreement, including Purchasing Entity Data, for any purpose other than fulfilling its obligations under this Master Agreement.

26. Records Administration and Audit.

a. The Contractor shall maintain books, records, documents, and other evidence pertaining to this Master Agreement and orders placed by Purchasing Entities under it to the extent and in such detail as shall adequately reflect performance and administration of payments and fees. Contractor shall permit no more frequently than once in any twelve month period the Lead State, a Participating Entity, a Purchasing Entity, the federal government (including its grant awarding entities and the U.S. Comptroller General), and any other duly authorized agent of a governmental agency, to audit, inspect, examine, copy and/or transcribe Contractor's books, documents, papers and records directly pertinent to this Master Agreement or orders placed by a Purchasing Entity under it for the purpose of making audits, examinations, excerpts, and transcriptions. Such audit request shall be made with at least thirty days in advance notice to Contractor. Each audit participant shall comply with Contractor confidentiality and site security policies, procedures and practices. This right shall survive for a period of six (6) years following termination of this Agreement or final payment for any order placed by a Purchasing Entity against this Agreement, whichever is later, to assure compliance with the terms hereof or to evaluate performance hereunder.

b. Without limiting any other remedy available to any governmental entity, the Contractor shall reimburse the applicable Lead State, Participating Entity, or Purchasing Entity for any overpayments inconsistent with the terms of the Master Agreement or orders or underpayment of fees found as a result of the examination of the Contractor's records.

c. The rights and obligations herein exist in addition to any quality assurance obligation in the Master Agreement requiring the Contractor to self-audit contract obligations and that permits the Lead State to review compliance with those obligations.

27. Administrative Fees: The Contractor shall pay to NASPO ValuePoint, or its assignee, a NASPO ValuePoint Administrative Fee of one-quarter of one percent (0.25% or 0.0025) no later than 60 days following the end of each calendar quarter. The NASPO ValuePoint Administrative Fee shall be submitted quarterly and is based on sales of the Services. The NASPO ValuePoint Administrative Fee is not negotiable. This fee is to be included as part of the pricing submitted with proposal.

Additionally, some states may require an additional administrative fee be paid directly to the state on purchases made by Purchasing Entities within that state. For all such requests, the fee level, payment method and schedule for such reports and payments will be incorporated into the Participating Addendum that is made a part of the Master Agreement. The Contractor may adjust the Master Agreement pricing accordingly for purchases made by Purchasing Entities within the jurisdiction of the state. All such

agreements shall not affect the NASPO ValuePoint Administrative Fee percentage or the prices paid by the Purchasing Entities outside the jurisdiction of the state requesting the additional fee. The NASPO ValuePoint Administrative Fee shall be based on the gross amount of all sales at the adjusted prices (if any) in Participating Addenda.

28. [Reserved]

29. Title to License: If access to the Product requires an application program interface (API), Contractor shall convey to Purchasing Entity an irrevocable and perpetual license to use the API during the term of the service to use the Services.

30. Data Privacy: The service provider must comply with all applicable laws related to data privacy and security, including IRS Pub 1075 that are applicable to it as an IT Service Provider. With respect to private and hybrid cloud services, prior to entering into a SLA with a Purchasing Entity, the Contractor and Purchasing Entity must cooperate and hold a meeting to determine the Data Categorization to determine whether the Contractor will hold, store, or process High Risk Data, Moderate Risk Data and Low Risk Data. The Contractor must document the Data Categorization in the SLA or Statement of Work.

31. Warranty: At a minimum the Contractor must warrant the following:

- a. Contractor has acquired any and all rights, grants, assignments, conveyances, licenses, permissions, and authorization for the Contractor to provide the Services described in this Master Agreement.
- b. Contractor will perform materially as described in this Master Agreement, SLA, Statement of Work, including any performance representations contained in the Contractor's response to the Solicitation by the Lead State.
- c. Unless otherwise provided for in the service terms, the Contractor will not interfere with a Purchasing Entity's access to and use of the Services it acquires from this Master Agreement.
- d. The Services provided by the Contractor are compatible with and will operate successfully with the environment (including web browser and operating system) specified either in the Solicitation response or in the service terms.
- e. The Contractor must use industry-leading technology to detect and remove worms, Trojans, rootkits, rogues, dialers, spyware, etc.

32. Transition Assistance:

- a. For private and hybrid cloud services, the Contractor shall reasonably cooperate with other parties in connection with all Services to be delivered under this Master Agreement, including without limitation any successor service provider to whom a Purchasing Entity's Data is transferred in connection with the termination or expiration of this Master Agreement. The Contractor shall assist a Purchasing Entity in exporting

and extracting a Purchasing Entity's Data, in a format available from the cloud provider and as agreed by a Purchasing Entity, at no additional cost to the Purchasing Entity for a period not to exceed 60 days from the date of termination or expiration. Any transition services requested by a Purchasing Entity involving additional knowledge transfer and support may be subject to a separate transition Statement of Work.

b. A Purchasing Entity and the Contractor shall, when reasonable, create a Transition Plan Document identifying the transition services to be provided and including a Statement of Work if applicable.

c. The Contractor must maintain the confidentiality and security of a Purchasing Entity's Data during the transition services and thereafter as required by the Purchasing Entity.

33. Waiver of Breach: Failure of the Lead State, Participating Entity, Purchasing Entity or Contractor to declare a default or enforce any rights and remedies shall not operate as a waiver under this Master Agreement or Participating Addendum. Any waiver by the Lead State, Participating Entity, Purchasing Entity or Contractor must be in writing. Waiver by the Lead State Participating Entity or Contractor of any default, right or remedy under this Master Agreement or Participating Addendum, or by Purchasing Entity with respect to any Purchase Order, or breach of any terms or requirements of this Master Agreement, a Participating Addendum, or Purchase Order shall not be construed or operate as a waiver of any subsequent default or breach of such term or requirement, or of any other term or requirement under this Master Agreement, Participating Addendum, or Purchase Order.

34. Assignment of Antitrust Rights: Contractor irrevocably assigns to a Participating Entity who is a state any claim for relief or cause of action which the Contractor now has or which may accrue to the Contractor in the future by reason of any violation of state or federal antitrust laws (15 U.S.C. § 1-15 or a Participating Entity's state antitrust provisions), as now in effect and as may be amended from time to time, in connection with any goods or services provided to the Contractor for the purpose of carrying out the Contractor's obligations under this Master Agreement or Participating Addendum, including, at a Participating Entity's option, the right to control any such litigation on such claim for relief or cause of action.

35. Debarment : The Contractor certifies, to the best of its knowledge, that neither it nor its principals are presently debarred, suspended, proposed for debarment, declared ineligible, or voluntarily excluded from participation in this transaction (contract) by any governmental department or agency. This certification represents a recurring certification made at the time any Order is placed under this Master Agreement. If the Contractor cannot certify this statement, attach a written explanation for review by the Lead State.

36. Performance and Payment Time Frames that Exceed Contract Duration: All maintenance or other agreements for services entered into during the duration of an SLA and whose performance and payment time frames extend beyond the duration of this Master Agreement shall remain in effect for performance and payment purposes

(limited to the time frame and services established per each written agreement). No new leases, maintenance or other agreements for services may be executed after the Master Agreement has expired. For the purposes of this section, renewals of maintenance, subscriptions, SaaS subscriptions and agreements, and other service agreements, shall not be considered as “new.”

37. Governing Law and Venue

a. The procurement, evaluation, and award of the Master Agreement shall be governed by and construed in accordance with the laws of the Lead State sponsoring and administering the procurement. The construction and effect of the Master Agreement after award shall be governed by the law of the state serving as Lead State (in most cases also the Lead State). The construction and effect of any Participating Addendum or Order against the Master Agreement shall be governed by and construed in accordance with the laws of the Participating Entity's or Purchasing Entity's State.

b. Unless otherwise specified in the RFP, the venue for any protest, claim, dispute or action relating to the procurement, evaluation, and award is in the Lead State. Venue for any claim, dispute or action concerning the terms of the Master Agreement shall be in the state serving as Lead State. Venue for any claim, dispute, or action concerning any Order placed against the Master Agreement or the effect of a Participating Addendum shall be in the Purchasing Entity's State.

c. If a claim is brought in a federal forum, then it must be brought and adjudicated solely and exclusively within the United States District Court for (in decreasing order of priority): the Lead State for claims relating to the procurement, evaluation, award, or contract performance or administration if the Lead State is a party; the Participating State if a named party; the Participating Entity state if a named party; or the Purchasing Entity state if a named party.

d. This section is also not a waiver by the Participating State of any form of immunity, including but not limited to sovereign immunity and immunity based on the Eleventh Amendment to the Constitution of the United States.

38. No Guarantee of Service Volumes: The Contractor acknowledges and agrees that the Lead State and NASPO ValuePoint makes no representation, warranty or condition as to the nature, timing, quality, quantity or volume of business for the Services or any other products and services that the Contractor may realize from this Master Agreement, or the compensation that may be earned by the Contractor by offering the Services.

39. NASPO ValuePoint eMarket Center: In July 2011, NASPO ValuePoint entered into a multi-year agreement with SciQuest, Inc. whereby SciQuest will provide certain electronic catalog hosting and management services to enable eligible NASPO ValuePoint's customers to access a central online website to view and/or shop the goods and services available from existing NASPO ValuePoint Cooperative Contracts. The central online website is referred to as the NASPO ValuePoint eMarket Center.

The Contractor will have visibility in the eMarket Center through Ordering Instructions. These Ordering Instructions are available at no cost to the Contractor and provided customers information regarding the Contractors website and ordering information.

At a minimum, the Contractor agrees to the following timeline: NASPO ValuePoint eMarket Center Site Admin shall provide a written request to the Contractor to begin Ordering Instruction process. The Contractor shall have thirty (30) days from receipt of written request to work with NASPO ValuePoint to provide any unique information and ordering instructions that the Contractor would like the customer to have.

40. Contract Provisions for Orders Utilizing Federal Funds: Pursuant to Appendix II to 2 Code of Federal Regulations (CFR) Part 200, Contract Provisions for Non-Federal Entity Contracts Under Federal Awards, Orders funded with federal funds may have additional contractual requirements or certifications that must be satisfied at the time the Order is placed or upon delivery. These federal requirements may be proposed by Participating Entities in Participating Addenda and Purchasing Entities for incorporation in Orders placed under this master agreement.

41. Government Support: No, materials, special access, or other obligations on behalf of the states or other Participating Entities, other than payment, are required under the Master Agreement.

42. NASPO ValuePoint Summary and Detailed Usage Reports: In addition to other reports that may be required by this solicitation, the Contractor shall provide the following NASPO ValuePoint reports.

a. Summary Sales Data. The Contractor shall submit quarterly sales reports directly to NASPO ValuePoint using the NASPO ValuePoint Quarterly Sales/Administrative Fee Reporting Tool found at <http://www.naspo.org/WNCPO/Calculator.aspx>. Any/all sales made under the contract shall be reported as cumulative totals by state. Even if Contractor experiences zero sales during a calendar quarter, a report is still required. Reports shall be due no later than 30 day following the end of the calendar quarter (as specified in the reporting tool).

b. Detailed Sales Data. Contractor shall also report detailed sales data by: (1) state; (2) entity/customer type, e.g. local government, higher education, K12, non-profit; (3) Purchasing Entity name; (4) Purchasing Entity bill-to and ship-to locations; (4) Purchasing Entity and Contractor Purchase Order identifier/number(s); (5) Purchase Order Type (e.g. sales order, credit, return, upgrade, determined by industry practices); (6) Purchase Order date; (7) and line item description, including product number if used. The report shall be submitted in any form required by the solicitation. Reports are due on a quarterly basis and must be received by the Lead State and NASPO ValuePoint Cooperative Development Team no later than thirty (30) days after the end of the reporting period. Reports shall be delivered to the Lead State and to the NASPO ValuePoint Cooperative Development Team electronically through a designated portal, email, CD-Rom, flash drive or other method as determined by the Lead State and

NASPO ValuePoint. Detailed sales data reports shall include sales information for all sales under Participating Addenda executed under this Master Agreement. The format for the detailed sales data report is shown in Attachment F.

c. Reportable sales for the summary sales data report and detailed sales data report includes sales to employees for personal use where authorized by the solicitation and the Participating Addendum. Report data for employees should be limited to ONLY the state and entity they are participating under the authority of (state and agency, city, county, school district, etc.) and the amount of sales. No personal identification numbers, e.g. names, addresses, social security numbers or any other numerical identifier, may be submitted with any report.

d. Contractor shall provide the NASPO ValuePoint Cooperative Development Coordinator with an executive summary each quarter that includes, at a minimum, a list of states with an active Participating Addendum, states that Contractor is in negotiations with and any PA roll out or implementation activities and issues. NASPO ValuePoint Cooperative Development Coordinator and Contractor will determine the format and content of the executive summary. The executive summary is due 30 days after the conclusion of each calendar quarter.

e. Timely submission of these reports is a material requirement of the Master Agreement. The recipient of the reports shall have exclusive ownership of the media containing the reports. The Lead State and NASPO ValuePoint shall have a perpetual, irrevocable, non-exclusive, royalty free, transferable right to display, modify, copy, and otherwise use reports, data and information provided under this section.

f. If requested by a Participating Entity, the Contractor must provide detailed sales data within the Participating State.

43. Entire Agreement: This Master Agreement, along with any attachment, contains the entire understanding of the parties hereto with respect to the Master Agreement unless a term is modified in a Participating Addendum with a Participating Entity.

44. Limitation of Liability: Except as otherwise set forth in the Indemnification Paragraphs above, the limit of liability shall be as follows:

a. Contractor's liability for any claim, loss or liability arising out of, or connected with the Services provided, and whether based upon default or other liability such as breach of contract, warranty Negligence, misrepresentation or otherwise, shall in no case exceed direct damages in: (i) an amount equal to two (2) times the charges specified in the purchase order for the Services, or parts thereof forming the basis of the Purchasing Entity's claim, (said amount not to exceed a total of twelve (12) months charges payable under the applicable Purchase order) or (ii) five million dollars (\$5, 000,000), whichever is greater.

b. The Purchasing Entity may retain such monies from any amount due Contractor as may be necessary to satisfy any claim for damages, costs and the like asserted against the Purchasing Entity unless Contractor at the time of the presentation of claim shall demonstrate to the Purchasing Entity's satisfaction that sufficient monies are set aside by the Contractor in the form of a bond or through insurance coverage to cover associated damages and other costs.

c. Notwithstanding the above, neither the Contractor nor the Purchasing Entity shall be liable for any consequential, indirect or special damages of any kind which may result directly or indirectly from such performance, including, without limitation, damages resulting from loss of use or loss of profit by the Purchasing Entity, the Contractor, or by others.

The limitation of liability in Section 43 will not apply to claims for bodily injury or death.

ATTACHMENT B – IDENTIFICATION OF SERVICE MODELS MATRIX

Offerors must complete the following form to identify the service models your firm offers under this RFP. You may provide a list of the different SaaS, IaaS, and/or PaaS services that you offer, including the Categorization of Risk that you have the ability to store and secure. This document is to provide purchasing entities and eligible users a quick snap shot of the cloud solutions your firm provides.

Unisys is pleased to offer all four cloud deployment models: public cloud, private cloud, community cloud, and hybrid cloud. Under these models, we are also pleased to offer all three cloud service delivery methods: IaaS, PaaS, and SaaS.

Attainment of certifications is a lengthy ongoing process for vendors. Additionally, the certification standards are constantly evolving. Our partners in this proposal are at a FISMA Moderate level. Therefore, across cloud deployment models and delivery options, we are positioned to secure and store FISMA Low and Moderate risk data types. For FISMA High, Unisys Stealth for data in motion, when combined with encryption technologies for data at rest, can meet the requirements of FISMA High risk data. For the private cloud model, we can design a cloud environment based on the required standards and controls defined by the client.

Service Model	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered
SaaS	Unisys AWS Microsoft Google	Unisys AWS Microsoft Google	See notes above	All
IaaS	Unisys AWS Microsoft Google	Unisys AWS Microsoft Google	See notes above	All
PaaS	Unisys AWS Microsoft Google	Unisys AWS Microsoft Google	See notes above	All

State of Utah NASPO Cloud Solutions Cost Proposal

Solicitation No: CH16012



UNISYS

ATTACHMENT G – COST SCHEDULE

Solicitation Number CH16012 NASPO ValuePoint Cloud Solutions RFP

Cloud Solutions By Category. Specify **Discount Percent %** Offered for products in each category. Highest discount will apply for products referenced in detail listings for multiple categories. Provide a detailed product offering for each category.

1. Unisys Partner: Amazon Web Services (AWS) Public and Community IaaS, PaaS, and SaaS

The AWS offerings below will be provided at a 1.5% discount off the AWS list price. A 15% maintenance support fee is also required. The support fees will be calculated based on consumption. This offering will include:

- Any of the product, solution sets below
- Business Level Support to the buyer from AWS
- Unisys provided monthly billing and reporting as required by the Master Agreement (MA), Participating Agreement (PA), and the requirements as stated in this RFP.
- Unisys provided customer enablement, contract management support, fee management, and fee disbursements as required by the MA, PA, and requirements as stated in this RFP.

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
	Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable computing capacity—literally, servers in Amazon’s data centers—that you use to build and host your software systems.	Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks. Database engines available through Amazon RDS include Amazon Aurora, MySQL, Oracle, Microsoft SQL Server, and PostgreSQL.	AWS Identity and Access Management (IAM) is a web service that enables AWS customers to manage users and user permissions in AWS. The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console. With AWS IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
	Unisys Stealth is software-defined security that delivers advanced protection for critical workloads in the Amazon Web Services cloud. Stealth is designed to create a secure communications tunnel, where the data is encrypted and split before transmission, and to cloak Stealth end points from users or devices except those who are preidentified as part of a secure community referred to as a Community of Interest (COI). Management of COIs is easy with Stealth because access is defined by identity, not physical topology.	Amazon EC2 Container Service is a highly scalable, high-performance container management service that supports Docker containers and allows you to easily run distributed applications on a managed cluster of Amazon EC2 instances.	users can access. AWS Directory Service is a managed service that allows you to connect your AWS resources with an existing on-premises Microsoft Active Directory or to set up a new, stand-alone directory in the AWS cloud. Connecting to an on-premises directory is easy; once this connection is established, users can access AWS resources and applications with their existing corporate credentials.
	Auto Scaling is a web service designed to launch or terminate Amazon EC2 instances automatically based on user-defined policies, schedules, and health checks.	Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. You can use Amazon DynamoDB to create a database table that can store and retrieve data, and serve many levels of request traffic. Amazon DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent and fast performance.	AWS Service Catalog is a service that allows administrators to create and manage approved catalogs of resources that end users can then access at a personalized portal. You can control which users have access to which applications or AWS resources to enable compliance with your business policies, while users can easily browse and launch products from the catalogs you create.
	Elastic Load Balancing automatically distributes your incoming application traffic	Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse	AWS Config is a fully managed service that provides you with an AWS

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
	across multiple Amazon EC2 instances. It detects unhealthy instances and reroutes traffic to healthy instances until the unhealthy instances are restored. Elastic Load Balancing scales its request handling capacity automatically in response to incoming traffic.	solution that makes it simple and cost-effective to efficiently analyze your data using your existing business intelligence tools.	resource inventory, configuration history, and configuration change notifications to enable security and governance. With AWS Config you can discover existing AWS resources, export a complete inventory of your AWS resources with configuration details, and determine how a resource was configured at a point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.
	Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you defined. This virtual network closely resembles a traditional network that you would operate in your own data center, with the benefits of using the scalable infrastructure of AWS.	Amazon ElastiCache is a web service that makes it easy to set up, manage, and scale distributed, in-memory cache environments in the cloud. It provides a high-performance, resizable, and cost-effective in-memory cache, while removing the complexity associated with deploying and managing a distributed cache environment.	AWS CloudHSM provides you with secure cryptographic key storage by making Hardware Security Modules (HSMs) available in the AWS cloud.
	Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service.	Amazon Elastic MapReduce (Amazon EMR) is a web service that makes it easy to process large amounts of data efficiently. Amazon EMR uses Hadoop processing combined with several AWS products to perform tasks such as web indexing, data mining, log file analysis, machine learning,	AWS Key Management Service (KMS) is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data and uses Hardware Security Modules (HSMs) to protect the security of your keys. AWS KMS is integrated with other AWS cloud services

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
		scientific simulation, and data warehousing.	including Amazon EBS, Amazon S3, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of key usage to help meet your regulatory and compliance needs.
	AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 GB or 10 GB Ethernet fiber optic cable. One end of the cable is connected to your router and the other to an AWS Direct Connect router. With this connection, you can create virtual interfaces directly to the AWS cloud and Amazon VPC, bypassing Internet service providers in your network path.	Amazon Kinesis is a managed service that scales elastically for real-time processing of streaming big data. The service takes in large streams of data records that can then be consumed in real time by multiple data processing applications that can be run on Amazon EC2 instances. The data processing applications use the Amazon Kinesis Client Library and are called Amazon Kinesis applications.	With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS Software Development Kits (SDKs), the command line tools, and higher level AWS cloud services. You can also identify which users and accounts called AWS APIs for services that support AWS CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate AWS CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn AWS CloudTrail logging on and off.
	Amazon Simple Storage Service (Amazon S3) is storage for the Internet. You can use Amazon S3 to store and retrieve data from anywhere on the web. You can accomplish these tasks using the simple and intuitive web interface of the AWS Management Console.	AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premises data sources at specified intervals. With AWS Data Pipeline, you can	Amazon CloudWatch is a web service that enables you to collect, view, and analyze metrics. Amazon CloudWatch lets you programmatically retrieve your monitoring data, view graphs, and set alarms to help you troubleshoot, spot

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
		regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to AWS cloud services such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon EMR.	trends, and take automated action based on the state of your cloud environment.
	Amazon Glacier is a storage service optimized for infrequently used data ("cold data"). The service provides secure, durable, and extremely low-cost storage for data archiving and backup. With Amazon Glacier, you can store your data cost effectively for months, years, or decades. Amazon Glacier enables you to offload the administrative burdens of operating and scaling storage to AWS, so you don't have to worry about capacity planning, hardware provisioning, data replication, hardware failure detection and recovery, or time-consuming hardware migrations.	Amazon Mobile Analytics is a service that lets you easily collect, visualize, and understand application usage data at scale. Many mobile application analytics solutions deliver usage data several hours after the events occur. Amazon Mobile Analytics is designed to deliver usage reports within 60 minutes of receiving data from an application so that you can act on the data more quickly.	AWS provides API-based cloud computing services with multiple interfaces to those services, including SDKs, IDE Toolkits, and command line tools for developing and managing AWS resources.
	Amazon Elastic Block Store (Amazon EBS) provides block-level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are highly available and reliable storage volumes that can be attached to at least one running instance that is in the same Availability Zone. Amazon EBS volumes that are attached to an Amazon EC2 instance are exposed as storage volumes that persist independently from the life of the instance. With Amazon EBS, you only pay for what you use.	With AWS Elastic Beanstalk, you can quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application	Amazon CloudFront is a content delivery web service. It integrates with other AWS cloud services to give developers and governments an easy way to distribute content to end users with low-latency, high data transfer speeds and no commitments.

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
	AWS Import/Export accelerates transferring large amounts of data between the cloud and portable storage devices that you mail to AWS. AWS transfers data directly onto and off your storage devices using Amazon's high-speed internal network. Your data load typically begins the next business day after your storage device arrives at AWS. After the data export or import completes, AWS returns your storage device. For large data sets, AWS Import/Export is significantly faster than Internet transfer and more cost effective than upgrading your connectivity.	health monitoring. AWS CloudFormation offers developers and system administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable way. You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and associated dependencies or runtime parameters, required to run your application.	AWS CodePipeline is a continuous delivery and release automation service that aids smooth deployments. You can design your development workflow for checking in code, building the code, deploying your application to staging, testing it, and releasing it to production. You can integrate third-party tools into the steps of your release process, or you can use AWS CodePipeline as an end-to-end solution.
	AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and the AWS storage infrastructure.	AWS CodeDeploy is a service that automates code deployments to Amazon EC2 instances. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate deployments, eliminating the need for error-prone manual operations, and the service scales with your infrastructure so you can easily deploy applications to one Amazon EC2 instance or thousands.	The Amazon AppStream web service deploys your application on AWS infrastructure and streams input and output between your application and devices such as personal computers, tablets, and mobile phones. Your application's processing occurs in the cloud, so it can scale to handle vast computational loads. Devices need only display output and return user input, so the client application on the device can be lightweight in terms of file size and processing requirements.
	Access and manage Amazon cloud services at a simple and intuitive web-based user	AWS CodeCommit is a secure, highly scalable, managed source control	Amazon CloudSearch is a fully managed service in the cloud that makes it easy to

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
	interface. You can also use the AWS Console mobile app to quickly view resources on-the-go.	service that hosts private Git repositories. AWS CodeCommit eliminates the need for you to operate your own source control system or worry about scaling its infrastructure. You can use AWS CodeCommit to store anything from code to binaries, and it supports the standard functionality of Git, allowing it to work seamlessly with your existing Git-based tools.	set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. With Amazon CloudSearch, you can quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.
	The AWS Command Line Interface (CLI) is a unified tool used to manage your AWS cloud services. With just one tool to download and configure, you can control multiple AWS cloud services from the command line and automate them through scripts.	AWS OpsWorks provides a simple and flexible way to create and manage stacks and applications. With AWS OpsWorks, you can provision AWS resources, manage their configuration, deploy applications to those resources, and monitor their health.	Amazon Simple Workflow Service (Amazon SWF) makes it easy to build applications that coordinate work across distributed components. In Amazon SWF, a task represents a logical unit of work that is performed by a component of your application. Coordinating tasks across the application involves managing intertask dependencies, scheduling, and concurrency in accordance with the logical flow of the application. Amazon SWF gives you full control over implementing tasks and coordinating them without worrying about underlying complexities such as tracking their progress

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
		Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity. Users can comment on files, send them to others for feedback, and upload new versions without having to email multiple versions of their files as attachments.	and maintaining their state. Amazon Elastic Transcoder lets you convert media files that you stored in Amazon S3 to media files in the formats required by consumer playback devices. For example, you can convert large, high-quality digital media files to formats that users can play back on mobile devices, tablets, web browsers, and connected TV sets.
		Amazon WorkSpaces is a fully managed desktop computing service in the cloud. Amazon WorkSpaces allows you to easily provision cloud-based desktops that allow end users to access the documents, applications, and resources you need with the device of your choice, including laptops, iPad, Kindle Fire, or Android tablets. With a few clicks in the AWS Management Console, you can provision a high-quality cloud desktop experience for users at a cost that is highly competitive with traditional desktops and half the cost of most Virtual Desktop Infrastructure (VDI) solutions.	Amazon Cognito is a simple user identity and data synchronization service that helps you securely manage and synchronize application data for your users across their mobile devices. You can create unique identities for your users through a number of public login providers (Amazon, Facebook, and Google) and support unauthenticated guests.
		Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles messages or workflows between other components in a system.	Amazon Flexible Payments Service facilitates the digital transfer of money between two entities – be they humans or computers.

Amazon Web Services Offering Details	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
		Amazon Simple Email Service (Amazon SES) is an outbound-only email-sending service that provides an easy, cost-effective way for you to send email.	AWS Marketplace is an online store that helps you find, buy, and immediately start using the software and services you need to build products and run your business.
		Amazon Simple Notification Service (Amazon SNS) is a web service that enables applications, end users, and devices to instantly send and receive notifications from the cloud.	
		AWS Lambda is a compute service that runs your code in response to events and automatically manages the compute resources for you, making it easy to build applications that respond quickly to new information. AWS Lambda starts running your code within milliseconds of an event such as an image upload, an in-app activity, a website click, or output from a connected device.	
Notes	- This entire set of offerings is consumption driven. - List price costs can be estimated using a live calculator available here: http://calculator.s3.amazonaws.com/index.html - A description of AWS Business level support information is available here: https://aws.amazon.com/premiumsupport/business-support/		

2. Unisys Partner: Microsoft Corporation Public and Community IaaS, PaaS, and SaaS:

The Microsoft offerings below will be provided at a 1.5% discount off the Microsoft list price. Optional Level 1, 2 and 3 support will be provided at 5% off support list price. If support is purchased, then there will be a separate cost for onboarding the Participating Entity on to the support platform. All support, maintenance costs and discounts will be determined based on the level and the scope of the support services requested. This offering will also include:

- Any of the product, solution sets below
- If support is purchased, Level 1 and level 2 supports will be provided by Unisys. Level 3 will be provided by Microsoft based on the discounting and onboarding costs mentioned above.
- Unisys provided monthly billing and reporting as required by this contract
- Unisys provided customer enablement, contract management support, fee management, and fee disbursements as required by the MA, PA, and requirements as stated in this RFP.

Microsoft Corporation	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
Offering Details	Microsoft IaaS includes network services and virtual machines. These services can be combined with PaaS services.	Microsoft Azure PaaS is a growing collection of integrated services – compute, storage, data, application, and networking.	Office 365: Enterprise Cloud productivity and collaboration services: Office, Exchange, SharePoint, Lync, Skype for Business, One Drive for Business, and Project
	Virtual Network: Provision and manage virtual networks in Azure and securely link to your on-premises IT infrastructure.	Azure Web Apps: Build websites with I.Net, PHP, Python, Java, or Node.js and deploy them in seconds.	Microsoft Dynamics: Enterprise Cloud customer relationship management
	Express Route: Connects on-premises infrastructure directly to Azure data centers without using the public Internet (will require telco connectivity).	Azure Storage: Massively scale storage in difference types.	Microsoft Intune: Enterprise Cloud PC and mobile device management
	Virtual Machines: Create new virtual machines or create and upload your own to create preconfigured virtual machines.	SQL Database as a Service	Azure Active Directory: Provides access and identity management solutions, directory services, identity governance, security, and application access and management.
	Traffic Manager: Load balance incoming global traffic across multiple services running in the	Azure Compute: Quickly deploy and manage multi-tier apps, achieving continuous	Azure Multifactor Authentication; Used with Azure Active directory, it

Microsoft Corporation	Infrastructure as a Service (IaaS)	Platform as a Service (PaaS)	Software as a Service (SaaS)
	same or different data centers.	availability.	allows you to safeguard access to data and applications while deploying a simple sign in process.
		Azure Media Services: Encode, store, and stream video and audio at scale.	
		Azure Mobile Services: Create highly functional mobile apps that can access backend capabilities.	
		Azure Stream Analytics: Perform real-time stream processing in the cloud for IoT solutions. Create dashboards and alerts.	
		Azure Logic Apps: Develop powerful integration solutions with SaaS and enterprise applications.	
		Azure API Apps: Expose your application's AP's to your SaaS and enterprise applications.	
Notes	- This entire set of offerings is consumption driven. - List price costs can be estimated using a live calculator available here: https://azure.microsoft.com/en-us/pricing/calculator/		

3. Unisys Partner: Google Public and Community SaaS

The Google offerings below will be provided at a 1.5% discount off the Google list price. Support will be provided by Google under the terms and conditions available here: <https://www.google.com/apps/intl/en/terms/tssg.html>. Additional support options are available and priced based on the scope of the engagement. This offering will also include:

- Any of the product, solution sets below
- Unisys provided monthly billing and reporting as required by the Master Agreement (MA), Participating Agreement (PA), and the requirements as stated in this RFP
- Unisys provided customer enablement, contract management support, fee management, and fee disbursements as required by the MA, PA, and requirements as stated in this RFP.

Google	Software as a Service (SaaS)
Notes: The entire set of offerings is consumption driven. List price costs can be located and estimated using a live calculator available here: https://cloud.google.com/products/calculator/	Google Apps Auth User includes the ability to define your Google Apps user identities for managing Android Applications and devices only.
	With Google Apps for Work, Education, or Unlimited, you can add on this service to set policies to route your outbound email to an encryption server for external recipients to have a way to securely receive and send email to your end users.
	Google Apps No Gmail includes Google Drive, Mobile Device Management, Groups, G+, Sites (unlimited), Videoconferencing, and VoIP.
	Google Apps for Work includes Mail, (30 GB of storage), Calendar, Google Drive, Mobile Device Management, Groups, G+, Sites (unlimited), Videoconferencing, and VoIP.
	Google Apps for EDU Unlimited, for qualifying K-12 and higher education organizations, includes Mail, Google Vault eDiscovery and Archiving (unlimited storage), Calendar, Google Drive, Mobile Device Management, Groups, G+, Sites (unlimited), Videoconferencing, and VoIP.
	Google Apps Deskless includes the same services and is intended for nontechnical workers and infrequent users.
	Google Apps Unlimited includes Mail, Google Vault eDiscovery and Archiving (unlimited storage), Calendar, Google Drive, Mobile Device Management, Groups, G+, Sites (unlimited), Videoconferencing, and VoIP
	Google Apps Unlimited Deskless customers can blend this product into their contract. This offering includes the same services and is intended for nontechnical workers and infrequent users.
	Google Vault eDiscovery and Archiving (unlimited storage)
	Google Apps Vault Deskless customers can blend this product into their contract. This offering includes the same services and is intended for nontechnical workers and infrequent users.

4. Unisys Corporation: Cloud Solution Offerings:

See itemized discounting below. Maintenance and support fees will be calculated and discounted based on the terms of the engagement with the Participating Entity.

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)	Discounting
Unisys Justice Law Enforcement and Border Security Solutions for the Cloud			
Secure Image Management Solution	Unisys developed a highly secure solution to meet the forensic integrity standards and policy requirements of several of the largest police departments in the world. The secure storage of images and other multimedia content occurs in a tamperproof information archive – in effect, a digital evidence vault.	All	20% off list
U-LEAF	U-LEAF provides a POLE-type data model (Person, Object, Location and Event) for the storage and recording of incidents and entities. The POLE model allows entities to be recorded in the system once. The recorded entities, however, can be linked to other entities and events as many times as necessary, to build the picture of an incident, or a network of associations.	All	20% off list
Unisys Social Services Solutions for the Cloud			
Unisys 311	The Unisys “311 solution” is a Multi- Channel Citizen Service Delivery and Engagement solution that can be used by a Local Government or any Governmental Agency to deliver and manage non-emergency services, respond to inquiries and engage with their constituents. This offers multi-channel interaction including social media and mobile applications, supported by a knowledgebase and GIS systems and also interfaces with existing systems to manage work orders and service requests. The solution can be configured to meet specific needs and extend its functionality and flexibility through other apps, and integration to mobile apps, social media and other systems.	All	20% off list
State/Local Government Enterprise Regulatory System (AMANDA)	The Unisys AMANDA Platform (from CSDC) is designed to provide a collection of back office functions such as Licensing, Permitting, Inspections, Land Use, Planning and a number of other functions in various form factors (desktop, mobile and tablet) and can also be deployed either on premise, hosted, or delivered via the Cloud. The solution includes the design, delivery, deployment and ongoing support and maintenance of the solution. The AMANDA platform can be configured specific to client needs in order to extend its functionality and interface with other systems as required.	All	20% off list
Unisys Enterprise Content Management for the Cloud			

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)	Discounting
Infolmage	Unisys Infolmage is the Enterprise Content Management (ECM), Business Process Management (BPM) and Record Management (RM) solution for organizations looking to significantly improve business processes that depend heavily high-volume paper documents, documents generated from Internet transactions, Office documents, and other electronic documents that need to be accessed for automated and manual processes. Infolmage can easily capture, manage, store, and access the content required for cases, inquiries, and process-centric work, regardless of data structure or document origination, with a single intuitive user interface. Unisys Infolmage brings together ECM, imaging, workflow, document management, record management and integration technologies to form an integrated end-to-end solution.	All	25% off list
Unisys Horizontal Solutions for the Cloud			
Stealth	The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. We built a better way to deal with advanced threats for our clients by applying novel approaches to new threats. By substituting traditional hardware topology for software-based cryptography, our Stealth Microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users.	Private and Hybrid Cloud Only. For Public and Community Cloud, see AWS offering above. IaaS, PaaS and SaaS	15% off list
VantagePoint	VantagePoint can also extend beyond the boundaries of IT with uses cases to support relevant business scenarios in, • Security: Strengthens the security posture of organizations through greater visibility into the Stealth enabled infrastructure. • Advanced Data Analytics: Cuts across data silos and sources relevant data to optimize the performance of Advanced Analytics solution. • Cloud and Infrastructure Services: Enables hybrid IT by seamlessly integrating with a variety of data sources and services and presenting them via a digital control panel, irrespective of technology underpinnings. • End User Services: Supports millennial users via personalized and secure access to relevant data and services, across endpoints. • Service Management: Cuts across ITSM platforms to improve service delivery through service automation, orchestration, and aggregation capabilities.	Private and Hybrid Only IaaS, PaaS and SaaS	25% off list

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)	Discounting
	Facilities/Crisis Management: Helps organizations monitor their facilities and keep their Facilities Managers abreast of any security threats and risks. Also acts as a standard communication vehicle for users at times of crisis. (178) With VantagePoint, our clients gain a personalized, secure, and intuitive data and service aggregation platform that cuts across strategic and operational dimensions of business and accelerates digital transformation.		
Unisys Cloud Hosting	Unisys Hosted Private Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can provide, integrate and scale the cloud infrastructure to meet client's needs. Our management platform enables a single pane of glass to effectively manage client's workloads. Unisys Hybrid Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can source, integrate and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our management platform enables a single pane of glass to effectively manage a variety of clouds within an organization, whether they be public or private.	All	20% off list
ServiceNow	ServiceNow offers enterprise service management software for human resources, law, facilities management, finance, marketing, and field operations in the cloud. ServiceNow has its STAR Self-Assessment available on the CSA's website. ServiceNow specializes in ITSM applications and provides forms-based workflow application development. ServiceNow has open integration options to variety platforms such as: Salesforce, SharePoint, and BMC Remedy Action Request System.	All	15% off list
Unisys ClearPath Forward!	Unisys ClearPath Forward is an Intel based fabric computing platform designed to run business critical applications that require predictable performance and low transport latency. The ClearPath Forward platform provides hardware partitioning technology similar to the Mainframes, but is designed to run Enterprise Windows and Linux operating environments on commodity server hardware. The fabric interconnect included with the ClearPath Forward platform provides contemporary transport speeds or higher. The ClearPath Forward	Private and Hybrid Only IaaS, PaaS and SaaS	3% off list

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)	Discounting
	platform fabric can be used in a hybrid mode with industry available virtualization technologies including VMware® and Docker®.		
Unisys White Label Offerings	Unisys has strategic, long lasting relationships in the industry and is pleased to offer our partner products surrounding the cloud market space. This is an evolving list. Our current partner list includes: Salesforce, Oracle, Verizon, SHI, Decision Lens, Aptio, NetApp, VMware, EMC, NetApp, Birst, Okta, Box, WatchDox. We have created webpage on our public site that is updated with our list of current partners. The page is available here: http://www.unisys.com/ms/wsca-cloud-hosting/	All	5% off list

5. Unisys Corporation: Professional Services

The following professional services are available to Participating Entities. These services are provided by Unisys in accordance with time and materials.

Unisys Professional Services	Description
Cloud Advisory Services	Unisys Cloud Advisory Services provide strategic and financial guidance on aligning IT with business objectives. This starts with a roadmap that outlines the vision of a Hybrid IT based on a combination of existing data center, internal cloud, and external cloud resources to provide agility, flexibility, and control.
Data Center Planning, Design, and Implementation Services	Unisys Data Center Planning, Design, and Implementation Services offer a complete range of services that delivers a cohesive, end-to-end optimization of data centers. With a wide range of services for discovery, analysis, optimization, virtualization, consolidation, and migration of data centers that can complement client efforts and fill gaps in skills and capacities and a combination of world-class people, processes, and technology with the program and project management expertise, Unisys transforms clients' existing data centers to a business engine that provides agility at a lower cost.
CloudBuild Services	Unisys CloudBuild Services enable organizations to successfully build a cloud that is integrated with the overall business process, transforming their existing infrastructure to an agile IT-as-a-service model. The Unisys "8-Tracks" model, a 360-degree cloud view approach, covers eight critical data center domains. Together with ConOps, which includes industry best practices, Unisys enables your cloud infrastructure to meet the security, regulatory, and compliance requirements that enable us to deliver the most secure and reliable cloud in the marketplace.
Hybrid Cloud Strategy	Unisys Hybrid Cloud strategy helps organizations overcome key challenges when planning to implement a cloud environment. We provide governments and businesses with a comprehensive cloud architecture that gives clients cloud on their terms. We can source, integrate, and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our homegrown VantagePoint management platform enables a single pane of glass to effectively manage a variety of clouds in an organization, whether they be public or private.
Unisys Platform Services	
Platform Assessment Services	Unisys Platform Assessment Services enable governments and businesses to foster their platform adoption initiatives. We analyze organizations' business goals and objectives along with their technology landscape of their enterprise application portfolio and recommend a best-fit Platform Suitability Analysis to help them make the right decision in their move to PaaS
Architecture Design Services	Unisys Architecture Design Services assist governments and businesses with designing an effective architecture suited for cloud. We provide several concepts and best practices that are essential to build highly scalable applications in the cloud – be it on-premises, public, or hybrid.
PaaS enablement Services	Unisys PaaS enablement Services analyze an organization's current set of applications and provide a strategy to use the best practices and engagement models of various platform providers such as Microsoft Azure, IBM Bluemix, SAP HCP, and SFDC AppCloud.
Application Migration Services	Unisys Application Migration Services assist organizations in their cloud migration activities by using tools and frameworks adhering to Unisys best practices that promise minimal downtime without affecting the day-to-day business processes.
Application Portability Services	Unisys Application Portability Services help organizations in developing cross-platform applications that can be scaled across multiple cloud platforms.
Application Development Services	Unisys Application Development Services assist organizations with building rich, interactive applications focused on business logic and workflows using visual tools as well as cloud-based tools, architectures, and services that make their applications cloud ready.
IoT Development Services	Unisys IoT Development Services enable organizations to use IoT Services across various cloud platforms that enable apps to communicate and consume data collected by the connected devices, sensors, and gateways.
API Management Services	Unisys API Management Services provide a solution that addresses the aspects of the application programming interface

Unisys Professional Services	Description
	(API) life cycle for on-premises and cloud environments and offer capabilities to create, run, manage, secure and monetize APIs and microservices that deliver an integrated user experience and enable rapid deployment and simplified administration of APIs.
Horizontal Technology Integration Services	
Horizontal Technology Integration Services	Unisys is a world recognized leader in integrating technology infrastructures across platforms and vendors. We particularly specialize in integrating solutions surrounding service desks, cloud provisioning, data center management, and the ITIL framework.
Unisys Service Management Services	
Maturity and Platform Technology Assessments	Unisys Service Management Consulting Services leverage subject matter experts with more than 15 years of experience in working and delivering services for ITIL and service management. We leverage that knowledge to provide best practices, lessons learned, and roadmap development for Service Management disciplines and benchmark them with industry and operational best practices.
Service Management Platform Implementation	Unisys brings experience in moving existing ITSM tool information, processes, and requirements to new cloud enterprise Service Management platforms. By leveraging industry best practices and structured methodologies, Unisys can establish transition and transformation plans that minimize risk to clients' ongoing operations and deliver speed to value with the new cloud-based Service Management platform.
Enterprise Cloud Service Management Platform Support	Unisys can provide ongoing day-to-day support in the ongoing management of a client's Service Management platform. This service allows the organization to focus on the business aspects instead of the routine activities necessary to support and maintain a platform. Unisys provides cost-effective solutions that leverage staff in many locations, driving the cost of servicing the environment to effective and efficient levels.

Unisys Corporation: Labor Rates

Unisys is pleased to provide the following labor categories to support the Participating Entities.

Title	Rate Per Hour	Role Descriptions
Cloud Practice Director	\$285.40	Manages the delivery of contracted services to clients to maintain time, quality, and cost of delivery. Maintains tight control over the project schedule, risks, scope of work, and budget; confirms that operational teams and subcontractors have a clear understanding of client requirements. Builds and maintains strong client relationships; provides day-to-day client advice and support. -Promotes the organization's capabilities to clients, identifies sales opportunities to be forwarded to the account managers, and achieves contract extensions or additional business in the accounts. Contracts may involve both short- and long-term commitment of service and vary significantly in value or strategic importance.
Cloud Program Manager	\$164.27	Manages very complex programs, high-risk programs, or both. May manage fixed price contracts. Oversees program budget and schedules. May direct staff. Has primary responsibility for program growth; may market new technology or acquire follow-on business. The total value of programs under this manager's responsibility is more than \$100 million (over the life of the contract, not annually). May be responsible for programs of a lower dollar value if they are more complex or developmental. This description does not include engineers or other individuals who are temporarily assigned program management responsibilities and technical functional managers for a program.

Title	Rate Per Hour	Role Descriptions
Transition Manager	\$124.45	Manages programs of moderate risk and complexity or may have deputy responsibility for a large program. Is frequently involved simultaneously in several programs. Oversees program budget and schedules prepared by subordinate staff. May have supervisory responsibilities, including hiring, firing, and salary and performance management. May have primary responsibility for program growth. Serves a primary client contact. The total value of programs under this manager's responsibility is \$10 million to \$50 million over the life of the contract, not annually. (For India, the incumbent manages a project team of 100 to 200 employees.) May be responsible for programs of a lower dollar value if they are more complex or developmental. This description does not include engineers or other individuals who are temporarily assigned program management responsibilities and technical functional managers for a program.
Project Manager	\$124.45	Oversees one or more of the following: hardware/software architecture, operating systems, system configuration and control, and Cloud Migration projects. Provides subordinates with guidance based on organizational goals and company policy. Accomplishes results through subordinates' meeting schedules and resolving technical or operational problems. Develops and administers budgets, schedules, and performance standards. Assists in the development of overall objectives and long-range goals for the assigned area. Manages employees with higher skills, more than two clients, projects that are medium to complex. Has more project management and transformation responsibilities.
Quality Assurance Manager	\$137.72	Provides data-driven analysis of business processes; helps to implement improvements in the client focus, efficiency, accuracy, and effectiveness of these processes. Delivers champion awareness training; implements business processes in organizations. Provides expertise in external standards and validations (e.g., ISO 9000). Proficient in a wide range of improvement tools, including statistics. Leads and directs a team of professional individual contributors with specific assigned program objectives. Interprets and applies policy; translates goals to programs and projects.
Lead Cloud Architect	\$202.43	This role is subject to Unisys billing rates as well as productivity, utilization, and total chargeability metrics. Only billable/direct client chargeable employees are to be assigned to this role. Works in and across practices and organizations to design leading-edge technology or application solutions that lead to profitable revenue growth for Unisys. Conceptualizes, architects, designs, implements, and supports integrated solutions for client engagements. Demonstrates extensive knowledge of industry, technology, and strategy trends; uses this knowledge to generate profitable revenue growth for Unisys. Builds relationships with product development organizations; is recognized by these organizations as a key product champion who adds value through exposure to client situations. Advises, analyzes, researches, designs, installs, and implements complex integrated solutions for an entire enterprise. Interacts extensively with the client's senior management team (e.g., IT directors) on business, systems, architectural, and technical issues. Provides other architects with mentoring and team leadership. Develops the integrated technology requirements project plan. Integrates complex solutions with new or existing client environments. Uses advanced diagnostic analytical and design capabilities. Confirms that the client is provided with the best solution (e.g., completes solution integrity engineering and development attributes).

Title	Rate Per Hour	Role Descriptions
Infrastructure Migration Consultant	\$202.43	Analyzes and evaluates major system project requirements of considerable complexity and that require a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements and provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex multiple interlinked systems. Defines system objectives and prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions by applying systems and methods of engineering techniques. Reviews alternate approaches and selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned problems. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development to a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as organization spokesperson for advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.
Monitoring Specialist	\$114.49	Performs Tier 1 operating system-specific tasks to support a host environment. Supports and maintains servers, networks, messaging, mainframes, database environments, and/or host infrastructures. Responds to service outages and other problems and issues; serves as a Tier 1 escalation point for well-defined incidents and basic problems, working under general supervision.
Storage Specialist	\$202.43	Delivers Level 2 remote hardware and software support services to clients to resolve product use and multi-product/platform problems, and/or questions relating to storage on enterprise systems across networks and in the cloud. Provides referrals, dispatches, or both to other service providers to confirm that the client's service level and technical requirements are met. Provides design and implementation support for storage innovations. Prepares and approves technical documentation; verifies that technical and client documentation is clear, accurate, and complete. Identifies training needs. May develop and conduct training sessions for other analysts and clients. May lead teams of analysts on defined projects.

Title	Rate Per Hour	Role Descriptions
Application Migration Architect	\$202.43	This role is subject to Unisys billing rates, productivity, utilization, and total chargeability metrics. Only billable/direct client chargeable employees are to be assigned to this role. Works in and across practices and organizations to design leading-edge technology or application solutions that lead to profitable revenue growth for Unisys. Conceptualizes, architects, designs, implements, and supports integrated solutions for client engagements. Demonstrates extensive knowledge of industry, technology, and strategy trends; uses this knowledge to generate profitable revenue growth for Unisys. Builds relationships with product development organizations; is recognized by these organizations as a key product champion who adds value through exposure to client situations. Advises, analyzes, researches, designs, installs, and implements complex integrated solutions for an entire enterprise. Interacts extensively with the client's senior management team (e.g., IT directors) on business, systems architectural, and technical issues. Provides other architects with mentoring and team leadership. Develops the integrated technology requirements project plan. Integrates complex solutions with new or existing client environments. Uses advanced diagnostic, analytical, and design capabilities. Confirms that the client is provided with the best solution (e.g., completes solution integrity engineering and development attributes).
Database Consultant	\$160.95	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.

Title	Rate Per Hour	Role Descriptions
Risk, Compliance and Security Consultant	\$160.95	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; provides direction in the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.

Title	Rate Per Hour	Role Descriptions
Network Specialist	\$160.95	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; provides direction for the identification of problems and their potential resolution. Provides analytical support in the conceptualization, development, and implementation of complex multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Acts as advisor to management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.
Service Delivery Manager	\$155.97	Is accountable for managing the delivery of contracted outsourced services such as business process and information technology to clients with account revenue of \$6 million to \$15 million per year (levels are based on the size of account revenue). Serves as the primary point of contact to one or more clients on overall and day-to-day service delivery. Verifies that SLAs (service level agreements) and KPIs (key performance indicators) defined in the relevant contracts are met or exceeded. To enable implementations and ongoing services to be delivered on time and meet client requirements, maintains tight control over the project's schedule, risks, scope of work, and budget. Builds and maintains strong client relationships; participates in client meetings on performance to confirm client satisfaction. Verifies that operational teams and subcontractors maintain a clear understanding of the client's needs; provides day-to-day client advice and support. Promotes the organization's capabilities; works to achieve contract extensions or win additional business in the accounts. This description does not include specific technical functional managers or other individuals who are temporarily assigned project management responsibilities.

Title	Rate Per Hour	Role Descriptions
Disaster Recovery Specialist	\$160.95	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; and provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex, multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.

Title	Rate Per Hour	Role Descriptions
AWS Specialist	\$202.43	Analyzes and evaluates major system project requirements of considerable complexity that require a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex, multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas and guides their development into a final product. Erroneous decisions or recommendations would typically result in failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as organization spokesperson on advanced projects and/or programs. Acts as advisor to management on advanced technical research studies and applications. May provide work leadership for lower level employees. Considered expert in field. This description does not include those with full supervisory responsibility.

Title	Rate Per Hour	Role Descriptions
Azure Specialist	\$202.43	Analyzes and evaluates major system project requirements of considerable complexity requiring a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements and provides direction in the identification of problem and potential resolution. Provides analytical support in the conceptualization, development and implementation of complex, multiple inter-linked systems. Defines system objectives and prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements and develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches and selects appropriate methodology. Addresses complex problems where analysis of situations or data requires and evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections in technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems and provides solutions which are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.

Title	Rate Per Hour	Role Descriptions
VMware Specialist	\$202.43	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; and provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex, multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.
ServiceNow Specialist	\$202.43	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; and provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex, multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological

Title	Rate Per Hour	Role Descriptions
		capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.
Google Specialist	\$202.43	Analyzes and evaluates major system project requirements of considerable complexity that requires a thorough understanding of the parameters affecting and interfacing with the system. Reviews user requirements; and provides direction for the identification of problems and their potential resolution. Provides analytical support for the conceptualization, development, and implementation of complex, multiple interlinked systems. Defines system objectives; prepares system design specifications to meet user requirements and satisfy interface problems. Formulates logical statements of user requirements; develops solutions through application of systems and methods of engineering techniques. Reviews alternate approaches; selects appropriate methodology. Addresses complex problems for which the analysis of situations or data requires an evaluation of often intangible factors. Reviews literature, patents, and current practices relevant to the solution of assigned projects. Recommends corrections to technical applications and analysis. Evaluates vendor capabilities to provide required products or services. Is accountable for technical contributions that lead to a profitable return on technical investment. Applies advanced technical principles, theories, and concepts. Contributes to the development of new principles and concepts. Works on unusually complex technical problems; provides solutions that are highly innovative and ingenious. Works under consultative direction toward predetermined long-range goals and objectives. Assignments are often self-initiated. Determines and pursues courses of action necessary to obtain desired results. Work is checked through consultation and agreement with others instead of a formal review by the leader. Develops advanced technological ideas; guides their development into a final product. Erroneous decisions or recommendations would typically lead to failure to achieve critical organizational objectives and affect the image of the organization's technological capability. Serves as the organization's spokesperson on advanced projects, programs, or both. Advises management on advanced technical research studies and applications. May provide lower level employees with work leadership. Considered expert in field. This description does not include those with full supervisory responsibility.
System Administrator	\$114.49	Delivers Level 2 remote hardware and software support services to clients to resolve product use, multiproduct/platform problems and/or questions relating to enterprise systems, networks, and application software, as well as desktop applications that are beyond the scope of the client's Level 1 help desk support function. Provides referrals, dispatches, or both to other service providers to confirm that the client's service level and technical requirements are met. Provides the client and Unisys management with alerts and information on a situation's status. Coordinates critical client issues and implementations of new products, systems, or both. Prepares and approves technical documentation; verifies that the documentation is clear, accurate, and complete. Identifies training needs. May develop and conduct training sessions for other analysts and clients. May lead teams of analysts on defined projects.

State of Utah NASPO Cloud Solutions

Technical Proposal

Solicitation No: CH16012



UNISYS

TABLE OF CONTENTS

1.0	(M) Signature page (rfp Section 5.1)	1- ~~32~~
2.0	(M) Executive Summary (RFP Section 5.4)	2-1
3.0	Mandatory Minimums (RFP Section 5).....	3-4
4.0	Business Profile (RFP Section 6).....	4-1
5.0	Organization Profile (Rfp section 7).....	5-1
6.0	Technical response – aws – ms included (RFP Section 8).....	6-1
7.0	Confidential, protected or proprietary information.....	7-1
8.0	Exceptions and/or Additions to the Standard Terms and Conditions	8-1
Appendix 3 – AWS CAIQ		
Appendix 4 – Microsoft O365		
Appendix 5 – Microsoft Azure		
Appendix 6 – Microsoft Dynamics		
Appendix 7 – Google CAIQ		
Appendix 8 – AWS SLA		
Appendix 9 – Microsoft SLA		
Appendix 10 – Google Apps SLA		
Appendix 11 – Unisys-1		
Appendix 12 – AWS-1		
Appendix 13 – Google-1		
Appendix 14 – AWS Risk and Compliance White Paper		
Appendix 15 – Windows Azure Security and Compliance		
Appendix 16 – Google DPA		
Appendix 17 – Microsoft Online Services Terms		
Appendix 18 – Google Compute (Cloud) SLA		
Appendix 19 _ Getting Started With VMs on Windows Azure		
Appendix 20 – AWS+Access+Policy+(State)		
Appendix 21 –AWS+Access+Policy+(State)		
Appendix 22 – Google Apps For work		
Appendix 23 – Google Apps Security		
Appendix 24 – Windows Azure-Secure Privacy Compliance		

1.0 (M) SIGNATURE PAGE (RFP SECTION 5.1)


Section Title: RFP Signature Page. The Lead State's Request for Proposal Signature Page completed and signed. See Section 5.1 of the RFP.

5.1 (M) SIGNATURE PAGE Proposals must be submitted with a vendor information form, located on Bidsync as an attachment to the RFP, which must contain an ORIGINAL HANDWRITTEN signature executed in INK OR AN ELECTRONIC SIGNATURE, and be returned with the Offeror's proposal.

Unisys Response: Our vendor information form is below.



State of Utah Vendor Information Form

Lead Company Name (include division if applicable) Unisys Corporation		Federal Tax Identification Number 38-0387840		State of Utah Sales Tax ID Number 12278699-002-STC	
Outgoing Address 801 Lakeview Drive Ste. 100		City Bluebell	State PA	Zip Code 19422	
Remittance Address (if different from ordering address) 810 Lakeview Drive Ste. 100		City Bluebell	State PA	Zip Code 19422	
Type: ☐ Proprietorship ☐ Partnership ☐ Government ☒ For-Profit Corporation ☐ Non-Profit Corporation		Company Contact Person Mr. Harsh Bajpai			
Telephone Number (include area code) 703-439-6200		Fax Number (include area code)			
Company's Internet Web Address http://www.unisys.com		Email Address Harsh.Bajpai@unisys.com			
Offeror's Authorized Representative's Signature 					
Type or Print Name Mr. Harsh Bajpai					
Position or Title of Authorized Representative Director, Cloud Solutions and Services					
Date MARCH 10, 2016					

12/15/2014

2.0 (M) EXECUTIVE SUMMARY (RFP SECTION 5.4)

- **Section Title: Executive Summary.** The one or two page executive summary is to briefly describe the Offeror's Proposal. This summary should highlight the major features of the Proposal. It must indicate any requirements that cannot be met by the Offeror. The Lead State should be able to determine the essence of the Proposal by reading the executive summary. See Section 5.4 of the RFP.
- **(M) EXECUTIVE SUMMARY**
- **Offerors must provide an Executive Summary of its proposal. An Executive Summary should highlight the major features of an Offeror's proposal. Briefly describe the proposal in no more than three (3) pages. The evaluation committee should be able to determine the essence of the proposal by reading the Executive Summary. Any requirements that cannot be met by the Offeror must be included.**

Unisys Response:

Unisys appreciates the importance of this innovative and flexible solicitation. We are pleased to offer the cloud deployment models (private, community, public, and hybrid) and the underlying cloud services models (IaaS, PaaS, and SaaS) requested in this solicitation. We will also provide architecture, design, implementation, and value-add services that will support the various cloud solution options. The following table illustrates the depth and breadth of the offerings in our proposal.

	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Private Cloud	Unisys	Unisys	Unisys
Hybrid Cloud	Unisys	Unisys	Unisys
Cloud Professional Services	Unisys	Unisys	Unisys
Cloud Value Add Services	Unisys	Unisys	Unisys


Our cloud partners in this proposal are Amazon Web Services (AWS), Microsoft and Google. Unisys will resell some of the offerings through our own partner agreements and fill orders for the others directly, as required by the Participating Entity.

Unisys is offering the State of Utah and the Participating Entities a set of world-class, best-of-breed cloud providers. Proven solution options will provide the Participating States with a diverse set of products, solutions and services that are typically needed for its evolving requirements. Our ability to deliver all types of cloud solutions across all partners; seamlessly and expeditiously with our own certified and experienced staff will provide the Participating Entities with a quick reach-back vehicle that they can leverage as needed. Our proven, past performance in supporting cloud transformations and cloud implementations for clients in state, local, and Federal governments will provide the States with a qualified level of confidence in Unisys. Our strong, formal partnerships allow us direct access to the latest resources of our partners, which we can leverage expeditiously. Finally, Unisys can provide the Participating Entities with a one-stop shop for all types of cloud solutions and services. This will minimize risk to the State and establish transparency and vendor accountability, thereby improving the State's overall effectiveness and efficiency in all cloud-related matters.

The Unisys Solution Approach for NASPO – Transforming Today's Enterprise to Tomorrow's Digital Government

As a trusted agent and partner, Unisys has provided value-added information technology services for more than 50 years to clients in state, local, U.S. Federal, and foreign governments in 180 countries. As the prime holder of the WSCA Public Cloud Hosting Service contract, Unisys exceeds current contract

requirements by providing best-in-class experience, quality, and timely service delivery. We pride ourselves as being an honest broker because we understand our clients' requirements, including those of foreign governments throughout the world.



Our ability to rapidly deploy needed resources and specialized services is unparalleled. Our efficient and effective approach focuses on providing precisely the services required at the best possible cost. Our subject matter experts power a portfolio of capabilities aligned with our clients' changing priorities with a focus on delivering next-generation Digital Government Solutions and Services.



UNISYS Digital Government Transformational Business Benefits

It is critical in today's digital environment that the State partners with the right organization that can bring the knowledge, skills, and domain expertise to the partnership. As the incumbent, our goal is to build on the quality service the State receives today. Unisys has a deep appreciation and understanding of providing mission-critical and innovative solutions to the public sector. Our combined cloud hosting and services capabilities in our proposal will exceed the State's key objectives as follows:

- Service enablement:** The pace of change is accelerating faster than public organizations can adjust their plans, investments, and—more important—their employee base. The pressure to adopt cloud services is intense; adopting cloud services haphazardly could be detrimental. There are persistent concerns for security, compliance, regulations, and cost. Unisys brings staff to deliver comprehensive consulting, integration, and support services that enable clients to migrate to cloud services while maintaining operational compatibility and transforming their existing infrastructure for the delivery of an agile IT-as-a-Service model. CloudBuild Services is a multivendor offering for the planning, design, and implementation of an enterprise cloud. These services will help the State to avoid the “cloud-in-a-corner” syndrome, decrease risk in migrating to cloud services, and decrease potential project failures and budget overruns.

- **Improved security and comfort level:** As more and more of an organization’s mission-critical and day-to-day operations move to the cloud, security becomes much more important. Unisys brings our award-winning Stealth security products, which provide a smarter way to manage risk in enterprises. Stealth is now fully integrated with the Amazon Web Services cloud infrastructure. Stealth on AWS will provide the State with a trusted path to extend its security protections to the cloud, unlocking huge cost savings and providing needed agility to respond to constituents’ concerns.
- **Cost efficiencies:** Unisys understands the financial benefits available to users of cloud computing services. In particular, the very competitive pay-as-you-go cloud computing model in our proposal, paired with full control over fine-grained characteristics of the State’s cloud-based infrastructures, will enable users to achieve significant savings in ongoing operating expenses.
- **Flexibility and scalability:** Orchestration in a cloud computing environment empowers Unisys clients to have complete control over the timing and availability of IT resources. We combine tasks into workflows so the provisioning and management of various IT components and their associated resources can be automated. The productivity gained from the ability to provision services in minutes (not days or weeks, as in many IT environments) enables users to respond to changing conditions and new requirements quickly without sacrificing service delivery quality. Our proposal exceeds the RFP requirements by enabling advanced automation capabilities not available in most cloud computing environments.
- **Significant reduction of support staff time:** By freeing resources currently performing easily automatable tasks, and adopting a “factory” model for service delivery, IT organizations can reallocate staff to more productive uses that improve the ability of states and other entities to meet their goals sooner or more effectively. Our proposal fully supports the needs of Participating Entities to offload activities to automated environments in the cloud computing ecosystem.

“Of the IT vendors, Unisys continues to have one of the clearest sets of offerings for secure private, public and hybrid models.” – 451 Research

3.0 MANDATORY MINIMUMS (RFP SECTION 5)

5.2 (M) COVER LETTER

Proposals must include a cover letter on official letterhead of the Offeror. The cover letter must identify the RFP Title and number, and must be signed by an individual authorized to commit the Offeror to the work proposed. In addition, the cover letter must include:

March 10, 2016

Christopher Hughes, Assistant Director
State of Utah, Division of Purchasing
3150 State Office Building, Capitol Hill
Salt Lake City, UT 84114-1061

Subject: Unisys Response to Utah Solicitation Number CH16012 – NASPO ValuePoint Cloud Hosting Services

Dear Mr. Hughes:


Unisys is pleased to provide our response to Solicitation Number CH16012 – NASPO ValuePoint Master Agreement for Public Cloud Hosting Services. We sincerely appreciate the importance of this solicitation.

Unisys is pleased to offer the requested cloud deployment models (private, community, public, and hybrid) and the underlying cloud services models (IaaS, PaaS, and SaaS) as part of our proposal. We will also provide architecture, design, implementation, and value-add services that will support the various cloud solution options. Our cloud service provider partners in this proposal are Amazon Web Services (AWS), Microsoft, and Google.

Unisys believes that offering the State of Utah and the Participating Entities a set of world-class, best-of-breed cloud providers is a win-win proposition for the State. Proven solution options will provide the State with a full set of services that are typically needed for its diverse IT ecosystems. Our ability to deliver cloud solutions seamlessly and expeditiously with our own certified and experienced staff will provide Participating Entities with a quick reach-back vehicle that they can leverage as needed. Our proven past performance in supporting cloud transformations and implementations for clients in state, local, and Federal governments will provide the State of Utah with a qualified level of confidence in Unisys. Our strong, formal partnerships with AWS, Microsoft and Google allow us direct access to the latest resources of our partners, which we can leverage expeditiously. Finally, Unisys can provide the Participating Entities with a one-stop shop for cloud solutions and services. This will minimize risk to the State and establish transparency and vendor accountability, thereby improving the State's overall effectiveness and efficiency in cloud-related matters.

Unisys, the prime contractor, is a public Fortune 500 corporation (NYSE: UIS) that operates in more than 100 countries and has more than 20,000 associates. For more than 50 years, we have provided the U.S. Federal and state governments with mission-critical information technology (IT) solutions and related consulting services, including small- and large-scale Cloud Services. More than 45 percent of our revenue of \$4 billion comes from our valued long-term government clients. Understanding our client needs by working closely with clients with our consultants is a cornerstone of our heritage and success in government and business in the United States and around the world.

Sincerely,



Harsh Bajpai, Unisys Corporation
Director, Cloud Solutions & Services

Email: Harsh.Bajpai@unisys.com; Telephone: 703-439-6200

5.2.1 A statement indicating the Offeror's understanding that they may be required to negotiate additional terms and conditions, including additional administrative fees, with Participating Entities when executing a Participating Addendum.

Unisys understands that we may be required to negotiate additional terms and conditions, including additional administrative fees, with Participating Entities when executing a Participating Addendum.

5.2.2 A statement naming the firms and/or staff responsible for writing the proposal.

Unisys is responsible for the writing of this proposal with information provided from our partners.

5.2.3 A statement that Offeror is not currently suspended, debarred or otherwise excluded from federal or state procurement and non-procurement programs.

Unisys and our partners for this proposal are not currently suspended, debarred, or otherwise excluded from Federal or state procurement and non-procurement programs.

5.2.4 A statement acknowledging that a 0.25% NASPO ValuePoint Administrative Fee and any Participating Entity Administrative fee will apply to total sales for the Master Agreement(s) awarded from the RFP.

Unisys acknowledges the statement above.

5.2.5 A statement identifying the service model(s) (SaaS, IaaS, and/or PaaS) and deployment model(s) that it is capable of providing under the terms of the RFP. See Attachment C for a determination of each service model subcategory. The services models, deployment models and risk categories can be found in the Scope of Services, Attachment D. Note: Multiple service and/or deployment model selection is permitted, and at least one service model must be identified. See Attachment H.

5.2.6 A statement identifying the data risk categories that the Offeror is capable of storing and securing. See Attachment D and Attachment H.

Unisys is pleased to offer all four cloud deployment models: public cloud, private cloud, community cloud, and hybrid cloud. Under these deployment models, we are also pleased to offer all three cloud service delivery methods: IaaS, PaaS, and SaaS.

Attainment of certifications is a lengthy ongoing process for vendors. Additionally, certification standards are constantly evolving. Our partners in this proposal are at a FISMA Moderate level. Therefore, across our cloud deployment models and delivery options, we are positioned to secure and store FISMA Low and Moderate risk data types. For FISMA High, Unisys Stealth for data in motion, when combined with encryption technologies for data at rest, can meet the requirements of FISMA High risk data. For the private cloud model, we can design a cloud environment based on the required standards and controls defined by the client.

In accordance with the requirements of this RFP, we completed Attachment H – Identification of Service Models Matrix and provided it in **Attachment H**.

Attachment H – Identification Service Models

Service Model:	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered:
SaaS	Unisys AWS Microsoft Google	Unisys AWS Microsoft Google	See response above	All
IaaS	Unisys AWS Microsoft	Unisys AWS Microsoft	See response above	All
PaaS	Unisys AWS Microsoft	Unisys AWS Microsoft	See response above	All

5.3 (M) ACKNOWLEDGEMENT OF AMENDMENTS

If the RFP is amended, the Offeror must acknowledge each amendment with a signature on the acknowledgement form provided with each amendment. Failure to return a signed copy of each amendment acknowledgement form with the proposal may result in the proposal being found non-responsive.

Unisys acknowledges the amendments released under this RFP. Unisys provides our signed copy of the amendment acknowledgement form below.

AMENDMENT ACKNOWLEDGEMENT FORM**ACKNOWLEDGEMENT OF AMENDMENTS TO RFP (SOLICITATION CH16012)**

This attachment represents that the Offeror has read, reviewed, and understands the totality of Solicitation CH16012, including the final RFP document posted on February 10, 2016.

By signing below, the Offeror attest to reviewing the documents listed above.

UNISYS CORPORATION
Offeror
[Signature]
Representative Signature

5.4 (M) EXECUTIVE SUMMARY

Offerors must provide an Executive Summary of its proposal. An Executive Summary should highlight the major features of an Offeror's proposal. Briefly describe the proposal in no more than three (3) pages. The evaluation committee should be able to determine the essence of the proposal by reading the Executive Summary. Any requirements that cannot be met by the Offeror must be included.

Unisys provides our executive summary in Section 2.0, Executive Summary.

5.5 (M) GENERAL REQUIREMENTS

5.5.1 *Offeror must agree that if awarded a contract it will provide a Usage Report Administrator responsible for the quarterly sales reporting described the Master Agreement Terms and Conditions, and if applicable Participating Addendums.*

If Unisys is awarded this contract, we will provide a Usage Report Administrator responsible for the quarterly sales reporting described in the Master Agreement Terms and Conditions, and if applicable, Participating Addendums.

5.5.2 *Offeror must provide a statement that it agrees to cooperate with NASPO ValuePoint and SciQuest (and any authorized agent or successor entity to SciQuest) with uploading an Offeror's ordering instructions, if awarded a contract.*

Unisys agrees with the statement above.

5.5.3 *Offeror must at a minimum complete, provide, and maintain a completed CSA STAR Registry Self-Assessment. Offeror must either submit a completed The Consensus Assessments Initiative Questionnaire (CAIQ), Exhibit 1 to Attachment B, or submit a report documenting compliance with Cloud Controls Matrix (CCM), Exhibit 2 to Attachment B. Offeror must also represent and warrant the accuracy and currency of the information on the completed. Offerors are encouraged to complete and submit both exhibits to Attachment B.*

5.5.4 *Offeror, as part of its proposal, must provide a sample of its Service Level Agreement, which should define the performance and other operating parameters within which the infrastructure must operate to meet IT System and Purchasing Entity's requirements.*

In the private cloud space, government clients contract Unisys to design and implement a cloud platform on their premises. We build this platform to order in accordance with the client's requirements for security, governance, compliance, certification, and accreditation. Unisys then can manage a custom environment for the client or return it to the client to manage it. In the hybrid cloud space, government clients run a private environment that mirrors an external environment, which could be public, community, or private. Therefore, we are providing responses to RFP requirements 5.5.3 and 5.5.4 as they apply to our cloud partners in this proposal.

Appendix 3 – AWS CAIQ

Appendix 4 – Microsoft Office365 CCM

Appendix 5 – Microsoft Azure CCM

Appendix 6 – Microsoft Dynamics CCM

Appendix 7 – Google CAIQ

The appendix to this proposal provides our collective responses to RFP requirement 5.5.3 here:

Appendix 8 – AWS SLA

Appendix 9 – Microsoft SLA

Appendix 10 – Google Apps SLA

5.7 RECERTIFICATION OF MANDATORY MINIMUMS AND TECHNICAL SPECIFICATIONS

Offeror must acknowledge that if it is awarded a contract under the RFP that it will annually certify to the Lead State that it still meets or exceeds the technical capabilities discussed in its proposal.

Unisys acknowledges that if we are awarded a contract under the RFP, we will annually certify to the Lead State that we still meet or exceed the technical capabilities discussed in our proposal.

4.0 BUSINESS PROFILE (RFP SECTION 6)

Section Title: Business Profile: This section should constitute the Offeror's response to the items described in Section 6 of the RFP. An Offeror's response must be a specific point-by-point response, in the order listed, to each requirement in the Section 6 of the RFP.

6 BUSINESS INFORMATION

6.1 (M)(E) BUSINESS PROFILE

Provide a profile of your business including: year started, organizational structure, client base (including any focus by region, market sector, etc.), growth over the last three (3) years, number of employees, employee retention rates (specific for employees that may be associated with the services related to the RFP) over the last two (2) years, etc. Businesses must demonstrate a minimum of three (3) years of experience providing cloud solutions for large scale projects, including government experience, to be eligible for award.

Unisys Response:

Unisys Corporation was founded in 1886 and is based in Blue Bell, Pennsylvania. We provide information technology services and products worldwide. Approximately 88 percent of our revenue is derived from services. The remaining 12 percent of our revenue is derived from the sale of technology products and solutions. Cloud and infrastructure services represent 52 percent of our total services revenue. Sales of commercial products and services to various agencies of the U.S. Federal government represented about 19 percent of our total consolidated revenue in 2015. Our market capitalization is more than \$500 million. Our revenue for the year ending in 2014 was \$3.4 billion. We employ more than 20,000 associates worldwide. Our services organization in North America alone is 3400 persons strong. Our attrition rate for cloud and data center infrastructure services within North America is less than 23%.

Section 6.2 provides our experience in working on large-scale projects to provide government agencies with cloud solutions.

6.2 (M)(E) Scope of Experience

Describe in detail the business' experience with government or large consortium contracts similar to the Master Agreements sought through this RFP. Provide the approximate dollar value of the business' five (5) largest contracts in the last two (2) years, under which the Offeror provided Solutions identical or very similar to those required by this RFP. Government experience is preferred.

Unisys Response:

Scope of Experience: 1 Contract Term: Current Segment: U.S. Federal Government Solution/Offering: Private Cloud Storage	U.S. Department of the Treasury, Internal Revenue Service (IRS) GSA's Alliant Government-wide Acquisition Contract
On July 26, 2012, Unisys was awarded the Enterprise Storage Acquisition task order from the Internal Revenue Service (IRS) to provide private cloud-based storage for the agency's information records. The task order, awarded against GSA's Alliant Government-wide Acquisition Contract, had a 1-year base period and nine 1-year options worth up to \$139 million over 10 years. Unisys is fulfilling our third contract option year. Under the task order, Unisys acquired storage assets owned and managed by IRS and now manages more than 6 PB of storage space in several IRS data centers and facilities. We planned, built, deployed, and now maintain a new IRS storage environment, transitioning those assets to a private cloud-based SaaS model. Under this model, IRS pays for storage capacity as needed instead of purchasing storage equipment necessary to meet its peak demands. Using virtualization technology, Unisys mitigated	

Scope of Experience: 1 Contract Term: Current Segment: U.S. Federal Government Solution/Offering: Private Cloud Storage	U.S. Department of the Treasury, Internal Revenue Service (IRS) GSA's Alliant Government-wide Acquisition Contract
overcapacity of storage without the “stovepiped” effect and was able to manage storage as a single entity. We migrated more than 90 percent of relevant IRS data to the private storage cloud we operate for the IRS. The current contract is supported by an on-site Program Manager who started his career in the U.S. Navy. He has 25 years of technology experience in the public sector. There is also a Deputy Program Manager with 25 years of experience. Six storage engineers, a security manager, a business analyst, and an engagement manager are also on site. They collectively have about 10 years of industry experience. Two storage architects serve as subject matter experts who have more than 20 years of experience in designing large-scale, complex storage systems.	
Scope of Experience: 2 Contract Term: Current Segment: U.S. Federal Government Solution/Offering: Cloud Hosting and Migration Services	U.S. Department of the Interior (DOI) DOI Foundation Cloud Hosting Service Contract
On August 15, 2013, the U.S. Department of the Interior (DOI) tapped 10 companies to compete for a series of cloud computing projects that could involve a total investment of up to \$10 billion. DOI structured the deal among 10 systems integrators—Unisys, IBM, Verizon, AT&T, Lockheed Martin, Aquilent, Autonomic Resources, CGI, GTRI, and Smartronix—to compete for cloud migration; cloud management; cloud hosting; and a series of projects surrounding public, private, and hybrid cloud. Under this master contract, Unisys won and delivered the first task order to move DOI's Financial and Business Management System (FBMS) to the cloud, a move that provides significant flexibility while meeting DOI's stringent requirements for the high availability of data and applications. This move was announced as completed on January 5, 2016. The project enables DOI to dynamically reallocate cloud computing resources overnight and evaluate service level agreement performance alongside monthly invoices. Unisys worked with a partner, Virtustream, to provide and Infrastructure as Service (IaaS) solution, including Virtustream's SAP in the cloud hosting services. The Unisys Team collaborated closely with DOI to orchestrate a coordinated migration to the cloud and to manage the services once live. The successful migration made DOI the first Federal organization to move its SAP-based financial management application to the cloud. The project also complies with the criteria established in the Office of Management and Budget's Cloud First policy, which mandates that federal agencies take advantage of the cost savings and efficiencies of cloud computing.	

Scope of Experience: 3 Contract Term: Completed in 2015 Segment: U.S. Federal Government Solution/Offering: Public Cloud	General Services Administration (GSA) Enterprise Email and Collaborative Services Contract
In 2010, Unisys won the Enterprise Email and Collaborative Services Contract from the GSA to upgrade and migrate the agency's email systems from Lotus Notes and Domino to a public cloud-based email service provided by Google. This \$9.8 million, 5-year contract ended in 2015. In the first migration of a Federal agency's employees to a cloud-based email environment, Unisys supported the GSA's migration from its legacy Lotus Notes and Domino software to Google Apps for Government Software as a Service (SaaS) based on the FISMA-accredited Google Apps Premier platform. Google Apps for Government includes features that support collaboration, office productivity, and email delivered as SaaS over the Internet. Unisys recognized the need to support a growing GSA workforce; therefore, we proposed a customizable solution to accommodate increases in email volume and the number of user mailboxes. Unisys provides incremental email backup with flexible retention capabilities. Our solution also provides the GSA with the ability to add, change, and delete change mailbox accounts and directory data. With the migration completed in July 2011, Unisys supports 18,000 GSA employee mailboxes, each containing 30 GB of email storage.	
Scope of Experience: 4 Contract Term: Current Segment: U.S. State Government Solution/Offering: Public Cloud Hosting and Services (IaaS, PaaS, and SaaS)	State of Utah Western States Contracting Alliance (WSCA) / National Association of State Procurement Officials (NASPO) / ValuePoint Contract
In January 2013, Unisys was one of the four vendors awarded the WSCA Public Cloud Hosting and Services Contract. Unisys, in partnership with Amazon Web Services (AWS), delivers public cloud offerings from Amazon, Salesforce, and Google to the states participating in the WSCA/NASPO/ValuePoint contract. Unisys also provides assessment, design, and implementation services through our own workforce. In 2015, we booked an order totaling approximately \$1.5 million through this contract vehicle. We are providing the following list of sample Unisys clients under the current WSCA/NASPO/ValuePoint contract with a brief description of each solution offering. 1. State of Washington a) Department of Natural Resources: IaaS via AWS. SAP upgrade project. Amazon S3 Storage. b) Department of Natural Resources: IaaS & PaaS via AWS. Application DevOps infrastructure in Amazon.	

Scope of Experience: 4 Contract Term: Current Segment: U.S. State Government Solution/Offering: Public Cloud Hosting and Services (IaaS, PaaS, and SaaS)	State of Utah Western States Contracting Alliance (WSCA) / National Association of State Procurement Officials (NASPO) / ValuePoint Contract
c) Department of Ecology: IaaS via AWS. Application Development and Testing. d) Washington State Attorney General's Office: IaaS. e) Department of Health: IaaS via AWS. f) OCIO Office: SaaS via Salesforce. g) Department of Transportation: IaaS via AWS. h) Department of Social and Health Services: IaaS via AWS. 1. State of Wyoming a) Office of the Chief Information Officer (OCIO) Application Development and Testing in Amazon, AWS 2. State of Arizona a) Department of Transportation: IaaS via AWS. b) CIO Office: Cloud Broker via Pantheon. c) Arizona Strategic Enterprise Technology Office: IaaS via AWS d) Department of Health Services: IaaS via AWS 3. State of Hawaii a) Cloud Broker between State of Hawaii and Pacxa, one of the largest local IT firms in Hawaii. 4. State of Utah a) Department of Technology Services: SaaS. 5. State of Iowa a) State of Iowa: IaaS via AWS. b) Iowa Workforce Development: IaaS via AWS. 6. State of Missouri a) Intelligent Application Alignment (IAA) via Unisys: This is a Cloud Readiness assessment service that does discovery, assessment, and cloud portability of workloads in a given environment. 7. State of Montana a) State of Montana Department of Justice: Public Cloud Hosting via AWS. 8. State of Wisconsin	

Scope of Experience: 4 Contract Term: Current Segment: U.S. State Government Solution/Offering: Public Cloud Hosting and Services (IaaS, PaaS, and SaaS)	State of Utah Western States Contracting Alliance (WSCA) / National Association of State Procurement Officials (NASPO) / ValuePoint Contract
a) Department of Administration: Intelligent Application Alignment (IaaS) via Unisys. This is a Cloud Readiness assessment service that does discovery, assessment, and cloud portability of workloads in a given environment.	

Scope of Experience: 5 Contract Term: Current Sector: U.S. State Government Solution/Offering: Hybrid (Public/Private) Cloud and Data Center Consolidation	Commonwealth of Pennsylvania (CoPA) Data PowerHouse (DPH) / Pennsylvania Compute Services (PACS)
In 1999, Unisys was contracted to consolidate CoPA’s data center environments to one main data center (the Data PowerHouse) with an associated disaster recovery center. We were also enlisted to provide ongoing data center services, including server management; system engineering; systems operations, maintenance, and support; data center operations; data storage and archiving; information security engineering; application hosting; and disaster recovery. We also provide hardware and software provisioning; data center infrastructure maintenance; physical and logical security; and data security services that include data loss protection, malware protection, intrusion prevention and scanning. Our support model was based on ITIL V3 processes. In 2014, the term of this engagement was extended to 2022, with the scope to include the design, build, implementation, and hosting of a hybrid cloud environment for CoPA and its more than 50 Commonwealth wide agencies. This follow-on contract is called PACS (PA Compute Services). The contract terms for this engagement also were extended to 2022. The total contract value is \$1.61 billion. Every year, this contract produces \$100 million in revenue. On this engagement, Unisys provides the full breadth of cloud brokering solutions, including design, implementation, hosting, managed services, and technology integration. As part of the extension, Unisys deployed the ServiceNow platform for Service Management including governance, policy and process deployment and the full ITIL complement of service delivery. As part of the effort, Unisys built a robust service catalog, B2B interfaces and full hybrid cloud infrastructure management and operations. The benefits to the state includes increasing the operational efficiency by 35%, reductions in labor to support the infrastructure by allowing operations administration for all cloud models in one virtual location and reductions in operating costs. Throughout the term, Unisys implemented numerous projects that gained efficiencies worth \$240 million for CoPA, including mainframe and server consolidations, virtualization, storage area network (SAN) consolidations, and an option for variable service levels. We also implemented technologies such as network core switches, mass storage devices, mainframes with virtual hosts, SAN switches, and tape devices that reduced operating costs and improved functionality. Unisys continues to support CoPA’s data center needs and is helping CoPA on an agency wide data center modernization initiative that includes a new data center service delivery model that creates a flexible, reliable, secure, and robust IT structure. This model’s foundation is built on the concept of on-demand computing consumption for efficiency and cost savings as well as a broad service catalog.	

6.3 (M) Financials

Offeror must provide audited financial statements, of the last two years, to the State that demonstrate that an Offeror meets at a minimum Dun and Bradstreet (D&B) credit rating of 3A2 or better, or a recognized equivalent rating. Please provide the Respondent’s D&B Number and the composite credit rating. The State reserves the right to verify this information. If a branch or wholly owned subsidiary is bidding on this RFP, please provide the D&B Number and score for the parent company that will be financially responsible for performance of the agreement.

Unisys Response:

Unisys financial statements for the past 2 years are available at our web URL:
<http://www.unisys.com/investor-relations/financials-filings>.

Our D&B number is: 00-535-8932.

Unisys is providing the following as-recognized equivalent ratings:

- Moody's rating of B1
- S&P rating of B+.

6.4 (E) GENERAL INFORMATION

6.4.1 Provide any pertinent general information about the depth and breadth of your Solutions and their overall use and acceptance in the cloud marketplace.

Unisys Response:

Unisys and its partners Amazon, Microsoft and Google are the industry leaders in the cloud marketplace. Our combined strength and worldwide presence across all sectors of the public and the private industry is well known and accepted. We offer all types of cloud delivery models and service options. We are constantly expanding and enhancing our cloud models and delivery options. Section 6.2 showcases Unisys strong experience in the public sector cloud space. Additionally, the graphic 2015 Gartner Magic Quadrant for Cloud demonstrates the strength of our partners in this RFP.

Exhibit 6.4.1.1: 2015 Gartner Magic Quadrant for Cloud



6.4.2 Offeror must describe whether or not its auditing capabilities and reports are consistent with SAS 70 or later versions including, SSAE 16 6/2011, or greater.

Unisys Response:

Unisys is subject to client audit requirements according to SSAE 16, SOC 1, Type 2. Our existing staff have more than 10 years of experience, including the previous SAS 70 auditing guidelines, and work very closely with our external auditors at KPMG to create three SSAE 16 Type 2 reports covering at least 12 service centers. For auditing our client-facing IT environments, we follow ISO 27007 and ISO 27008. For auditing our financial aspects (related to SSAE 16 and otherwise), we follow GAAP under the AICPA's Code of Professional Ethics under Rule 203, Accounting Principles.

6.5 (E) BILLING AND PRICING PRACTICES

DO NOT INCLUDE YOUR PRICING CATALOG, as part of your response to this question.

6.5.1 Describe your billing and pricing practices, including how your billing practices are transparent and easy to understand for Purchasing Entity's.

Unisys Response:

Unisys follows a reliable and transparent approach with billing, pricing, and collections to enable our clients to run their businesses effectively. We will provide the Purchasing Entity with invoices for services provided in a timely way as and when the work is completed. Unisys is always open for support; billing-related queries can be forwarded directly to our representatives, who will work toward resolving issues in an appropriate method. We will send invoices to the Purchasing Entity at its designated address without fail. Unisys can also email a soft copy to a designated email address when required.

Unisys will provide access to a Consumption Reporting Solution that will enable the Purchasing Entities to track their cloud usage and charges. The Consumption Reporting Solution will provide a tabular view to monitor cloud usage in real time. A graphical view provides entities a comparison of their past usage with the current month that will help in forecasting their costs in their cloud consumption initiatives. The entities can download usage and billing reports in appropriate formats such as excel and pdf as and when required through this solution.

6.5.2 Identify any typical cost impacts that a Purchasing Entity might need to consider, if any, to implement your cloud solutions.

Unisys Response:

For the costs pertaining to billing, see our Cost Proposal.

6.5.3 Offeror must describe how its Solutions are NIST compliant, as defined in NIST Special Publication 800-145, with the service models it offers.

Unisys Response:

Unisys and our providers adhere to the NIST essential characteristics for Cloud Computing. For more information, refer to our response to Section 8.1.2.

6.6 (E) SCOPE AND VARIETY OF CLOUD SOLUTIONS

Specify the scope and variety of the Solutions you offer under this solicitation. You may provide a list of the different SaaS, IaaS, and/or PaaS services and deployment models that you offer.

Unisys Response:

Unisys is pleased to offer the requested cloud deployment models (private, community, public, and hybrid) and the underlying cloud services models (IaaS, PaaS, and SaaS) as part of our proposal. We will also provide architecture, design, implementation, and value-add services to support the various cloud solution options. Our cloud service provider partners in this proposal are Amazon Web Service (AWS), Microsoft, and Google. Exhibit 6.6.1 captures our offerings for this proposal.

Exhibit 6.6.1 Unisys Cloud Solution Offerings for NASPO.

	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Private Cloud	Unisys	Unisys	Unisys
Hybrid Cloud	Unisys	Unisys	Unisys
Cloud Professional Services	Unisys	Unisys	Unisys
Cloud Value-Add Services	Unisys	Unisys	Unisys

6.7 (E) BEST PRACTICES

Specify your policies and procedures in ensuring visibility, compliance, data security and threat protection for cloud-delivered services; include any implementations of encryption or tokenization to control access to sensitive data.

Unisys Response:

Section 8.6 and Section 8.8 of our proposal outlines specific details on compliance, data security and threat protection, including certifications and accreditations as they apply to this solicitation.

Unisys is pleased to provide the following, additional relevant examples of our additional past performance on access security and securing the cloud.

Brief Background Unisys Stealth

The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. Stealth can substitute traditional hardware topology for software-based cryptography. Stealth microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users (including database administrators).

- Unisys Stealth is an innovative, software-based security solution that does the following:
- Conceals end points, making them undetectable to unauthorized parties inside and outside the enterprise
- Tightens access control by focusing on user identity instead of physical devices so that security moves with the user and is easier to manage
- Protects sensitive data in motion from potential compromise through encryption
- Reduces costs by allowing clients to consolidate and virtualize networks, servers, and cloud architectures
- Supports regulatory compliance requirements (PCI) that draw resources away from a client's core business.

Unisys Stealth software has a legacy of being ready for government. Stealth received the National Information Assurance Partnership's (NIAP's) coveted Evaluation Assurance Level 4+ (EAL4+) certification, a Common Criteria international standard in effect since 1999. Unisys is pursuing the new standard of certification called Protection Profiles, as acknowledged by the U.S. Federal Government.

Relevant past performance of implementations of encryption or tokenization to control access to sensitive data:

State of New York: Unisys is implementing Stealth for Enterprise for the State of New York internal private cloud environment. Unisys has been engaged on the contract for the past 24 months, and the work is under way. We have helped New York to migrate and onboard about 35 of the state's 53 agencies to the multi-tenant data center in Albany running on Stealth.

Amazon: Unisys Stealth Cloud is an offering on Amazon Web Services (AWS) available directly at the AWS Marketplace. Stealth Cloud protects data in motion and isolates end-point workloads in a multitenant environment.

Commonwealth of Pennsylvania: as already referenced above, Unisys designed, integrated, and implemented several of Pennsylvania's own security and access control technologies into the Commonwealth's centralized data center.

5.0 ORGANIZATION PROFILE (RFP SECTION 7)

Section Title: Organization Profile: This section should constitute the Offeror's response to the items described in Section 7 of the RFP. An Offeror's response must be a specific point-by-point response, in the order listed, to each requirement in the Section 7 of the RFP.

7 ORGANIZATION AND STAFFING

7.1 (ME) Contract Manager

The Offeror must provide a Contract Manager as the single point of contact for management of the NASPO ValuePoint Master Agreement, administered by the State of Utah. The Contract Manager must have experience managing contracts for cloud solutions.

Harsh Bajpai, Unisys Director of Cloud Solutions and Services, is the current contract manager for the WSCA/NASPO/ValuePoint Cloud Hosting Contract with the State of Utah. If Unisys is awarded this contract, he will continue to manage it. For information on his qualifications, see Section 7.1.2.

7.1.1 Provide the name, phone number, email address, and work hours of the person who will act as Contract Manager if you are awarded a Master Agreement.

Mr. Harsh Bajpai

Unisys Corporation

Director, Cloud Solutions and Services

Phone Number: 703-439-6200

Email Address: Harsh.Bajpai@unisys.com

Work Hours: 8:00 a.m. to 8:00 p.m. Eastern Standard Time

7.1.2 Describe in detail the Contract Manager's experience managing contracts of similar size and scope to the one that will be awarded from this RFP. Provide a detailed resume for the Contract Manager.

Mr. Bajpai has supported the following government contracts for Unisys:

- WSCA/NASPO/ValuePoint: 2012 to present; expires in 2017
- SSA ITSSC: 2012 to 2014
- NASA SWEP IV&V: 2010 to 2015
- FEMA COMMIT: 2010 to 2011
- DHS EAGLE: 2009 to 2011.

Unisys provides the following resume for Mr. Bajpai.

Profile Summary:

Mr. Bajpai is an information technology professional with 15 years of experience in the following market segments:

- U.S. Federal and state-level public sector enterprises
- Global news and information service providers as well as media organizations
- Global investment banking.

Skills Summary:

Corporate Functions: Information Technology, Business Development, Sales, Sales Operations, Client Engagements and Implementations, and Marketing

Business: Transformation and Alignment, Strategy Formulation and Execution, Professional Services Management, Program Management, Outsourcing and Offshoring, Process Improvement and Automation, Governance, Practice Development, and Product and Services Management

Technology: Technology Infrastructure and Data Center Transformation and Optimization (DCTO), Virtualization, Cloud & Managed Services, Program Management, Information Security and Compliance, Software Development Lifecycle (SDLC), and End User Computing Platforms

Professional Experience:

08/2008 – Present	Director	Unisys Corporation
Market Space: USDHS (HQ, USCIS, USCG, FEMA, NPPD and CBP), GSA, U.S. Treasury, USPTO, USDA, SSA, FDIC, DHA (MHS & VA), and USSOCOM		
Role Summary - Develop and manage delivery of technology product solutions to the U.S. Federal markets. Support an annual revenue target of \$44 million. - Manage the Unisys portion of high-value services for existing strategic accounts. Provide SME-based technology consultation and integration services. Average annual and total task order value is \$120 million. - Provide SME-based technology consultation and integration services for U.S. Federal, state, local, and key global accounts. - Establish and maintain channel partnerships, vendor relationships, and client relationships across the GS, SES, and executive levels.		
Key Achievements - Key active member of the team that won the following contracts: COMMIT, ITSSC, TASPO, SWEP IV and 5, BASICS, Eagle 1, Networks, GSA EEAS, and WSCA/NASPO/ValuePoint. Post award of these contracts, program managed and executed task orders under these large federal contract vehicles. - Developed the Data Center Transformation and Outsourcing Practice (DCTO) for the Unisys Federal and Commercial markets. This framework today is the Unisys DCTO methodology. - Developed and delivered the Data Center Migration strategy for the Social Security Administration (SSA) under the ITSSC contract. - Developed and delivered a secure telework solution framework for the U.S. Coast Guard. This engagement realized the first reference implementation for the Unisys Stealth platform in the U.S. Federal market. - Developed and delivered the <i>Forward!</i> by Unisys solution to FDIC. This engagement realized the first reference implementation for the Unisys Stealth solution in the U.S. Federal market. - Initiated and won a \$30 million, 3-year contract with a large northeastern state for data center consolidation, modernization, security, and managed security services. - Won the USSOCOM medallion from GDIT for leading and representing Unisys on a very large IDIQ procurement. - Developed and delivered the first Unisys private cloud solution to HSBC Bank in Hong Kong. This solution realized the first reference implementation for the Unisys Private Cloud offering. - Strategic team member behind the \$100 million SITA win for AirCore application and baggage claims systems. - As the manager of the WSCA/NASPO/ValuePoint contract, generated \$5 million in Unisys sales revenue under the public cloud hosting and services vertical in the State and Local		

08/2008 – Present	Director	Unisys Corporation
market segment.		
04/2006 – 08/2008	Senior Director, VP	Thomson Reuters
Position Summary & Functional Responsibilities • Management of the global infrastructure, application support, and I-Dev teams • End-to-end responsibility for the primary and secondary data centers of Reuters Media • Outsourcing and offshoring • Global infrastructure availability • Infrastructure platform and service delivery as well as project management and compliance functions for Reuters Media.		
Key Achievements • Outsourced the Reuters Media primary data center to a managed service provider on a utility computing platform. • Program managed the migration of the Reuters Media platform from .NET to the LAMJ stack. • Outsourced the QA and development infrastructure to a hosting provider. • Built the software development facility of Reuters Media in Beijing, China. • Outsourced the software development production support function of Reuters Media to Ukraine. • Built the globally load-balanced, live/live DR/COOP facility for Reuters Media in the United Kingdom with a managed service provider.		
01/2001 – 03/2006	Various roles – see below	Dow Jones & Co. Inc.
Position Summary & Functional Responsibilities 2004 – 2006: E-Commerce Product Manager for WSJ.com and Barrons.com 2003 – 2004: Business Information Security and Compliance Officer 2001 – 2003: Information Security Architect		
Key Achievements E-Commerce Product Manager for WSJ.com and Barrons.com • Delivered an e-commerce platform for the \$10 million WSJ.com and LexisNexis partnership. • Delivered the e-commerce platform for WSJ.com and <i>The Wall Street Journal</i> combined product offering in Asia and Europe. • Owned the WSJ.com and Barrons.com e-commerce platform. Business Information Security and Compliance Officer • Was accountable for the Sarbanes-Oxley program for all external business units of Dow Jones & Co. • As a technology SME for the core M&A team of Dow Jones & Co., executed M&A activities for ONI, Factiva.com, and VentureWire. Information Security Architect • Oversaw information security and assurance of Factiva.com, a Reuters and Dow Jones		

01/2001 – 03/2006	Various roles – see below	Dow Jones & Co. Inc.
Company. - Delivered forensic analysis for the legal department and external law enforcement agencies. - Introduced and deployed Open Source production-class tools such as intrusion detection, vulnerability analysis, and configuration management.		

10/1998 – 01/2001	Technical Architect	SnS Medical Systems
- Assumed end-to-end architectural and operational ownership of the hosted software billing system. - Assumed service delivery ownership of client-owned and -managed systems. - Owned the software service delivery and security for the enterprise. - Managed deployment and field service consultants.		

08/1995 – 10/1998	General Manager, Operations	Accor Inc.
- Canvassed and presented quarterly competitive market analysis reports for the regional district and for the 15 properties in the district. - Successfully turned around operationally challenged and P&L-stagnant locations across NJ, PA, and CT. - Designed and Implemented operational metrics and reporting for the regional district. - Oversaw annual revenue of \$2 million.		

Education:

MBA, Business Administration, University of Maryland University College, Adelphi, MD, 2005

M.S., e-Commerce, University of Maryland University College, Adelphi, MD, 2003

B.A., Marketing & Finance, California State University, Fullerton, Fullerton, CA, 1995

A.A., Social & Behavioral Sciences, Citrus College, Glendora, CA, 1993

A.S., Business Administration, Citrus College, Glendora, CA, 1993

Certifications, Affiliations, and other Professional Work Experience

CISSP, MCSE, MSCE+I, and MCT

Middlesex County Court, Juvenile Conference Committee Member, 2004 – 2006

Middlesex County College, Adjunct Professor, Department of Computer Science, 2004 – 2006

Princeton University, Guest Lecturer, 03/2013

Baltimore National Aquarium, Volunteer Diver, 06/2013 – Present

7.1.3 Describe in detail the roles and responsibilities of the Contract Manager as they apply to the NASPO ValuePoint Master Agreement that will be awarded from this RFP.

Harsh Bajpai is our current contract manager for the WSCA/NASPO/ValuePoint Cloud Hosting Contract with the State of Utah. If Unisys is awarded this contract, he will continue to manage it and serve as the single point of contact to Utah and the Participating States.

Mr. Bajpai will continue to facilitate onboarding of new customers and respond to questions and concerns about new and existing procurements. He also will continue to coordinate customer enablement through our partners. He will provide the collaboration and support required from our partners in this proposal.

The following functions are required to support a contract of this scope and size. The current WSCA/NASPO/ValuePoint contract has talented individuals who support the functions stated below. If Unisys is awarded this contract, we are committed to provide this support structure.

Contracts Desk: A team of talented individuals supports the day-to-day operations of current and new contracts. The Contract Operations Specialist supports Mr. Bajpai in writing statements of work (SOWs); resolving billing-related issues; managing the calculation and disbursement of fees, credits, etc.; and overseeing terms and conditions as required by this solicitation. If Unisys is awarded this contract, we will assign a dedicated Contract Operations Specialist to it.

Attorney: If Unisys is awarded this contract, we will dedicate an attorney to support Mr. Bajpai on required legal matters.

Enablement, Billing, and Reporting: If Unisys is awarded this contract, we will dedicate a team of billing and reporting staff to Mr. Bajpai. This team will provide customer enablement with our various partners on this contract. This team will also provide the billing details and reports required to manage and execute this contract.

Certified Architects: Unisys has geographically positioned certified architects in and near the states that are our current clients. These architects can deliver full end-to-end solutions. They also engage frequently with our clients as consultants and advisors. If Unisys is awarded this contract, we will continue to support and expand this function. These architects will support Mr. Bajpai.

Client Executives: Unisys has geographically positioned Client Executives in and near the states that are our current clients. Our Client Executives are account managers dedicated to a given state. A Client Executive's main function is to provide our clients with onsite presence and support. If Unisys is awarded this contract, we will continue to support and expand this function. These Client Executives will support Mr. Bajpai.

Subject Matter Experts: Unisys maintains a very strong partnership with our partners in this proposal. Formal partnerships require us to maintain a certain number of certified personnel to fulfill this partnership's terms and conditions. Unisys will continue to strengthen our bench based on the required demands. Additionally, our partnerships give us direct access to technical and nontechnical resources in our partner organizations. Unisys bench staff and our partners' resources support Mr. Bajpai under the current WSCA/NASPO/ValuePoint contract. If Unisys is awarded this contract, we will continue to support and expand this resource base to support Mr. Bajpai.

6.0 TECHNICAL RESPONSE – AWS – MS INCLUDED (RFP SECTION 8)

Section Title: Technical Response. This section should constitute the Technical response of the proposal and must contain at least the following information:

A complete narrative of the Offeror's assessment of the Cloud Solutions to be provided, the Offerors ability and approach, and the resources necessary to fulfill the requirements. This should demonstrate the Offeror's understanding of the desired overall performance expectations and clearly indicate any options or alternatives proposed.

A specific point-by-point response, in the order listed, to each requirement in the Section 8 of the RFP. Offerors should not provide links to a website as part of its response.

Offeror's should focus their proposals on the technical qualifications and capabilities described in the RFP. Offerors should not include sales brochures as part of their response.

8 TECHNICAL REQUIREMENTS

If applicable to an Offeror's Solution, an Offeror must provide a point by point responses to each technical requirement demonstrating its technical capabilities. If a technical requirement is not applicable to an Offeror's Solution then the Offeror must explain why the technical requirement is not applicable.

If an Offeror's proposal contains more than one Solution (i.e., SaaS and PaaS) then the Offeror must provide a response for each Solution. However, Offerors do not need to submit a proposal for each Solution.

8.1 (M)(E) TECHNICAL REQUIREMENTS

8.1.1 Offeror must identify the cloud service model(s) and deployment model(s) it intends to provide to Eligible Users. See Attachment D.

Unisys Response:

Unisys has the ability and the qualifications to provide the four types of cloud deployment methods (Public, Private, Community, and Hybrid) defined in Attachment D. Unisys also has the ability and the qualifications to provide the three types of cloud service models (IaaS, PaaS, and SaaS) defined in Attachment D.

To provide this wide and deep cloud offering catalog, Unisys partnered with Amazon Web Services (AWS), Microsoft, and Google for this proposal. Unisys believes that offering a selection of world-class, best-of-breed cloud providers will provide the Participating Entities with the needed flexibility, agility, and technological options in a competitive environment. We also believe that offering these solutions and services from a cloud-neutral systems integrator like Unisys minimizes risk, promotes accountability, and improves efficiency. With this approach, the Participating Entities can choose their cloud solutions from a competitive set of vendors that are backed by a single, experienced cloud-neutral systems integrator. Exhibit 1 summarizes the Unisys offering for this proposal.

Exhibit 1: Unisys Cloud Solutions Offerings to NASPO.

	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Private Cloud	Unisys	Unisys	Unisys
Hybrid Cloud	Unisys	Unisys	Unisys
Cloud Professional Services	Unisys	Unisys	Unisys
Cloud Value Add Services	Unisys	Unisys	Unisys

Our partners in this proposal provide qualified offerings in the public and the community (government-specific) cloud deployment space. Our partners also provide the three cloud service (IaaS, PaaS, and SaaS) models for the U.S. public sector. Unisys is a recognized leader in providing hybrid and private cloud deployment methods in the public sector. We continue to build more private and hybrid cloud environments for federal, state, and local governments. Unisys is also a recognized leader in SaaS-based applications services platforms. Our main presence is in the Health and Human Services, Public Safety and Justice, Law Enforcement, Border Security, Digital Government, and Citizen Services markets.

As the incumbent, Unisys is fully aware of what is needed to fulfill the requirements and the overall performance expectations necessary for this contract vehicle. We believe that our end to end cloud expertise and experience along with strong partnerships will provide unique capabilities for our government clients.

8.1.2 For the purposes of the RFP, meeting the NIST essential characteristics is a primary concern. As such, describe how your proposed solution(s) meet the following characteristics, as defined in NIST Special Publication 800-145:

Refer to the point-by-point responses below.

8.1.2.1 NIST Characteristic - On-Demand Self-Service: Provide a brief written description of how the cloud solutions proposed satisfies this individual essential NIST Characteristic. Attest capability and briefly describe how self-service technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of On-demand Self Service, one of the 5 essential characteristics of the NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Examples of these technologies include those from vendors such as BMC, Microsoft, VMware, HP, ServiceNow, Open Source Foundation, IBM, and Dell. Unisys has leveraged combinations of these technologies to create customized, automated self-service portals that can manage a workload's entire life cycle in an internal, private cloud environment or in a hybrid environment. Unisys has relevant past performance in this space in the U.S. Federal, state, and local markets. We provide some of these past performance details in Section 4.0, Business Profile.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Amazon Web Services (AWS) provides customers with on-demand access to a wide range of cloud infrastructure services, charging them only for the resources they actually use. AWS will enable the State to eliminate the need for costly hardware and the administrative pain that goes along with owning and operating it. Instead of the weeks and months it takes to plan, budget, procure, set up, deploy, operate, and hire for a new project, AWS customers can simply sign up for AWS and immediately begin deployment in the cloud with the equivalent of 1, 10, 100, or 1,000 servers. Whether an organization needs to prototype an application or host a production solution, AWS makes it simple for customers to get started and be productive.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Windows Azure is an internet-scale, high-availability cloud fabric operating on globally-distributed Microsoft data centers. Windows Azure and related tools support the development and deployment of applications into a hosted environment that extends the on-premises data center. On-demand self-service refers to the service provided by Microsoft Azure that enables the provision of resources on demand whenever required. In on-demand services can be enabled from the HTML Portal, using Azure API, using CLI for Mac, Linux, and Windows with Azure Service Management. Azure on-demand self-service resource sourcing programmatically is a prime feature allowing the user to scale the infrastructure. Microsoft provides a full-featured portal that enables customers to order the hundreds of cloud services available from Azure. Azure can scale from one to thousands of virtual machine instances with the benefit of built-in virtual networking and load balancing. Microsoft bills customers per minute, allowing them to pay only for what they use.

Google Public and Community Cloud for SaaS: Google Apps administrators have access to an Admin Panel that allows them to perform user administration, service configurations, reporting on demand. Additional tools that support this characteristic include the APIs in the Admin SDK that can automate many common tasks as well as an extensive set of third-party tools developed by trusted independent software vendors.

8.1.2.2 NIST Characteristic - Broad Network Access: Provide a brief written description of how the cloud solutions proposed satisfies this individual essential NIST Characteristic. Attest capability and briefly describe how network access is provided.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Broad Network Access, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Vendor-based examples of technologies include those of vendors such as Cisco, Juniper, Nortel, Foundry, and F5. Technology-based examples include leased lines, Internet, VPN (SSL and IPsec), web applications, thick clients, thin clients, and mobile clients. Using these technologies, Unisys has provided our Federal, state, and local clients with low-latency, redundant broad network access methods.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS provides a simple way to access servers, storage, databases, and a broad set of application services over the Internet. Cloud computing providers such as AWS own and maintain the network-connected hardware required for these application services; customers provision and use what they need from a web application, a mobile client, or programmatically at published and well-documented APIs.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Azure allows customers to access their Web, Mobile, IaaS, SaaS, and PaaS platforms easily. Network access to the cloud is available over VPN or dedicated high-speed connections between the customer and the Microsoft cloud network. Connectivity between Microsoft data centers is delivered by one of the largest fiber networks in the world, establishing high levels of performance and security. Microsoft Azure enables the enterprise to create an on-premises network route from on-premises VPN device and the Azure virtual network. Configure on-premises hardware or software VPN device to terminate the VPN tunnel, which uses Internet Protocol security (IPsec). Gateway connects on-premises to Microsoft Azure through many connection bandwidth using software only or hardware based connectivity. Basic VPN connection, Standard VPN, ExpressRoute gateway connections.

Google Public and Community Cloud SaaS: Google's network delivers speeds of 1 petabyte per second of total bidirectional bandwidth, has dozens of network points of presence, and offers a clear and mature peering/content delivery policy so that customers are connected to Google's infrastructure, quickly negating the need for high-cost dedicated access lines. Google Apps (SaaS) is device, platform, and network neutral. As a pure-play cloud-based solution, the services can be accessed from modern browsers from a network access point. Customers that have requirements to limit access to trusted devices or trusted access points have access to Device Management tools from the Admin Panel at no additional cost.

8.1.2.3 NIST Characteristic - Resource Pooling: Provide a brief written description of how the cloud solutions proposed satisfies this individual essential NIST Characteristic. Attest capability and briefly describe how resource pooling technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Resource Pooling, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Technology-based examples include Unisys Stealth, VMware, Hyper-V, and Xen-based virtualization; Unisys ClearPath Forward, a hardware partitioning

utility compute platform; storage systems from EMC and NetApp; etc. Using these technologies, we created shared services environments for various Federal and state entities. Examples include the Internal Revenue Service; the Commonwealth of Pennsylvania; and the States of New York, California, and Michigan. Section 4.0, Business Profile contains the solution sets that Unisys has created.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: The AWS environment is a virtualized, multitenant environment. AWS has implemented security management processes, PCI controls, and other security controls designed to isolate each customer from other customers. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture has been validated by an independent PCI Qualified Security Assessor (QSA) and was found to comply with the requirements of PCI DSS version 2.0, published in October 2010.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Azure resource pooling supports scalable systems involved in cloud computing and software as a service (SaaS) and enable near-infinite growth with immediate availability. The kinds of services that can apply to a resource pooling strategy include data storage services, processing services, bandwidth provided services, and other Azure-related compute elasticity.

Google Public and Community Cloud for SaaS: Google's global infrastructure is a shared pool of resources that dynamically serve each end user with a primary data center access point that may rotate throughout the session without the end user being aware and at the same time replicate data across at least two additional geographically dispersed data centers that also may rotate through the session.

8.1.2.4 NIST Characteristic - *Rapid Elasticity*: Provide a brief written description of how the cloud solutions proposed satisfies this NIST Characteristic. Attest capability and briefly describe how rapid elasticity technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Rapid Elasticity, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Examples of these technologies include those from vendors such as Microsoft, Amazon, VMware, Dell, Chef and Puppet. Unisys has leveraged combinations of these technologies to create customized deployment configurations with pre-defined business rules enabling auto-scaling that can seamlessly scale resources up or down by provisioning and de-provisioning resources automatically as the workloads change. The resources shall in turn be handled by a load balancer to distribute the workloads to different resources in an internal, private cloud environment or in a hybrid environment.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS provides a massive global cloud infrastructure that will allow the State to quickly innovate, experiment, and iterate. Instead of waiting weeks or months for hardware, the State can instantly deploy new applications, instantly scale up as its workload grows, and instantly scale down according to demand. Customers need to be confident that their existing infrastructure can handle a spike in traffic and that the spike will not interfere with normal business operations. Elastic Load Balancing and Auto Scaling can automatically scale a customer's AWS resources up to meet unexpected demand and scale those resources down as demand decreases.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Azure supports rapid elasticity allowing automated requests for additional resources (i.e. compute, disk space, connectivity and other types of services). The State can automatically scale resources according to a variety of metrics. The ability to scale resources up or down, combined with per minute billing, makes Azure the most efficient computing platform on the market. The automatic scaling feature at the Azure portal is very intuitive and

easy to use. This feature is used by thousands of customers needing the ability to turn compute resources up or down automatically.

Google Public and Community Cloud for SaaS: Google Apps (SaaS) can scale from one user to tens of thousands of users. The use of the services included is designed to scale horizontally in the per user storage allocation. Google Apps can be purchased with an allocation of 30 GB per user or with unlimited storage.

8.1.2.5 NIST Characteristic - Measured Service: Provide a brief written description of how the cloud solutions proposed satisfies this NIST Characteristic. Attest capability and briefly describe how measured service technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Measured Service, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Examples include VantagePoint, a solution offered by Unisys that we implemented at the Commonwealth of Pennsylvania. For the IRS, we provide a customized private storage environment that is based on a measurable SLA down to the data types, which maps to storage offerings categorized into Platinum, Gold, Silver, and Bronze Storage.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS uses automated monitoring systems to provide a high service performance and availability. Proactive monitoring is available from a variety of online tools for internal and external use. Systems in AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so that personnel are always available to respond to operational issues. This schedule includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Azure supports NIST's principal areas of measured services such as measured service setup, which allows Azure to control a system, user, or tenant's usage of resources with metering capability. Azure supports automated remote services measurement tools to provide auditing and accountability for utilization. Azure-measured service confirms that even when there is no specific interaction for a service change, service change is still audited to support billing cycles.

Google Public and Community Cloud for SaaS: Google Apps (SaaS) is already optimized for unlimited use by end users. Google maintains high availability, low latency, and fault tolerance as a part of its contracted services. No metering or rate limiting will be required because Google does not charge for bandwidth use

8.1.3 Offeror must identify for each Solution the subcategories that it offers for each service model. For example if an Offeror provides a SaaS offering then it should be divided into education SaaS offerings, e-procurement SaaS offerings, information SaaS offering, etc.

Unisys Private and Hybrid Cloud for IaaS and PaaS: These are custom solutions that Unisys can design, implement, and manage for clients according to their needs. As a systems integrator, Unisys has integrated many hardware and software solution sets to meet the requirements of our clients and NIST. Our custom examples include private storage as a service for the IRS, hybrid storage as a service for the Commonwealth of Pennsylvania, and private cloud for the U.S. Department of the Interior.


Unisys Public and Community, Private, and Hybrid Cloud for SaaS: Unisys is a leader in providing SaaS for the Health and Human Services, Public Safety and Justice, Law Enforcement, Border Security, Digital Government, and Citizen Services markets. We provide custom integrated applications that can be hosted in the public and community cloud environments as an IaaS or a PaaS; or these solutions can be hosted in a hybrid or a private cloud. A recent example includes 311 Service for the City of Philadelphia. The 311 application is a Unisys application that is hosted on Salesforce.com. Unisys customized and

integrated the application for the City of Philadelphia and delivered it on SalesForce.com on a consumption basis. **Appendix 11 – Unisys-1** contains the definition summary of each cloud offering provided by Unisys.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Exhibit 2 illustrates each service models offered by Amazon in the public and community cloud deployment methods. In the point-by-point responses, below as required by this RFP, we have defined each service that is illustrated in Exhibit 2. **Appendix 12 – AWS-1** contains the definition summary of each one of the services in Exhibit 2 in a consolidated table.

Exhibit 2.AWS Offerings.

Aw



UT003

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

IaaS is most like traditional IT delivery. Customers provision their own virtual machines, define their own networks, and allocate their own virtual hard disks. IaaS shifts the burden of operating datacenters, virtualization hosts, and hypervisors. In addition, the business continuity and disaster recovery infrastructure is shifted from the enterprise to the service provider.

Deploying an application and managing an IaaS environment provides the most flexibility that Azure has to offer. With any deployment choice, there will be pros and cons that must be considered. The greatest benefit of an IaaS implementation is that it offers the greatest amount of control from the operating system to manage access to the application.

IaaS Services by Microsoft includes:

Systems Center Virtual Machine Manager: System Center Virtual Machine Manager provides centralized administration and management of a virtual environment.

Systems Center Operations Manager: System Center Operations Manager 2007 provides end-to-end monitoring for the enterprise IT environment.

Systems Center Configuration Manager: Systems Center Configuration Manager 2007 provides a comprehensive solution for change and configuration management for the Microsoft platform.

PaaS extends IaaS further by providing multitenant services that customers subscribe to. Platform services are a transformational computing model that can dramatically reduce the costs and increase the agility of delivering applications to end users internally and externally. PaaS users bring their own application code but leverage robust platforms, which they do not need to maintain.

With PaaS applications, many of the layers of management are removed and more flexibility is provided than an application running on IaaS instances. Specifically, there is no need to manage the operating system, including patching, which reduces some of the complexity of designing the deployment.

A significant benefit of deploying an application running in a PaaS environment is the ability to quickly and automatically scale up the application to meet the demand when traffic is high, and inversely scale down when the demand is less. Deploying an application in the PaaS model is very cost effective from a scalability and manageability perspective.

PaaS Services by Microsoft includes:

Microsoft SQL Azure Database: a fully relational cloud database solution built on SQL Server that offers highly available, scalable, multi-tenant database services.

Windows Azure: delivers on-demand compute and storage to host, scale, and manage web applications through Microsoft data centers.

SaaS is the real promise of cloud computing. By integrating applications from one or multiple vendors, customers need to bring only their data and configurations. They can eliminate the costs of building and maintaining applications and platform services and still deliver the secure, robust solutions to the end users.

Choosing an Azure SaaS offering provides the least amount of responsibility on the customer's side. At the same time, providing a lesser amount of flexibility in comparison with an IaaS or PaaS approach.

SaaS Services by Microsoft includes:

Office 365: a group of software plus services subscriptions that provides productivity software and related services to its subscribers. Office 365 offers plans that include e-mail and social networking services through hosted versions of Exchange Server, Skype for Business Server, SharePoint, and Office Online integration with Yammer, as well as access to Office software.

Microsoft Dynamics: with minimal configuration, offers constituent relationship management (CRM) and other extended CRM solutions to help automate workflow and centralize information.

Microsoft Exchange Online: delivers email with protection, plus calendar and contacts.


Microsoft Live Meeting: delivers hosted web conferencing.

Microsoft SharePoint: creates a highly secure, central location for collaboration, content, and workflow.

Microsoft Communications Online: provides real-time, person-to-person communication with text, voice, and video across agencies.


Many scenarios need to implement a blend of Azure offerings to meet the needs of their organization and application requirements. The following diagram (Exhibit 3) highlights the main differences from a manageability perspective, when using public cloud SaaS, PaaS, IaaS and On-Premises implementations

Exhibit 3: Microsoft IaaS/PaaS/SaaS differentiation



Google Public and Community Cloud for SaaS: Exhibit 4 and Exhibit 5 illustrate each service model offered Google in the public and community cloud deployment models. **Appendix 13 – Google-1** contains the definition summary of each services in these exhibits in a consolidated table.

Exhibit 4. Google SaaS.



8.1.4 As applicable to an Offeror's proposal, Offeror must describe its willingness to comply with, the requirements of Attachments C & D.

Unisys is willing to comply with the 'NIST Service Models' and 'Scope of Services'.

8.1.5 As applicable to an Offeror's proposal, Offeror must describe how its offerings adhere to the services, definitions, and deployment models identified in the Scope of Services, in Attachment D.

Unisys and our partners for this proposal have a proven track record in providing, building, deploying, and managing public, community, private, and hybrid cloud environments that deliver IaaS, PaaS, and SaaS services. We independently and jointly delivered and built these environments for U.S. Federal, state, and local governments that also follow the requirements defined in the State provided Attachment D. Our collective and individual past performances are stated in Section 4.0, Business Profile.

Additionally, our certifications and accreditations that are listed throughout this proposal attest to our understanding and our adherence to the services, definitions, and deployment models identified in the State provided Attachment D.

8.2 (E) SUBCONTRACTORS

8.2.1 Offerors must explain whether they intend to provide all cloud solutions directly or through the use of Subcontractors. Higher points may be earned by providing all services directly or by providing details of highly qualified Subcontractors; lower scores may be earned for failure to provide detailed plans for providing services or failure to provide detail regarding specific Subcontractors. Any Subcontractor that an Offeror chooses to use in fulfilling the requirements of the RFP must also meet all Administrative, Business and Technical Requirements of the RFP, as applicable to the Solutions provided. Subcontractors do not need to comply with Section 6.3.

8.2.2 Offeror must describe the extent to which it intends to use subcontractors to perform contract requirements. Include each position providing service and provide a detailed description of how the subcontractors are anticipated to be involved under the Master Agreement.

Unisys does not intend to use subcontractors. We intend to use our highly qualified, world-recognized, strategic, and strong industry partners (Amazon, Microsoft, and Google) to provide some of the cloud solutions detailed in our proposal. Unisys will resell our partners' cloud models and services using our agreements with our partners for the delivery of IaaS, PaaS, and SaaS from public and community clouds.

Unisys will directly provide private and hybrid cloud models that will be capable of delivering IaaS, PaaS, and SaaS. Additionally, we will offer free professional services and value-add services.

The Master Agreement will be an agreement between the State and Unisys.

8.2.3 If the subcontractor is known, provide the qualifications of the subcontractor to provide the services; if not, describe how you will guarantee selection of a subcontractor that meets the experience requirements of the RFP. Include a description of how the Offeror will ensure that all subcontractors and their employees will meet all Statement of Work requirements.

Not applicable based on the response above. Unisys is not using subcontractors. We are using our partners. Unisys has partnership agreements in place with the partners in this proposal.

8.3 (E) WORKING WITH PURCHASING ENTITIES

8.3.1 Offeror must describe how it will work with Purchasing Entities before, during, and after a Data Breach, as defined in the Attachments and Exhibits. Include information such as:

- **Personnel who will be involved at various stages, include detail on how the Contract Manager in Section 7 will be involved;**
- **Response times;**
- **Processes and timelines;**
- **Methods of communication and assistance; and**
- **Other information vital to understanding the service you provide.**

Unisys follows a Security Incident Response Process that details the routine handling of intentional or inadvertent information security events affecting the integrity, confidentiality, authentication, non-repudiation, and availability of information, and the information technology infrastructure of Unisys. The Unisys representative will be our Contract Manager, Harsh Bajpai. He will be supported by a core incident response team that includes representatives from our IT/IS, legal, communications, program management, and partner organizations as well as subject matter experts from various areas of technology as needed. This team from Unisys and our partners will closely coordinate with the State's incident response teams to contain, assess, mitigate, and prevent incidents. Once engaged for a confirmed breach, the members of our core incident response team will remain with the incident until it is resolved to the State's satisfaction.

The process of identification will be a joint responsibility between the State as well as Unisys and our partners because incidents can be detected by the parties involved. If the State detects an incident, the State can contact our Contract Manager using the information in Section 5.0, Organization Profile. If Unisys or our partners detect an incident, we will contact the State. Our Contract Manager will assemble the required incident response team, establish an internal and an external communications plan, and summon an investigation team of technology, legal, public relations, and communications experts. Unisys will work with the State to provide required response times based on the incident's severity. We will cooperate with the involved parties from Federal, State, and local authorities under the State's guidance and Federal law.

8.3.2 Offeror must describe how it will not engage in nor permit its agents to push adware, software, or marketing not explicitly authorized by the Participating Entity or the Master Agreement.

Unisys does not push adware, software, or marketing that is not explicitly authorized by the Participating Entity or the Master Agreement.

8.3.3 Offeror must describe whether its application-hosting environments support a user test/staging environment that is identical to production.

In the public and the community IaaS cloud space for application hosting, clients create, manage, and develop solutions in their own environment. These environments are created in accordance with their functional requirements and specifications. Vendors provide these environments that are bound by SLAs. Vendors will provide test and development environments for PaaS and SaaS delivery models. For those environments, the underlying infrastructure, features, functionality, and specifications are also bound by the governing SLAs. In the hybrid and private cloud models, IaaS, PaaS, and SaaS environments are built to client requirements and the governing SLAs.

8.3.4 Offeror must describe whether or not its computer applications and Web sites are be accessible to people with disabilities, and must comply with Participating entity accessibility policies and the Americans with Disability Act, as applicable.

The solutions in this proposal are accessible to people with disabilities and complies with Participating Entity accessibility policies and the Americans with Disabilities Act.

8.3.5 Offeror must describe whether or not its applications and content delivered through Web browsers are be accessible using current released versions of multiple browser platforms (such as Internet Explorer, Firefox, Chrome, and Safari) at minimum.

Unisys and our partners support current versions of Internet Explorer, Firefox, Safari, and Google Chrome. Unisys and our partners constantly track and move with the industry when it comes to browser compatibilities and browser technologies in the marketplace. This also includes browsers that run on computers, laptops, tablets, mobile devices, etc.

8.3.6 Offeror must describe how it will, prior to the execution of a Service Level Agreement, meet with the Purchasing Entity and cooperate and hold a meeting to determine whether any sensitive or personal information will be stored or used by the Offeror that is subject to any law, rule or regulation providing for specific compliance obligations.

As an incumbent that has fulfilled the requirements for nine Participating Entities under the current contract, Unisys takes our interaction with the Purchasing Entities seriously and continuously strives to enhance this experience. A purchasing entity usually contacts Unisys directly or through a third party. Upon mutual determination, our Contract Manager, Harsh Bajpai, and a Unisys Client Executive conduct an initial client call with the Purchasing Entity. This call focuses mainly on technical and functional requirements as well as regulatory compliance requirements and obligations. If a follow-up, deep-dive call is required, Unisys provides the appropriate certified subject matter experts (SMEs). At this stage, the Purchasing Entity brings in its SMEs. If further discussions are required, Unisys has designated staff, as described in Section 5.0, Organization Profile. Unisys and our SME resources understand the workload and the governing requirements and can suggest an appropriate cloud destination to the Participating Entity.

8.3.7 Offeror must describe any project schedule plans or work plans that Offerors use in implementing their Solutions with customers. Offerors should include timelines for developing, testing, and implementing Solutions for customers.

Public and Community Cloud Offerings for IaaS, PaaS, and SaaS from Amazon, Microsoft and Google: Unisys will provide a Statement of Work (SOW) within 5 business days of receiving the firm intent to purchase from the customer. Upon receiving the signed copy of the SOW, we will provision the account for the customer and provide login details within 5 business days.

Private and Hybrid Cloud Offerings for IaaS, PaaS, and SaaS from Unisys: Because these are customized solutions of varying sizes and requirements, Unisys is providing example transition, design, build, and migration times for our Federal and State government clients. The example client list for reference includes USDA, IRS, and Commonwealth of Pennsylvania. They each have a footprint of 1,000 or more servers. Our typical transition time for their environments is 30, 60, or 90 days. We can design and build a new environment in 3 to 6 months. Workload migration depends on the participation and the prioritization from the client, workload complexity, mission criticality, current redundancies, and migration provisions in place. Therefore, it is difficult to estimate these migrations. Although Unisys has

migrated some of our smaller state-level clients in 3 to 6 months, some of our engagements with larger Federal and state-level clients have lasted up to 24 months.

8.4 (E) CUSTOMER SERVICE

8.4.1 Offeror must describe how it ensure excellent customer service is provided to Purchasing Entities. Include:

- **Quality assurance measures;**
- **Escalation plan for addressing problems and/or complaints; and**
- **Service Level Agreement (SLA).**

Problem Reporting and Escalation

Unisys Technology and Support Services are available from Unisys to our clients 24x365. Our Technology and Support Services organization responds to a client's request for support. Usually, the client calls a central toll-free number or enters a service incident at our website. Before contacting Unisys, the client is required to consolidate and vet incidents and then document the suspected fault, specifying its severity and the impact the problem causes. This step is known as Level 1 Support, in which Unisys verifies the client's entitlement to support.

If a call or a website contact cannot be directly answered and resolved, our Level 1 Support team triages it and sends it to the appropriate Level Support 2 team that knows the specific target application. This team can be an engineering support team at Unisys or a third-party technology vendor with which Unisys partners for this solution. Our Level 2 Support team analyzes the incident and rejects problems not directly related to a fault in the application. Our Level 2 Support team may need to work with our Level 1 Support team to call the client several times to seek information to help isolate the problem or to develop a fix.

Once a fault is isolated, our Level 2 Support team submits a fault report to our development team (Unisys or the third-party technology partner). This is known as Support Level 3, which is responsible for the final fix. Our Level 3 development team then fixes and tests the fault. The fix is then made available to the client. To escalate an issue to the Unisys Support system at the same toll-free number, the client can ask to speak to a Duty Manager.

The SLAs are available in **Appendices 8 – 10.**

Appendix 8 – AWS SLA

See Appendix 8

Appendix 9 – Microsoft SLA

See Appendix 9

Appendix 10 – Google Apps SLA

See Appendix 10

8.4.2 Offeror must describe its ability to comply with the following customer service requirements:

a. You must have one lead representative for each entity that executes a Participating Addendum. Contact information shall be kept current.

If Unisys is awarded the contract, we will provide one lead representative for each entity that executes a Participating Addendum and keep the contact information current.

b. Customer Service Representative(s) must be available by phone or email at a minimum, from 7AM to 6PM on Monday through Sunday for the applicable time zones.

Unisys customer service and the customer service of our partners in this proposal operate a help desk 24x365.

c. Customer Service Representative will respond to inquiries within one business day.

The Unisys contract manager identified in Section 5.0, Organization Profile will respond to queries within 1 business day. Our standard customer service is available 24x365. Our technical support response times will be based on the support model purchased. For this proposal, the level of support that we propose comes with a response time of 60 minutes or shorter.

d. *You must provide design services for the applicable categories.*

Unisys is positioned to provide design services for the applicable categories. The Cloud Solutions section of our Cost Proposal lists the various services and their descriptions.

e. *You must provide Installation Services for the applicable categories.*

Unisys is positioned to installation services for the applicable categories. The Cloud Solutions section of our Cost Proposal lists the various services and their descriptions.

8.5 (E) SECURITY OF INFORMATION

8.5.1 *Offeror must describe the measures it takes to protect data. Include a description of the method by which you will hold, protect, and dispose of data following completion of any contract services.*

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

Unisys and our partners implemented a wide range of controls to protect the data. Each of us is certified and accredited to many industry and government standards. The responses to RFP requirements 8.6.1 through 8.6.6 list those standards. In general, data protection controls have been implemented at multiple layers across people, process, and technology. The responses to RFP requirements 8.6.1 through 8.6.6 provide these details.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS customers retain control and ownership of their data, and all data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers should consider the sensitivity of their data and decide if and how they will encrypt data while it is in transit and while it is at rest.

There are several options for encrypting data at rest, ranging from completely automated AWS encryption solutions to manual, client-side options. Choosing the right solutions depends on which AWS cloud services are being used and customer requirements for key management. Information on protecting data at rest using encryption can be found in the Protecting Data Using Encryption section of the Amazon Simple Storage Service (Amazon S3) Developer Guide available on the Internet.

For Securing Data in Transit services from AWS provide support for both Internet Protocol Security (IPSec) and Secure Sockets Layer/Transport Layer Security (SSL/TLS) for protection of data in transit. IPSec is a protocol that extends the IP protocol stack, often in network infrastructure, and allows applications on upper layers to communicate securely without modification. SSL/TLS, on the other hand, operates at the session layer, and while there are third-party SSL/TLS wrappers, it often requires support at the application layer as well. Unisys Stealth available as an AWS service provides end-to-end, IPSec based protection for data in-transit as well as workload isolation within the cloud and/or the private enterprise without the need for application modification.

It is important that customers understand some important basics regarding data ownership and management in the cloud shared responsibility model:

1. Customers continue to own their data.
2. Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
3. Customers can download or delete their data whenever they like.
4. Customers should consider the sensitivity of their data, and decide if and how to encrypt the data while it is in transit and at rest.

AWS provides customers with the ability to delete their data. However, AWS customers retain control and ownership of their data, and it is the customer's responsibility to manage their data

Data Recovery/Transfer

AWS allows customers to move data as needed on and off AWS storage using the public Internet or AWS Direct Connect (which lets customers establish a dedicated network connection between their network and AWS).

AWS Import/Export accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS transfers customer data directly onto and off of storage devices using Amazon's high-speed internal network and bypassing the Internet. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than customers upgrading their connectivity. With Import/Export encryption is mandatory, and AWS will encrypt customer data using the password they specified and transfer it onto the device

Deleting Data

Customers can use Multi-Object Delete to delete large numbers of objects from Amazon S3. This feature allows customers to send multiple object keys in a single request to speed up their deletes. Amazon does not charge customers for using Multi-Object Delete.

Customers can use the Object Expiration feature to remove objects from their buckets after a specified number of days. With Object Expiration customers can define the expiration rules for a set of objects in their bucket through the Lifecycle Configuration policy that they apply to the bucket. Each Object Expiration rule allows customers to specify a prefix and an expiration period.

Archiving Data

With Amazon S3's lifecycle policies, customers can configure their objects to be archived to Amazon Glacier or deleted after a specific period of time. Customers can use this policy-driven automation to quickly and easily reduce storage costs as well as save time. In each rule customers can specify a prefix, a time period, a transition to Amazon Glacier, and/or an expiration. For example, customers could create a rule that archives all objects with the common prefix "logs/" 30 days from creation, and expires these objects after 365 days from creation. Customers can also create a separate rule that only expires all objects with the prefix "backups/" 90 days from creation. Lifecycle policies apply to both existing and new S3 objects, ensuring that customers can optimize storage and maximize cost savings for all current data and any new data placed in S3 without time-consuming manual data review and migration.

AWS Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft data centers are in nondescript buildings that are physically constructed, managed, and monitored 24 hours a day to protect data and services from unauthorized access as well as environmental threats. The data centers are surrounded by a fence with access restricted by badge controlled gates. Preapproved deliveries are received in a secure loading bay and monitored by authorized personnel. Loading bays are physically isolated from information processing facilities.

CCTV is used to monitor physical access to data centers and the information systems. Cameras are positioned to monitor perimeter doors, facility entrances and exits, interior aisles, caged areas, high-

security areas, shipping and receiving, and facility external areas such as parking lots and other areas of the facilities.

Microsoft data centers receive SSAE16/ISAE 3402 Attestation and are certified to ISO 27001. Microsoft GFS and Azure maintain a current, documented, and audited inventory of equipment and network components for which they are responsible. GFS uses automated mechanisms to detect discrepancies of device configuration by comparing them with the defined policies. To prevent unauthorized access, GFS turns the unused ports off by default.

Windows Azure Fabric Controlled Hardware Device Authentication maintains a set of credentials (keys, passwords, or both) used to authenticate itself to various Windows Azure hardware devices under its control. The system used for transporting, persisting, and using these credentials is designed to make it unnecessary for Windows Azure developers, administrators, and backup services and personnel to be exposed to secret information. The Azure platform is designed and architected specifically to prevent the possibility of moving or replicating production data outside the Azure cloud environment. Controls for protecting production data include the following:

- Physical and logical network boundaries with strictly enforced change control policies
- Segregation of duty requiring a business need to access an environment
- Highly restricted physical and logical access to the cloud environment
- Strict controls based on SDL and OSA that define coding practices, quality testing, and code promotion
- Ongoing awareness and training on security, privacy, and secure coding practices
- Continuous logging and audit of system access
- Regular compliance audits to maintain control effectiveness.

Microsoft Azure customers must define policies and establish controls for the replication or high availability of production data and the demarcation of their production environment.

Internally, Microsoft tracks data flows and network connectivity among its facilities worldwide. Microsoft will not transfer Customer Data outside the geographies the customer specifies (for example, from Europe to the United States or from the United States to Asia) except when necessary for Microsoft to provide customer support, troubleshoot the service, or comply with legal requirements; or when the customer configures the account to enable the transfer of customer data, including through the use of the following features and functionality:

- Features that do not enable geographic selection such as Content Delivery Network (CDN) that provides a global caching service
- Web and worker roles, which back up software deployment packages to the United States regardless of deployment geography
- Preview, beta, or other prerelease features that may store or transfer customer data to the United States regardless of deployment geography
- Azure Active Directory (except for Access Control), which may store Active Directory data globally except for the United States (when Active Directory data remains in the United States) and Europe (when Active Directory Data is in Europe and the United States)
- Azure Multi-Factor Authentication, which stores authentication data in the United States
- Azure RemoteApp, which may store end user names and device IP addresses globally, depending on where the end user accesses the service.

Microsoft established policies, procedures, and mechanisms for effective key management to support encryption of data in storage and in transmission for the Azure service's key components. Azure provides each subscription with an associated logical certificate store that enables automatic deployment of service-specific certificates, and to which customers can upload their own.

Certificates used in Azure are x.509 v3 certificates and can be signed by another trusted certificate, or they can be self-signed. The certificate store is independent of hosted services, so it can store certificates regardless of whether those services currently use them. These certificates and other credentials uploaded to Azure are stored in encrypted form.

Microsoft Azure does not encrypt tenant data in storage. However, tools in Azure and third-party tools allow encryption of data in Azure storage. To implement encryption at rest, customers use .NET cryptographic services.

For customers using virtual machines, additional options are available, including the Encrypting File System in Windows Server 2008 R2 (and newer), Azure Rights Management Services, and Transparent Data Encryption (TDE) in SQL Server 2008 R2 (and newer).

When using Azure SQL Database, externally encrypted records cannot be queried by T-SQL (other than “retrieve all”) and may require a schema change such as the introduction of surrogate keys to enable retrieval of specific records or ranges of records.

Microsoft uses best practice procedures and a wiping solution that complies with NIST 800-88, Guidelines on Media Sanitization. For hard drives that cannot be wiped, Microsoft uses a destruction process that destroys it (shredding) and renders the recovery of information impossible (e.g., disintegrate, shred, pulverize, or incinerate). The appropriate means of disposal is determined by the asset type. Records of the destruction are retained. Microsoft Azure services use approved media storage and disposal management services. Paper documents are destroyed by approved means at the predetermined end-of-life cycle.

Data destruction techniques vary depending on the type of data object being destroyed, whether it be subscriptions, storage, virtual machines, or databases. In Azure’s multitenant environment, careful attention is taken to prevent one customer’s data from leaking into another customer’s data; or when a customer deletes data, no other customer (usually including the customer that once owned the data) from gaining access to that deleted data.

Azure follows NIST 800-88 guidelines, which address the principal concern of guarding data from unintentional release. These guidelines encompass electronic and physical sanitization.

Google Public and Community Cloud for SaaS: Google has deployed encryption services for data at rest and data in flight as one aspect of data protection. Data protection is a primary design consideration for all of Google’s infrastructure, applications and personnel operations. These factors are audited against several Industry Security Standards such as SOC2, ISO27001 and FedRAMP. The key security controls include multi-point perimeter protection, real time scanning for malware, systems designs that block or disable unnecessary services and protocols, least privilege access models, extensive logging and log review cycles. As data is written to Google’s infrastructure it is obfuscated, sharded, encrypted and replicated in a unique manner that not only ensure data availability but also security. As no single file for one single user at one single customer account is stored all in one place, any effort to access the data through unauthorized channels is extremely complex and highly unlikely.

Data disposal at the completion of any contract services is largely the responsibility of the customer. They would exercise an exit strategy that begins with migrating the data to a different service provider, deleting the data from Google, deleting the User accounts and closing or cancelling the service. Google agrees, via the Data Processing Agreement, that all customer data shards will be off all servers and backup tapes within 180 days of deletion.

8.5.2 Offeror must describe how it intends to comply with all applicable laws and related to data privacy and security.

In accordance with mandated policy, Unisys and our partners in this proposal comply with regulatory requirements for data privacy and security. Designated staff such as the contract manager, security

officers, and subject matter experts coordinate with the legal offices and external enforcement agencies to facilitate the requirements for data privacy and security.

8.5.3 *Offeror must describe how it will not access a Purchasing Entity's user accounts or data, except in the course of data center operations, response to service or technical issues, as required by the express terms of the Master Agreement, the applicable Participating Addendum, and/or the applicable Service Level Agreement.*

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

In accordance with established security policy, standards, and procedures, Unisys and our partners limit access to data relevant to job responsibilities, segregation of duties, and access controls. Technology controls are established to monitor, enforce, and report on violations. Unisys also provides provisions for a customer to enforce available security policies around their data as and when applicable. Unisys understands and respects the sensitivity of customer data and will not access a Purchasing Entity's user accounts or data.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS does not access customer data, and customers are given the choice as to how they store, manage and protect their data. AWS customers retain control and ownership of their data, and all data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers should consider the sensitivity of their data and decide if and how they will encrypt data while it is in transit and while it is at rest.

AWS Regions and Availability Zones: The AWS cloud infrastructure is built around regions and Availability Zones. A region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity and housed in separate facilities. These Availability Zones offer you the ability to operate production applications and databases that are more highly available, fault tolerant and scalable than would be possible from a single data center. AWS currently has 12 regions and 32 Availability Zones throughout the world: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), South America (Sao Paulo), and China (Beijing).. ~~Figure 1~~ depicts the current AWS regions and Edge Locations, along with new regions that are coming soon.

Formatte
Not Bold




Figure 1 – Global Map of AWS Regions and Edge Locations

Figure 2 illustrates the relationship between regions and Availability Zones.




Figure 2 – Regions and Availability Zones

Shared Responsibility: As cloud computing customers are building systems on top of cloud infrastructure, the security and compliance responsibilities are shared between the Cloud Service Providers (CSP) and cloud customers. In an Infrastructure as a Service (IaaS) model, customers control how they architect and secure their applications and data put on the infrastructure, while CSPs are responsible for providing services on a highly secure and controlled platform and providing a wide array of additional security features. The level of CSP and customer responsibilities in this shared responsibility model depends on the cloud deployment. Customers should be clear as to their responsibilities in each model. AWS's shared responsibility/security model is depicted in **Figure 3**.




Figure 34 – AWS Shared Responsibility Model

- **AWS Responsibility:** AWS operates, manages, and controls the infrastructure components, from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate.
- **Customer/Partner Responsibility:** Customers/partners assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, configuration of the AWS-provided security group firewalls, and other security, change management, and logging features.

AWS does not access customer data, and customers are given the choice as to how they store, manage and protect their data. There are four important basics regarding data ownership and management in the shared responsibility model:

- 1) Customers continue to own their data.
- 2) Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
- 3) Customers can download or delete their data whenever they like.
- 4) Customers should consider the sensitivity of their data and decide if and how to encrypt the data while it is in transit and at rest.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Privacy: Customers own and control their data.

For more than 20 years, Microsoft has been a leader in creating robust online solutions designed to protect its customers' privacy. Microsoft's time-tested approach to privacy and data protection is grounded in its commitment to organizations' ownership of and control over the collection, use, and distribution of their information.

Microsoft strives to be transparent in its privacy practices, offer customers meaningful privacy choices, and responsibly manage the data Microsoft stores and processes. One measure of Microsoft's commitment to the privacy of customer data is its adoption of the world's first code of practice for cloud privacy, ISO/IEC 27018.

Customers own their own data. With Azure, they have ownership of customer data, including text, sound, video, or image files and software, that are provided to Microsoft provides to them, or on their behalf, through the use of Azure. Customers can access their any time and for any reason without assistance from Microsoft. Microsoft does not use customer data or derive information from it for advertising or data mining.

Customers are in control of their data. Because the data customers host on Azure belongs to them, they have control over where their data is stored and how it is securely accessed and deleted.

Microsoft responds as follows to government and law enforcement requests to access data. When a government wants customer data—including for national security purposes—it must follow the applicable legal process, serving Microsoft with a court order for content or a subpoena for account information. If compelled to disclose customer data, Microsoft will promptly notify customers and provide a copy of the demand, unless legally prohibited from doing so. Microsoft does not provide governments with direct or unfettered access to customer data except as a customer directs or where required by law.

Security: Microsoft keeps customer data safe.

Microsoft leveraged its decades-long experience building enterprise software and running some of the world's largest online services to create a robust set of security technologies and practices. These help to confirm that Azure infrastructure is resilient to attack, safeguard user access to the Azure environment, and keep customer data secure through encrypted communications as well as threat management and mitigation practices, including regular penetration testing.

To manage and control identity and user access to customers' environments, data, and applications, Microsoft federates user identities to Azure Active Directory and enables multifactor authentication for more secure sign-in.

Microsoft encrypts communications and operation processes as follows. For data in transit, Azure uses industry-standard transport protocols between user devices and Microsoft data centers and in data centers themselves. For data at rest, Azure offers a wide range of encryption capabilities up to AES-256, giving customers the flexibility to choose the solution that best meets their needs.

To secure networks, Azure provides the infrastructure necessary to securely connect virtual machines to one another and to connect on-premises data centers with Azure virtual machines. To block unauthorized traffic to and in Microsoft data centers, Azure uses a variety of technologies. Azure Virtual Network extends a customer's on-premises network to the cloud over a site-to-site VPN.

To protect against online threats, Azure offers Microsoft Antimalware for cloud services and virtual machines. To help to mitigate threats to the Azure platform, Microsoft uses intrusion detection, denial-of-service (DDoS) attack prevention, regular penetration testing, and data analytics and machine learning tools.

To help organizations comply with national, regional, and industry-specific requirements governing the collection and use of individuals' data, Microsoft offers the most comprehensive set of certifications and attestations. Microsoft accomplishes this breadth of compliance offerings with a two-pronged approach:

- First, a team of Microsoft experts works with its engineering and operations teams, as well as external regulatory bodies, to track existing standards and regulations, developing hundreds of controls for the product teams to build into Microsoft's cloud services.
- Second, because regulations and standards are always evolving, Microsoft compliance experts anticipate upcoming changes to help maintain continuous compliance—researching draft regulations, assessing potential new requirements, and developing corresponding controls.

To demonstrate that these controls deliver compliance customers can rely on, Microsoft enterprise cloud services are independently validated through certifications and attestations as well as third-party audits.

In-scope services in the Microsoft Cloud meet key international and industry-specific compliance standards such as ISO/IEC 27001 and ISO/IEC 27018, FedRAMP, and SOC 1 and SOC 2. They also meet regional and country-specific standards and contractual commitments, including the EU Model Clauses, UK G-Cloud, Singapore MTCS, and Australia CCSL (IRAP). Additionally, rigorous third-party audits, such as by the British Standards Institution and Deloitte, validate the adherence of Microsoft cloud services to the strict requirements these standards mandate.

Google Public and Community Cloud for SaaS: Google will take and implement appropriate technical and organizational measures to protect Customer Data against accidental or unlawful destruction or accidental loss or alteration or unauthorized disclosure or access or other unauthorized processing. Google may update or modify the Security Measures from time to time provided that such updates and modifications do not result in the degradation of the overall security of the Services. Customer agrees that it is solely responsible for its use of the Services, including securing its account authentication credentials, and that Google has no obligation to protect Customer Data that Customer elects to store or transfer outside of Google's and its Subprocessors' systems (e.g., offline or on-premise storage). Google commits via terms of service to abide by all applicable laws and to hold all of customer's data as confidential and solely owned by the customer.

8.6 (E) PRIVACY AND SECURITY

8.6.1 Offeror must describe its commitment for its Solutions to comply with NIST, as defined in NIST Special Publication 800-145, and any other relevant industry standards, as it relates to the Scope of Services described in Attachment D, including supporting the different types of data that you may receive.

8.6.2 Offeror must list all government or standards organization security certifications it currently holds that apply specifically to the Offeror's proposal, as well as those in process at time of response. Specifically include HIPAA, FERPA, CJIS Security Policy, PCI Data Security Standards (DSS), IRS Publication 1075, FISMA, NIST 800-53, NIST SP 800-171, and FIPS 200 if they apply.

8.6.3 Offeror must describe its security practices in place to secure data and applications, including threats from outside the service center as well as other customers co-located within the same service center.

8.6.4 Offeror must describe its data confidentiality standards and practices that are in place to ensure data confidentiality. This must include not only prevention of exposure to unauthorized personnel, but also managing and reviewing access that administrators have to stored data. Include information on your hardware policies (laptops, mobile etc).

8.6.5 Offeror must provide a detailed list of the third-party attestations, reports, security credentials (e.g., FedRamp), and certifications relating to data security, integrity, and other controls.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4 and 8.6.5:

Unisys Cloud Services align with the NIST definition of cloud computing (NIST 800-145). Additionally, Unisys is certified to ISO 27001 for information security, which meets the security requirements for cloud computing. Along with ISO 27001, Unisys also aligned our security organization with other standards and frameworks such as HIPAA, HITECH, GLBA, ISAE 3402/SSAE 16, SOGP, NIST, Data Protection Acts, COBIT, PCI-DSS, FedRAMP, FISMA, CSA, FFIEC, ISO (27000, 9000, 20000, and 22301), ENISA, and various other country- and region-specific standards and requirements.

Unisys delivery centers are certified to ISO 27001 (Information Security), ISO 20000 (Service Management), ISO 9001 (Quality Management), and ISO 22301 (Business Continuity). Several of our client environments are certified to PCI-DSS standards. Additionally, our delivery centers are assessed annually for SOC 1 reporting. Independent third-party agencies audit these certifications every year.

Where required, physical segregation of the work space is implemented, and the workspace is secured through access controls. Networks are segregated by VLANs and protected by firewalls. Access control mechanisms are implemented to access data and applications.

Unisys adopted a data protection policy and standards for securing data. These standards define controls such as data classification, system maintenance, access controls, data storage, external data transfer,

media handling, and media disposal. These controls are followed across Unisys and audited regularly for compliance. To complement these controls, Unisys also has policies for Internet and email use, mobile devices, and acceptable use, supplemented by various security standards.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4, and 8.6.5:

The AWS cloud infrastructure is designed and managed in alignment with the following regulations, standards, and best practices:

- Federal Risk and Authorization Management Program (FedRAMP)
- Service Organization Controls (SOC) 1/American Institute of Certified Public Accountants (AICPA): AT 801 (formerly Statement on Standards for Attestation Engagements [aud] No. 16)/International Standard on Assurance Engagements (ISAE) 3402 (formerly Statement on Auditing Standards [SAS] No. 70)
- SOC 2
- SOC 3
- Payment Card Industry Data Security Standard (PCI DSS)
- International Organization for Standardization (ISO) 27001
- ISO 27017
- ISO 27018
- ISO 9001
- Department of Defense (DoD) Security Requirements Guide (SRG) security impact levels 2 and 4
- Federal Information Security Management Act (FISMA)
- US Health Insurance Portability and Accountability Act (HIPAA)
- FBI Criminal Justice Information Services (CJIS)
- National Institute of Standards and Technology (NIST) 800-171
- International Traffic in Arms Regulations (ITAR)
- Federal Information Processing Standard (FIPS) 140-2
- Family Educational Rights and Privacy Act (FERPA)
- Information Security Registered Assessors Program (IRAP) (Australia)
- IT-Grundschutz (Germany).

The AWS white paper in **Appendix 14 AWS Risk and Compliance Whitepaper** contains further details:

See Appendix 14 AWS Risk and Compliance Whitepaper

Additionally, the AWS white paper in **Appendix 15 AWS Security Whitepaper** also contains relevant information surrounding privacy and security.

See Appendix 15 AWS Security Whitepaper

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4 and 8.6.5:

Microsoft agrees that, during the term of a Purchasing Entity's subscription for its Government Community Cloud Services (as defined in Microsoft's terms and conditions), those services will be operated in accordance with a written data security policy and control framework that is consistent with the requirements of NIST 800-53 Revision 4, or successor standards and guidelines (if any), established to support Federal Risk and Authorization Management Program (FedRAMP) accreditation at a Moderate Impact level. Microsoft intends for Government Community Cloud Services to support FedRAMP Authority to Operate (ATO). Microsoft will use commercially reasonable efforts to obtain an ATO from a Federal agency, and to maintain this ATO through continuous monitoring processes and by conducting

regular FedRAMP audits. The Microsoft Trust Center available on line contains detailed information on Microsoft's compliance and adherence to other standards, such as CJIS, IRS 1075, HIPAA, FERPA, ISO/IEC 27001 and 27018, SOC 1, and SOC 2. The **Windows Azure Security and Compliance** white paper in **Appendix 24** contains further details:

See Appendix 24 **Windows Azure Security and Compliance White Paper**

Microsoft Dynamics CRM Online Government has received an Agency ATO from HUD; and Microsoft Office 365 U.S. Government has received an Agency ATO from DHHS.

The US Food and Drug Administration (FDA) Code of Federal Regulations (CFR) Title 21 Part 11 lists requirements for the security of electronic records of companies that sell food and drugs manufactured or consumed in the United States

The Defense Information Systems Agency (DISA) Cloud Service Support has granted a DISA Impact Level 2 Provisional Authorization (PA) to Microsoft Azure, Microsoft Azure Government, Microsoft Office 365 MT, and Microsoft Office 365 U.S. Government, based on their FedRAMP authorizations.

Microsoft contractual commitments, customers that are subject to FERPA can use Microsoft Azure, Microsoft Dynamics CRM Online, and Microsoft Office 365 and comply with FERPA.

NIST publishes a list of vendors and their cryptographic modules validated for FIPS 140-2. Rather than validate individual components and products, Microsoft certifies the underlying cryptographic modules used in Microsoft products, including Microsoft enterprise cloud services.

Microsoft engaged outside assessors to validate that Microsoft Azure and Microsoft Office 365 meet the FISC Version 8 requirements.

Microsoft enterprise cloud services offer customers a HIPAA Business Associate Agreement (BAA) that stipulates adherence to HIPAA's security and privacy provisions.

Microsoft Azure and Microsoft Office 365 were among the first cloud services to achieve this certification for the storage and processing of unclassified (DLM) data.

Microsoft Azure Government and Microsoft Office 365 U.S. Government cloud services provide a contractual commitment that they have the appropriate controls in place, and the security capabilities necessary for customers to meet the substantive requirements of IRS 1075.

The ISO/IEC 27001 certificate validates that Microsoft enterprise cloud services have implemented the internationally recognized information security controls defined in this standard, including guidelines and general principles for initiating, implementing, maintaining, and improving information security management within an organization.

Microsoft was the first cloud provider to adhere to the ISO/IEC 27018 code of practice, covering the processing of personal information by cloud service providers.

Azure complies with Payment Card Industry (PCI) Data Security Standards (DSS) Level 1 version 3.0, the global certification standard for organizations that accept most payment cards and store, process, or transmit cardholder data.

A Voluntary Product Accessibility Template, or VPAT, is a standardized form developed by the Information Technology Industry Council to document whether a product meets key regulations of Section 508, an amendment to the Rehabilitation Act of 1973. Microsoft offers detailed VPATs for many of its core cloud services, describing the accessibility features of those services.

Service Organization Controls (SOC) are a series of accounting standards that measure the control of financial information for a service organization. Azure's SOC 1 and SOC 2 Type 2 audit reports attest to the effectiveness of the design and operation of its security controls.

Other country-specific standards and contractual commitments, including the EU Model Clauses, UK G-Cloud, Singapore MTCS, and Australia CCSL (IRAP).

As shown in Exhibit 5, Microsoft cloud services have the largest compliance portfolio in the industry, the United States, and the world.

Exhibit 5. Regulatory Compliance of Microsoft Cloud Services.

Regulatory Compliance

Microsoft cloud services have the largest compliance portfolio in the industry

Industry	 ISO 27001	 SOC 1 Type 2	 SOC 2 Type 2	 PCI DSS Level 1	 Cloud Controls Matrix	 ISO 27018	 Content Delivery and Security Association	 Shared Assessments			
United States	 FedRAMP JAB P-ATO	 HIPAA / HITECH	 FIPS 140-2	 21 CFR Part 11	 FERPA	 DISA Level 2	 CJIS	 IRS 1075	 ITAR ready	 Section 508 VPA	
Regional	 European Union Model Clauses	 EU Safe Harbor	 GOV.UK	 China Multi Layer Protection Scheme	 China GB 18030	 China CCCPF	 Singapore MTCS Level 3	 Australian Signals Directorate	 New Zealand GCIO	 Japan Financial Services	 ENISA IAF

Data in Government Clouds are maintained in Azure Datacenters in the United States. Azure Government is a government-community cloud that you can trust, enabling US government agencies and their partners to transform their mission-critical workloads to the cloud. A physical and logical network-isolated instance of Azure, operated by screened US persons, Azure Government has been specifically designed to meet the needs of US government agencies and their partners.

Security Practices in Place to Secure Data and Applications:

The Microsoft Azure trustworthy foundation concept maintains application security through a process of continuous security improvement with its Security Development Lifecycle (SDL) and Operational Security Assurance (OSA) programs using Prevent Breach and Assume Breach security postures.

Prevent Breach works through the use of ongoing threat modeling, code review, and security testing; Assume Breach uses Red Team Blue Team exercises, live site penetration testing, and centralized security logging and monitoring to identify and address potential gaps; test security response plans; reduce exposure to attack; and reduce access from a compromised system, periodic postbreach assessment, and clean state.

To validate services, Azure uses third-party penetration testing based on the OWASP (Open Web Application Security Project) top 10 and CREST-certified testers. The outputs of testing are tracked on the risk register, which is audited and reviewed regularly to confirm compliance with Microsoft security practices.

Microsoft maintains and regularly updates the Azure Information Security Management Policy and information security guidelines, standard operating procedures for data security, and contractual commitments to international data protection directives that apply across Azure services.

Additionally, Microsoft Azure software updates are reviewed for unauthorized changes through the Security Development Lifecycle (SDL) change and release management processes. Automated mechanisms are used to perform integrity scans at least every hour to detect system anomalies or unauthorized changes. Microsoft applies SDL to design, develop, and implement Microsoft Azure services. SDL helps to confirm that communication and collaboration services are highly secure, even at the foundation level, and align with other industry standards, including FedRAMP, ISO, and NIST.

Data Confidentiality Standards and Practices That Are in Place to Maintain Data Confidentiality

Microsoft Azure adopted applicable corporate and organizational security policies, including an Information Security Policy. The policies were approved, published, and communicated across Azure teams. The Information Security Policy requires that access to Microsoft Azure assets to be granted based on business justification, with the asset owner's authorization and limits based on "need-to-know" and "least privilege" principles. Additionally, the policy also addresses requirements for the access management life cycle, including access provisioning, authentication, access authorization, removal of access rights, and periodic access reviews.

Access to Microsoft buildings is controlled, and access is restricted to those with a card reader (swiping the card reader with an authorized ID badge) or biometrics for entry into data centers. Front-desk personnel are required to positively identify full-time employees or authorized contractors without ID cards. Staff must always wear identity badges and are required to challenge or report individuals without badges. Authorized Microsoft personnel must escort guests.

Password policies for corporate domain accounts are managed through Microsoft's corporate Active Directory policy that specifies minimum requirements for password length, complexity, and expiration. Temporary passwords are communicated to users using Microsoft's established processes. Azure services and infrastructure must at least meet Microsoft corporate requirements, but an internal organization can increase the strength beyond this standard, at their own discretion and to meet their security needs. Technological safeguards, such as encrypted communications and operational processes, enhance the security of our customers' data. For data in transit, the Microsoft Cloud uses industry-standard encrypted transport protocols between user devices and Microsoft datacenters, and within datacenters themselves. For data at rest, the Microsoft Cloud offers a wide range of encryption capabilities up to AES-256, giving you the flexibility to choose the solution that best meets your needs.

Microsoft Azure uses Active Directory (AD) to manage user accounts. Azure Active Directory is a comprehensive identity and access management cloud solution that helps secure access to your data and on-premises and cloud applications, and simplifies the management of users and groups. It combines core directory services, advanced identity governance, security, and application access management, and is a key component of Microsoft Cloud services, including Microsoft Azure, Office 365, Microsoft Dynamics CRM Online, and Intune, as well as thousands of third-party SaaS apps. Azure Active Directory also makes it easy for developers to build policy-based identity management into their applications. The designated security group owners must approve security group membership in Microsoft Azure. Automated procedures are in place to disable AD accounts on the user's leave date. Domain-level user accounts are disabled after 90 days of inactivity.

Mobile and wireless devices are not permitted in Microsoft data centers where customer data is stored. Wireless access to Azure production environments, customer environments, or both is prohibited.

Employees, contractors, and third-party users are formally notified to destroy or return, as applicable, physical material that Microsoft provided to them during their employment or period of the contractor agreement; electronic media must be removed from contractor or third-party infrastructure. Microsoft may also conduct an audit to verify that data is removed in an appropriate way.

Google Public and Community Cloud for SaaS Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4 and 8.6.5:

Google has an FedRAMP ATO at the Moderate impact baseline. FedRAMP incorporates many NIST SPs and FIPS including 800-53, FIPS 199, FIPS 200), and has a specific offering. Google Apps for Education that is FERPA and COPPA compliant. Other compliance standards such as HIPAA and CJIS don't offer certification per se, but are commonly accommodated (i.e. Google will sign a BAA to meet HiTECH/HIPAA requirements, and has numerous customers who bear responsibility for meeting CJIS processing requirements). PCI DSS is generally not applicable to SaaS systems (though we can do email

hygiene processing to protect against incidental usage), but Google IaaS/PaaS does meet PCI DSS v3 standards. Google also holds and is committed to maintaining SOC2 and ISO27001 certifications.

At Google, members of the information security team review security plans for networks, systems, and services. They provide project-specific consulting services to Google's product and engineering teams. They monitor for suspicious activity on Google's networks, address information security threats, perform routine security evaluations and audits, and engage outside experts to conduct regular security assessments. Google specifically built a full-time team, known as Project Zero, which aims to prevent targeted attacks by reporting bugs to software vendors and filing them in an external database. Google has automated mechanisms that scan the infrastructure continuously to detect and correct deviations from the desired security configuration of its infrastructure. Google's security monitoring program is focused on information gathered from internal network traffic, employee actions on systems, and outside knowledge of vulnerabilities. At many points across Google's global network, internal traffic is inspected for suspicious behavior, such as the presence of traffic that might indicate botnet connections. This analysis is performed using a combination of Open Source and commercial tools for traffic capture and parsing. A proprietary correlation system built on top of Google technology also supports this analysis. Network analysis is supplemented by examining system logs to identify unusual behavior, such as attempted access of customer data. Google security engineers issue standing search alerts on public data repositories to look for security incidents that might affect Google's infrastructure. They actively review inbound security reports and monitor public mailing lists, blog posts, and wikis. Automated network analysis helps determine when an unknown threat may exist and escalates the threat to Google security staff; network analysis is supplemented by automated analysis of system logs. Google uses a proprietary storage and processing mechanism that isolates processing of data in change-rooted (chrooted) jails in a physical server. Access permissions restrict the ability for processes to interact between jails. Additionally, tenant data is stripped in chunks across many different drives; each chunk has its own access control list. This helps confirm that data is logically isolated between customers in storage and during processing. In addition to the processing and storage mechanisms already described, Google's security controls, including least privilege rights, logging, and auditing, were implemented consistently with FedRAMP requirements.

A copy of Google's Data Processing Agreement (DPA) is included in this proposal as **Appendix 16-Google-DPA**. See Appendix 16

In this DPA, the obligations of Google to hold customer data as confidential and wholly owned by the customer is detailed. Google's internal data access processes and policies are designed to prevent unauthorized persons, systems, or both from gaining access to systems used to process personal data. Google aims to design its systems to: (1) only allow authorized persons to access data they are authorized to access; and (2) confirm that personal data cannot be read, copied, altered, or removed without authorization during processing, use, and after recording. The systems are designed to detect inappropriate access. Google uses a centralized access management system to control personnel access to production servers and provides access only to a limited number of authorized personnel. LDAP, Kerberos, and a proprietary system using RSA keys are designed to provide Google with secure and flexible access mechanisms. These mechanisms are designed to grant only approved access rights to site hosts, logs, data, and configuration information. Google requires the use of unique user IDs, strong passwords; two-factor authentication, and carefully monitored access lists to minimize the potential for unauthorized account use. The granting or modification of access rights is based on the authorized personnel's job responsibilities, job duty requirements necessary to perform authorized tasks, and the need to know; they also must be in accordance with Google's internal data access policies and training. Approvals are managed by workflow tools that maintain audit records of changes. Access to systems is logged to create an audit trail for accountability. Where passwords are employed for authentication (e.g., login to workstations), password policies that follow at least industry-standard practices are implemented.

These standards include password expiration, restrictions on password reuse, and sufficient password strength. For access to extremely sensitive information (e.g., credit card data), Google uses hardware tokens. **Appendix 23 contains the Google Apps Security and Compliance** whitepaper contains additional details surrounding privacy and Security. See Appendix 23.

8.6.6 Offeror must describe its logging process including the types of services and devices logged; the event types logged; and the information fields. You should include detailed response on how you plan to maintain security certifications.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

Unisys has a Unisys Cloud Provisioning System logs all actions by both administrators and users. This logging system will be accessible by authorized Unisys personnel who can manage and monitor data for specific log types. Administrators can build templates and generate alerts on custom correlation events. All logs are text-based and depending on the log types can be customized between a syslog or log4j. This gives Unisys Cloud Provisioning System the ability to centralize log data and use specific tools such as splunk, q1 labs or SIEM tools to access logs.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS CloudTrail

AWS CloudTrail is a web service that records API calls to supported AWS services in an AWS account, delivering a log file to an Amazon Simple Storage Service (Amazon S3) bucket. To alleviate common challenges experienced in an on-premises environment, AWS CloudTrail makes it easier for customers to enhance security and operational processes while demonstrating compliance with policies or regulatory standards.

With AWS CloudTrail, customers can get a history of AWS API calls for their account, including API calls made from the AWS Management Console, AWS SDKs, command line tools, and higher level AWS services (such as AWS CloudFormation). The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

- For information on the services and features supported by AWS CloudTrail, visit the AWS CloudTrail FAQs at the AWS website.
- The AWS white paper Security at Scale: Logging In AWS provides an overview of common compliance requirements related to logging, detailing how AWS CloudTrail features can help satisfy these requirements.
- The AWS white paper Auditing Security Checklist for Use of AWS provides customers with a checklist to assist in evaluating AWS for an internal review or external audit.

AWS CloudTrail includes the following features:

- **Increased visibility:** AWS CloudTrail provides increased visibility into user activity by recording AWS API calls. Customers can answer questions such as: What actions did a given user take over a given period? For a given resource, which user took actions on it over a given period? What is the source IP address of a given activity? Which activities failed because of inadequate permissions?
- **Durable and inexpensive log file storage:** AWS CloudTrail uses Amazon S3 for log file storage and delivery, so that log files are stored durably and inexpensively. Customers can use Amazon S3 life cycle configuration rules to further reduce storage costs. For example, customers can define rules to automatically delete old log files or archive them to Amazon Glacier for additional savings.
- **Easy administration:** AWS CloudTrail is a fully managed service; customers simply turn on AWS CloudTrail for their account at the AWS Management Console, the Command Line Interface, or the AWS CloudTrail SDK to start receiving AWS CloudTrail log files in the specified Amazon S3 bucket.

- **Notifications for log file delivery:** AWS CloudTrail can be configured to publish a notification for each log file delivered, thereby enabling customers to automatically take action upon log file delivery. AWS CloudTrail uses the Amazon Simple Notification Service (Amazon SNS) for notifications.
- **Choice of partner solutions:** Many partners, including AlertLogic, Boundary, Loggly, Splunk, and Sumologic, offer integrated solutions to analyze AWS CloudTrail log files. These solutions include features such as change tracking, troubleshooting, and security analysis. For more information, see the AWS CloudTrail partners section of Amazon's home page.
- **Log file aggregation:** AWS CloudTrail can be configured to aggregate log files across multiple accounts and regions so that log files are delivered to a single bucket. For detailed instructions, refer to the Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket section of the user guide.

Amazon CloudWatch

Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications that run on AWS. Customers can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by customer applications and services, and log files that application generate. Customers can use Amazon CloudWatch to gain systemwide visibility into resource utilization, application performance, and operational health, using these insights to react and keep their application running smoothly.

Customers also can use CloudWatch Logs to monitor and troubleshoot systems and applications using their existing system, application, and custom log files. They can send their existing system, application, and custom log files to CloudWatch Logs and monitor these logs in near real time. This helps customers to better understand and operate their systems and applications; they can move their logs to highly durable, low-cost storage for later access.

LogAnalyzer for Amazon CloudFront

LogAnalyzer allows customers to analyze their Amazon CloudFront Logs using Amazon Elastic MapReduce (Amazon EMR). Using Amazon EMR and the LogAnalyzer application, customers can generate usage reports containing total traffic volume, object popularity, a breakdown of traffic by client IPs, and edge location. Reports are formatted as tab delimited text files and delivered to the Amazon S3 bucket that customers specify.

Amazon CloudFront's Access Logs provide detailed information on requests made for content delivered through Amazon CloudFront, the AWS content delivery service. The LogAnalyzer for Amazon CloudFront analyzes the service's raw log files to produce a series of reports that answer business questions commonly asked by content owners.

Reports Generated

This LogAnalyzer application produces four sets of reports based on Amazon CloudFront access logs. The Overall Volume Report displays the total amount of traffic delivered by CloudFront during the specified period. The Object Popularity Report shows how many times each customer object is requested. The Client IP report shows the traffic from each different client IP that made a request for content. The Edge Location Report shows the total number of traffic delivered through each edge location. Each report measures traffic in three ways: the total number of requests, the total number of bytes transferred, and the number of request broken down by HTTP response code. The LogAnalyzer is implemented using Cascading and is an example of how to construct an Amazon Elastic MapReduce application. Customers can also customize reports generated by the LogAnalyzer.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Azure enables customers to perform security event generation and collection from Azure Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) roles to central storage in their subscriptions. Customers

can then use HDInsight to aggregate and analyze the collected events. Additionally, these collected events can be exported to on-premises security information and event management (SIEM) systems for ongoing monitoring.

The Azure security logging, analysis, and monitoring life cycle includes the following tasks:

- Generation: Instrument applications and the infrastructure to raise events.
- Collection: Configure Azure to collect the various security logs in a storage account.
- Analysis: Use Azure tools such as HDInsight and on-premises SIEM systems to analyze the logs and generate security insights.
- Monitoring and reporting: Azure offers centralized monitoring and analysis systems that provide continuous visibility and timely alerts.

Security events are flagged on the Windows Event Log for the System, Security, and Application channels on virtual machines. To prevent events from being logged without potential data loss, it is important to appropriately configure the size of the event log. The size of this log must be based on the number of events that auditing policy settings generate and the event collection policies defined.

For Cloud Services applications that are deployed in Azure and virtual machines created from the Azure Virtual Machines Marketplace, a set of operating system security events are enabled by default. Customers can add, modify, or delete events to be audited by customizing the operating system audit policy. For more information, see Security Policy Settings Reference.

Customers can use the following methods to generate additional logs from operating system (such as audit policy changes) and Windows components (such as IIS):

- Group Policy to roll out policy settings for virtual machines in Azure that are domain joined
- Desired State Configuration (DSC) to push and manage policy settings; for more information, see Azure PowerShell DSC
- Service Deployment role startup code to roll out settings for Cloud Services (PaaS scenario).

Collection of security events and logs from Cloud Services or virtual machines in Azure occurs through two main methods:

- Azure Diagnostics, which collects events in a customer's Azure storage account
- Windows Event Forwarding (WEF), a technology on computers that run Windows.

Exhibit 6 lists several key differences between these two technologies. In accordance with a customer's requirements and these key differences, the appropriate method must to be chosen to implement log collection.

Exhibit 6. Key Differences between Azure Diagnostics and Windows Event Forwarding.

Azure Diagnostics	Windows Event Forwarding
Supports Azure Virtual Machines and Azure Cloud Services	Supports domain-joined Azure Virtual Machines only
Supports a variety of log formats, such as Windows event logs, Event Tracing for Windows (ETW) traces, and IIS logs. For more information, see Azure Diagnostics supported data sources	Supports Windows event logs only
Pushes collected data to Azure Storage	Moves collected data to central collector servers

Security Event Data Collection with Windows Event Forwarding

For domain-joined Azure Virtual Machines, customers can configure WEF by using Group Policy settings in the same way as for on-premises domain-joined computers. Using this approach, an organization could purchase an IaaS subscription, connect it to its corporate network by using ExpressRoute or site-to-site VPN, and then join the virtual machines that a customer has in Azure to the corporate domain. Afterwards, the customer can configure WEF from the domain-joined machines.

Event forwarding is broken into two parts: the source and the collector. The source is the computer in which the security logs are generated. The collector is the centralized server that collects and consolidates the event logs. IT administrators can subscribe to events so that they can receive and store events that are forwarded from remote computers (the event source). Collected Windows events can be sent to on-premises analysis tools, such as a SIEM, for further analysis.

Security Data Collection with Azure Diagnostics

Azure Diagnostics enables customers to collect diagnostic data from a cloud service worker role or web role, or from virtual machines running in Azure. It is a predefined guest agent extension that needs to be enabled and configured for data collection. A customer's subscription can include pushing the data to Azure Storage.

The data is encrypted in transit (by using HTTPS). The examples provided in this response use Azure Diagnostics 1.2. Microsoft recommends that customers upgrade to version 1.2 or higher for security data collection.

Exhibit 7 shows a high-level dataflow for security data collection that uses Azure Diagnostics and further analysis and monitoring.

Exhibit 7. High-level Dataflow for Collection of Security Data.

Google Public and Community Cloud for SaaS: Exhibit 8 lists Google's logging and reporting capability for the various Google offerings.

Exhibit 8. Google Reports.

Category	Report	Information Fields
Highlights	Apps Usage Activity	Total Email Files Owned Video Hangouts By Week, Month, Quarter, Half Year
	User Status	Total Number of Users Blocked Users Suspended Users Active Users
	Security	External Apps Installed Users not enrolled in Two-Step Verification Use of Less Secure (Third-Party) Apps
	Document Link Shared Status	Count of the following Shared Files: • Visible to Public • Visible to Anyone with Link • Visible to people at Domain • Visible to people at Domain with link • Private
Aggregate Reports 6 months of rolling data	GMail	Inbound Email Delivery Count of Delivered, Rerouted, Rejected Inbound Email Encryption Outbound Email Encryption Inbound Email Spam Outbound Email Delivery Total Emails
	Accounts	Two-step verification enrollment Two-step verification enforcement User account status Less secure apps access Admin status shows the number of users assigned to each Admin role defined
	Drive	Counts by day of the following files: • External Link Shared Files • Internal Link Shared Files • Total Files Owned
	Chrome Devices	Count of Weekly Chrome Device logins Count of devices by release channel Count of devices by boot mode Count of devices by version
	Mobile Managed Devices requires end	Count of 7-day-active Managed Devices by mobile OS (Android, iOS, and Active Sync)

Category	Report	Information Fields
	user configuration of mobile device management App. Management of Android, iOS and Active-Sync supported devices is included.	Count of 30-day active Managed Devices 7- and 30-day active managed Users Managed Android Devices by Android version Managed iOS Devices by iOS version
	Google+	Count of number of Video Hangouts Activity report on Chromebox for Meeting Devices
Security	General	Same as Accounts reports listed above.
	GMail	(POP) Last time used (IMAP) Last time used (Web) Last time used
	Drive	Same as Drive reports listed above.
Apps Usage Activity	General	Count by User of the following storage: • Total Storage used • GMail Storage used • Drive Storage used • Photos Storage used
	GMail	Same as GMail reports listed in Security and Aggregate Reports section except this is a count by user versus count by domain.
	Files Owned	Count by User of the following files: • Files Uploaded • Google Docs • Google Sheets • Google Slides • Google Forms • Google Drawings.
Account Activity	User Account Status	Record by user instead of by domain that covers the same attributes as the Security/General Report data, the Security/GMail data, and the Security and Apps Usage Drive data.
Audit Reports	Admin Activity	Log-level details of authorized administrator actions, including the following information: • Event Name • Event Description • Administrator Name • Data • IP Address.
	Login Report	End User Login Report that lists the following information: • User Name • Event Name - Failed Login - Login Challenge - Logout - Successful Login

Category	Report	Information Fields
		Suspicious Login • IP Address • Date • Login Type - Google Password - SAML
	Drive Audit Log ¹	A Drive activity report that lists the following information: • Document Title • Document ID • Document Type • Event Name • Event Description - Add to Folder - Create - Delete - Download - Edit - Move - Preview - Print - Remove from Folder - Rename - Restore - Trash - Upload - View - Editor Settings Change - Link Sharing Access Type Change • Date • Owner • User • IP Address.
	Calendar Audit Log	A Calendar activity report that lists the following information: • Activity Name - Access Level Change - Country Change - Calendar created - Calendar deleted - Description changed - Location changed - Time zone changed - Title changed - Subscription added or deleted - Event created - Event deleted - Guest added

¹ Available only with Google Apps Unlimited

Category	Report	Information Fields
		Guest removed Guest response change Event modified Event removed from trash Event restored Start time changed Title changed • User • Calendar ID • Date • Event Title • Event ID • API Kind • User Agent • IP Address.
	Token Report	A log of the following information on user activities tied to authorizing OAuth tokens for third-party apps: • Event Name Authorize Revoke • Application Name • User • Scope • Date • Client ID • IP Address.
	Groups Log	A log of the following information on user activities related to Google Groups for Business: • Event Name Group permission changed User accepted invitation to a group Owner approved request to join from Another user User joined a group User requested to join a group Basic settings changed Group created Group deleted Identity setting changed Information setting changed Information setting removed New member restriction changed Post replies setting changed Spam moderation setting changed Topic setting changed Message moderated Post from users will always be posted User added to group User banned from group User invitation revoked User invited

Category	Report	Information Fields
		Join request rejected User reinvited User removed from group • Event Description Includes the names of user and Group name • User • Date
	Email Log Search	A query engine to look up log detail on email delivery by selecting the following attributes: • Date • Sender • Recipient • Sender IP • Recipient IP • Message ID
Alerts	Predefined Alerts	Drive Settings Change TLS Failure User Deleted User Granted Admin Permission Exchange Journal Failure Suspended User made Active User Admin Privilege revoked Smarthost Failure GMail Settings Change User's Password Changed User Suspended Suspicious Login Activity Mobile Settings Change Calendar Settings Change Apps Outage Alert New User Added
	Custom Alerts	Custom alerts can be created against the following audit stream types: • Admin Audit • Token Audit • Calendar Audit • Login Audit • Drive Audit.

8.6.7 Offeror must describe whether it can restrict visibility of cloud hosted data and documents to specific users or groups.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys can restrict visibility of cloud hosted data and documents to specific users or groups. Unisys can accomplish this in multiple ways:

Unisys Stealth: The Unisys Stealth™ software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. By substituting traditional hardware topology for software-based cryptography, our Stealth Micro segmentation solutions prevent unauthorized access to sensitive information and reduce the

attack surface, thereby making endpoints invisible to certain users or user groups. Unisys Stealth tightens access control by focusing on user identity rather than physical devices, so security moves with the user and is easier to manage. It protects sensitive data in motion from potential compromise through encryption.


AD/LDAP: Unisys has in past achieved data access restriction to specific users or groups by synchronizing a user database with an external Lightweight Directory Access Protocol (LDAP) or Active Directory (AD). User roles and groups can be maintained in LDAP and AD which will authenticate and authorize each user based on its privileges. Access rights are determined by role. System administrators can create their own roles and map them to groups to suit a customer's business needs. System administrators can also define permissions for each role at a very granular level.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS Identity and Access Management (IAM) is a web service that enables AWS customers to manage users and user permissions in AWS. The service is targeted at organizations with many users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console. With AWS IAM, system administrators can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.

Permissions let system administrators specify who has access to AWS resources and which actions they can perform on those resources. Every AWS Identity and Access Management (IAM) user starts with no permissions. In other words, by default, users can do nothing—not even view their own access keys. To give a user permission to do something, system administrators can add the permission to the user (attach a policy to the user), or add the user to a group that has the desired permission.

Exhibit 9 shows the AWS IAM tab on the AWS Management Console.

Exhibit 9. AWS IAM Tab.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft Azure can restrict visibility of cloud-hosted data and documents to specific users or groups. Customer environments and data in Azure are isolated using numerous mechanisms, technologies, policies, processes, and architectural elements. Several of them are listed as follows:

- Virtual Networks – customer tenants and virtual machine deployments are kept logically separated through VNets that define DNS, security policies, and IP routing rules. Firewalls, ACLs, network security groups, IP filters, virtual appliances, load balancers, and network policies work together to prevent unauthorized traffic from entering or leaving a customer's tenant, either across network boundaries or between the virtualization host and guest.
- Encryption – customer data is encrypted in transit and at rest through configurable and standards-based providers using a variety of protocols such as BitLocker, AES-256 (in Azure Media Services), and IPsec (VNets).
- Access Control – Azure Storage, the Azure Portal, and other service components provide role-based access controls and key-based authentication to help confirm that only authorized entities can gain access to tenant data.

The concept of tenant containers is maintained in the Azure Active Directory service at several layers, from portals to persistent storage. These boundaries prevent a query scoped to a given tenant from returning directory data for another tenant. Front ends (Azure AD Sync, PowerShell, and Graph) store and retrieve data through an internal directory services API (DSAPI) that calls an authorization layer to confirm the data requested is allowed for the user requesting the data. These capabilities are available to customers for isolating and protecting their data, gaining access to only their data and no others.

Azure Resource Management RBAC roles have support for Azure Service Management API (Classic) resources using the following RBAC roles:

- Classic Network Contributor
- Classic Storage Contributor
- Classic Virtual Machine Contributor

Using these RBAC roles, it is possible to assign limited access to classic resources in the ARM Azure portal. The access is restricted to the abilities in the ARM portal for management of the resources.

RBAC is supported on classic Compute, Storage, and Networking objects. Compute includes IaaS VMs and PaaS Web/Worker roles. Networking includes vNets and subnets (NSGs are currently not supported). Storage includes storage accounts. Only classic resources in these three roles are supported.

Azure Resource Manager provides the ability to restrict operations on resources through resource management locks. Locks are policies which enforce a lock level at a particular scope. The scope can be a subscription, resource group or resource.

Google Public and Community Cloud for SaaS:

Restricting the visibility of cloud-hosted data and documents is mainly the customer's responsibility. As described in the collective response to RFP requirement 8.6.4, Google is obliged to confirm there is no unauthorized access of customer data. Customers must confirm that their end users use the service according to their acceptable use policies. During the onboarding process, customers will be assisted in the configuration controls that are included to enforce acceptable use and security policies and trained on how to oversee the ongoing use of the services in scope.

8.6.8 Offeror must describe its notification process in the event of a security incident, including relating to timing, incident levels. Offeror should take into consideration that Purchasing Entities may have different notification requirements based on applicable laws and the categorization type of the data being processed or stored.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

The Unisys proposed solution includes HP ArcSight to provide a robust SIEM security platform that the Unisys Security Officer in the Project Management Office (PMO) will use to monitor security incidents in the hosting services. To offer comprehensive cross-enterprise protection, our proposed solution converges monitoring, management, compliance, analytics, and response mechanisms for enhanced IT security. By unifying Security Event Management (SEM) and Security Information Management (SIM) capabilities, SIEM provides greater visibility to IT security; enhances protection of information, assets, and processes; and helps organizations to comply with GNU Privacy Guard (GPG)-13 (Protective Monitoring) and GPG-18 (Forensic Readiness). The Unisys SIEM solution leverages our best-in-class analytics engine based on ArcSight, the Unisys Noise Cancellation Advanced Analytics Platform (UNCAAP), and the PMO-dedicated Security Officer to minimize false alarm rates and offer advanced forensics and analytics with benefits of actionable intelligence, decision support, and security strategy insights.

With new application services and process improvement, the configuration of the rule base that operates the Unisys proposed solution will be aligned with the change management process, confirming the appropriate level of monitoring is used for each monitored system. Unisys will also work with customer proprietary network information (CPNI) to monitor current threat levels in the industry and gain early intelligence on potential upcoming attack vectors and changes to best practice as driven by security professionals who work in IT. Together with an effective communication strategy, technical and procedural services will provide defense in depth across the entire hosted service.

Unisys will use the Qualys Vulnerability Scanner to monitor the IP addresses in this engagement's scope at the mutually agreed intervals. The preferred method of scanning an environment is to perform a discovery scan before completing the vulnerability scan. This type of scan will provide an accounting of systems in a specific IP range and show what ports are open, what systems were discovered, and what services may be running on those ports. Typically, a nonintrusive scan will not have adverse effects on a system, though a slight risk does exist. This risk is higher when running an intrusive scan. System owners need to understand the risks that exist before the scan is run. The mutually agreed change control processes will be followed for systems that will be scanned.

After a scan is completed, analyzed, and published, the vulnerabilities must be remediated. This effort will be led by the Security Officer, who oversees the vulnerability scanning effort. The Security Officer should use the scan reports to create tickets for risks related to a system. The Security Officer will create tickets for vulnerabilities based on priority (Critical, High, Medium, etc.). The tickets will be routed to the Unisys Support team for remediation in the timelines agreed with the State.

A facility that allows for exemptions to remediate the identified vulnerability should be in place to account for situations to which one of the following factors apply:

- System stability
- Lack of an available patch for the operating system or application involved
- Any other business reason.

Unisys will also take over and manage the existing server antivirus infrastructure to reduce disruption to the State's transition to a new security provider model in the SIEM.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the AWS customer support team to be notified of operational issues that affect the customer experience. A Service Health Dashboard is available and maintained by the AWS customer support team to alert customers to issues that may be of broad impact. The AWS Security Center is available to provide customers with security and compliance details on AWS. Customers can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to issues that affect customers.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

If Microsoft becomes aware of unlawful access to customer data stored on its equipment or in its facilities, or unauthorized access to this equipment or facilities resulting in loss, disclosure, or alteration of customer data (each a security incident), Microsoft will promptly (1) notify the customer of the security incident; (2) investigate the security incident and provide the customer with detailed information on the security incident; and (3) take reasonable steps to mitigate the effects and to minimize damage resulting from the security incident.

Notifications of security incidents will be delivered to one or more of the customer's administrators by the means Microsoft selects, including email. It is the customer's sole responsibility to confirm that its administrators maintain accurate contact information on each applicable Online Services portal. Microsoft's obligation to report or respond to a security incident in this section is not an acknowledgement by Microsoft of fault or liability for the security incident.

The customer must notify Microsoft promptly of possible misuse of its accounts or authentication credentials or security incidents related to an Online Service.

Electronic Notices

Microsoft may provide the customer with information and notices about online services electronically, including by email, at the portal for the online service, or at a website that Microsoft identifies. Notice is given as of the date that Microsoft makes the information and notices available.

Incident Response Process

Microsoft maintains a record of security breaches with a description of the breach, the period, the consequences of the breach, the name of the reporter, to whom the breach was reported, and the procedure for recovering data.

For each security breach that is a security incident, notification by Microsoft (as described in "Electronic Notices") will be made without unreasonable delay and within 30 calendar days.

Microsoft tracks, or enables the customer to track, disclosures of customer data, including what data was disclosed, to whom, and at what time.

Microsoft security personnel verify logs at least every 6 months to propose remediation efforts if necessary.

Microsoft Security Development Lifecycle (SDL)

Microsoft SDL is supported in Azure; part of that development processes is notification. Azure has security controls in place to implement threat mitigation and to assist customers with mitigating potential threats in their environments. To detect threats and prevent exploits, Microsoft maintains continuous monitoring across servers, networks, and applications. Automated alerts notify administrators of anomalous behaviors, allowing them to take corrective action on internal and external threats.

Microsoft Operations Centers (MOCs) are globally distributed and work around the clock in a “follow-the-sun” model to enable Microsoft’s cloud services to be persistently available. Each MOC is staffed with a team of incident management professionals; collectively they monitor service health, process automation, infrastructure operations, event and crisis management, and communications across the business. They are responsible for more than 500 service components and monitor the servers and devices for the services Microsoft provides. Most critically, MOCs identify and resolve service incidents and outages when things go wrong.

Google Public and Community Cloud for SaaS: If Google becomes aware of a data incident, Google will promptly notify the customer of the incident and take reasonable steps to minimize harm and secure customer data. Notifications of data incidents will be sent to the notification email address provided by the customer in connection with the agreement or, at Google’s discretion, by direct communication (e.g., by a phone call or at an in-person meeting). The customer must confirm the contact information for the Notification Email Address is current and valid and must establish third-party notification obligations. Data incidents do not include unsuccessful access attempts or similar events that do not compromise the security or privacy of customer data, including pings, port scans, denial-of-service attacks, and other network attacks on firewalls or networked systems. Data incidents also do not include accidental loss or disclosure of customer data caused by the customer’s use of the services or the customer’s loss of account authentication credentials. Google’s obligation to report or respond to a data incident will not be construed as an acknowledgement by Google of faults or liabilities for the data incident.

8.6.9 Offeror must describe and identify whether or not it has any security controls, both physical and virtual Zones of Control Architectures (ZOCA), used to isolate hosted servers.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys provides security controls through its Stealth Offering:

The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures.. By substituting traditional hardware topology for software-based cryptography, our Stealth Micro segmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making endpoints invisible to unauthorized users. Unisys Stealth conceals endpoints making them undetectable to unauthorized parties inside and outside the enterprise. It tightens access control by focusing on user identity rather than physical devices, so security moves with the user and is easier to manage. Stealth protects sensitive data in motion from potential compromise through encryption and reduces costs by allowing you to consolidate and virtualize networks, servers, and cloud architectures.

The Unisys Stealth software has a legacy of being "Ready for Government" having received the National Information Assurance Partnership's (NIAP's) coveted Evaluation Assurance Level 4+ (EAL4+) certification, a Common Criteria international standard in effect since 1999. Stealth has been sold and

installed in US DoD and allied defense forces, as well as a variety of local and federal governments and commercial entities around the world.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Amazon Virtual Private Cloud (Amazon VPC) lets customers provision a logically isolated section of the Amazon Web Services (AWS) Cloud, where customers can launch AWS resources in a virtual network that they define. Customers have complete control over their virtual networking environment, including selection of their own IP address range, creation of subnets, and configuration of route tables and network gateways.

Customers can easily customize the network configuration for their Amazon Virtual Private Cloud. For example, they can create a public-facing subnet for their web servers that has access to the Internet, and move their backend systems such as databases or application servers to a private subnet with no Internet access. Customers can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

It is important for customers to realize the scale of engineering in the Microsoft Azure strategy. The designs are not for a single data center or a few dozen data centers, so the zoning must be completed at a hyperscale level. With more than 100 data centers worldwide, Microsoft has data centers in every region, connected by one of the largest cloud networks in the world. Only a few cloud providers are individually enterprise grade, hybrid, or hyperscale; Microsoft's cloud is the only one that offers all three.

The distributed and virtual networks in Azure help to isolate each customer's private network traffic logically from traffic belonging to other customers. A customer subscription can contain two types of isolated private network (and include firewall, load-balancing, and network address translation):

- **Deployment network:** Each deployment is isolated from other deployments at the network level. Several virtual machines in a deployment are allowed to communicate with each other at private IP addresses.
- **Virtual network:** Each virtual network is isolated from other virtual networks. Several deployments (inside the same subscription) can be moved to the same VNET and allowed to communicate at private IP addresses.

By default, virtual machines inside the private network do not receive inbound traffic from outside the deployment. The administrator defines an input end point that specifies which ports on which virtual machines should receive inbound traffic initiated from outside a deployment's isolated network—enabling traffic from the Internet and other deployments or customers inside Azure.

Microsoft Azure uses many safeguards to protect customer and enterprise data. These security practices and technologies include the following:

- **Identity and access management** – Azure Active Directory helps to enable only authorized users to access customers' environments, data, and applications; it provides multifactor authentication for highly secure sign-in.
- **Encryption** – Azure uses industry-standard protocols to encrypt data as it travels between devices and Microsoft data centers, and crosses in data centers
- **Secure networks** – Azure infrastructure relies on security practices and technologies to connect virtual machines to each other and to on-premises data centers while blocking unauthorized traffic. Azure Virtual Networks extend a customer's on-premises network to the cloud over a site-to-site virtual private network (VPN). Customers can also use ExpressRoute to create a cross-premises connection when they have to use the Internet.
- **Threat management** – Microsoft Antimalware protects Azure services and virtual machines. Microsoft also uses intrusion detection, denial-of-service (DoS) attack prevention, penetration testing, data analytics, and machine learning to constantly strengthen its defense and reduce risks.

- Compliance – Microsoft complies with international and industry-specific compliance standards and participates in rigorous third-party audits that verify Microsoft’s security controls.

Physical zoning and virtual zoning of host systems are supported through Microsoft Azure’s extensive use of software-defined networks that span across the data center from edge computing, not to the interval fiber backbones.

Physical access to facilities. Microsoft limits access to facilities that house information systems that process customer data to identified authorized individuals.

Physical access to components. Microsoft maintains records of incoming and outgoing media containing customer data, including the kind of media, the authorized sender and recipients, date and time, the number of media, and the types of customer data they contain.

Protection from disruption. Microsoft uses a variety of industry standard systems to protect against loss of data due to power supply failure or line interference.

Component disposal. Microsoft uses industry-standard processes to delete customer data when it is no longer needed.

With software enablement for restricting access, security is crucial to provide customer data and privacy. Microsoft goes beyond ISO and NIST standards to maintain physical and virtually safe environments. In the Azure infrastructure, for example, CPUs, storage, and SQL services are separated to support greater zoning and security requirements.

Software-defined Networks (to support zoning)

To improve flexibility and accelerate the adoption of advanced technologies into its network, Microsoft broadly adopted software-defined networking (SDN).

SDN provides the ability to dynamically modify Microsoft’s network using automated management tools to move data and resources to an area where they are best served. In an SDN environment, Microsoft can extract and separate the application, the control plane, and the transport of the data. This allows Microsoft to insert its own APIs to gain visibility of how the data flows and gain better control, and also allows Microsoft to upgrade network performance outside the hardware refresh cycle. Microsoft’s large, geographically distributed footprint of data centers and networks enables Microsoft to be located close to its customers to reduce network latency and allow for georedundant backup and failover.

Google Public and Community Cloud for SaaS:

Google’s focus on security and protection of data is among its primary design criteria. Google data center physical security features a layered security model, including safeguards like custom-designed electronic access cards, alarms, vehicle access barriers, perimeter fencing, metal detectors, and biometrics. The data center floor features laser-beam intrusion detection. Google’s data centers are monitored 24x7 by high-resolution interior and exterior cameras that can detect and track intruders. Access logs, activity records, and camera footage are available if an incident occurs. Data centers are also routinely patrolled by experienced security guards who underwent rigorous background checks and training. As visitors come closer to the data center floor, security measures increase. Access to the data center floor is possible only in a security corridor that implements multifactor access control using security badges and biometrics. Only approved employees with specific roles may enter. Less than one percent of visitors will ever step into a Google data center.

Google operates a global, multitenant environment running the world’s second largest IP data network providing customers with a low-latency, high-performing platform that runs 24x7.

Hard disks are assets that are tracked throughout their life cycle at Google from arrival to final destruction. These disks are components that Google uses to build its own servers from other component parts, including motherboards and a hardened, highly customized version of Linux.

Google uses a proprietary storage and processing mechanism that isolates processing of data in change-rooted (chrooted) jails in a physical server. Access permissions restrict the ability for processes to interact between jails. Additionally, tenant data is stripped in chunks across many different drives; each chunk has its own access control list. This helps to confirm that data is logically isolated between customers in storage and during processing.

In addition to the processing and storage mechanisms already described, Google security controls, including least privilege rights, logging, and auditing, was implemented to be consistent with FedRAMP requirements.

8.6.10 Provide Security Technical Reference Architectures that support Infrastructure as a Service (IaaS), Software as a Service (SaaS) & Platform as a Service (PaaS).

Illustrated below are several examples from all the providers in this RFP. The exhibit titles are self-explanatory.

Exhibit 10 is an example Unisys hybrid cloud model that is applicable across all three cloud services models.



Exhibit 11 is an example Unisys hybrid cloud model that is more specific to IaaS.

Establishing a Platform that enable Client Centric IaaS Solutions and a Cloud Ecosystem




Exhibit 12 below is an example view of the full-scale hybrid enterprise environment

Integrating Clients Cloud Environment within the context of the full IT Landscape




Exhibit 13 is an example of highly available, fault-tolerant reference architecture for IaaS on AWS.




Exhibit 13. AWS Fault Tolerance and High Availability.

Exhibit 14 is an example of reference architecture for SharePoint on AWS as PaaS.




Exhibit 15 is an example of reference architecture for an e-commerce website on AWS as SaaS.




Exhibit 16 is an example of reference architecture in the Windows Azure Cloud environment representing IaaS.




Exhibit 17 is an example of reference architecture for SharePoint in the Windows Azure Cloud environment representing PaaS.

SHAREPOINT 2013: APP OVERVIEW FOR IT PRO


APP AND SOLUTION COMPARISON

	CONTENTS	HOW TO GET	HOW TO USE/INSTALL	HOW TO MANAGE
APPS FOR SHAREPOINT Are stand-alone applications that provide specific information or functionality to a SharePoint site. Apps for SharePoint are easy for users to install, use, manage, upgrade, and remove. Apps can be SharePoint-hosted (reside and execute in SharePoint) or cloud-hosted (Windows Azure or other systems), or both.	APP (.app) 	Apps can be downloaded from: - Internal App Catalog (contains apps approved and uploaded by the organization) - Public SharePoint Store	Apps can be used: - In your hosted or on-premises SharePoint site App code is installed: - On a separate web site from your other sites in its own, isolated, domain - In the cloud (cloud-based apps)	Apps can be managed and monitored by: - Site administrators - SharePoint Online Service administrators for a tenancy (sandbox) - Farm administrators
SOLUTIONS Are small to large scale packages used to customize or enhance SharePoint sites. Full trust solutions need a farm or SharePoint Online Service administrator to deploy, manage, and remove. Partial trust solutions must be installed to a sandbox.	SOLUTION (.wsp) 	Solutions can be acquired from: - Third-party developers - Your own development team	Solutions can be used: - In your hosted (sandbox solutions only) or on-premises SharePoint site Solution code is installed: - As a full trust solution (in the global assembly cache) - As a sandboxed solution	Solutions can be installed and monitored by: - SharePoint Online Service administrators for a tenancy (sandbox) - Farm administrators (sandbox or full trust)

WHY USE APPS?


APP ISOLATION

Isolation prevents unauthorized access to users' data. The isolation level determines the limit to potential breaches in security. By default, apps are deployed to their own web site in a special, isolated domain, rather than to your farm or a sandbox. Processes run in that domain.



HOSTING OPTIONS

Apps for SharePoint can be hosted by a provider, autohosted by Windows Azure, hosted by SharePoint, or a combination of these. All can leverage SharePoint components. Custom code can only be run in provider-hosted or Windows Azure autohosted options. The following hosting options are supported:



MONITOR AND MANAGE APPS

Farm administrators and SharePoint Online Service Administrators can monitor apps for SharePoint and respond to errors and issues. Site owners can manage apps for their sites.

Task	Site level	SharePoint Online Tenancy level	Farm level
Add, delete and view app details Add or delete an app in a site View details about an app	✓	✓	✓
App Catalog Configure the App Catalog Manage the App Catalog		✓	✓
Monitor apps Specify apps to monitor View install locations, manage resources		✓	✓
Errors View and troubleshoot errors		✓	✓
Licenses View and manage licenses		✓	✓

To monitor and manage apps:

- Farm administrators: General Application Settings in Central Administration
- SharePoint Online Service Administrators: Apps pages in SharePoint Online Administration Center
- Site owners: All Site Content page

APP LIFECYCLE


INSTALL AND UPDATE

Users can add and update apps for SharePoint themselves. Installation and update/upgrade can happen at any time, initiated by the user. Users can get apps for SharePoint from the SharePoint Store, or from an App Catalog that is set up for their organization.

For on-premises deployments, IT can determine whether users can download and install apps for SharePoint and can restrict access to the SharePoint Store.


FIND

A user finds an app in the SharePoint Store or App Catalog.




INSTALL

A user adds an app to his or her site. Other users do the same.



UPDATE (DEV)

The app developer releases an update to an app.



UPGRADE

Users are notified of an update and decide whether to upgrade the app on their sites.




UNINSTALL AND DISABLE/REMOVE

Users can uninstall an app completely, which removes it from their sites.

The SharePoint Store administrators can disable an app that is unsafe and remove it from the Store.


UNINSTALL

A user decides not to continue using the app and uninstalls it. The app instance and all data is deleted.




DISABLE/REMOVE (STORE)

SharePoint Store administrators discover an issue with the app and want to remove it.



Store administrators disable and remove the app. Data from the app is retained for recovery purposes.



© 2012 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at ITSPDocs@microsoft.com.



Exhibit 17. Reference Architecture for SharePoint in an Azure Cloud Environment Representing PaaS.

Exhibit 18 is an example of reference architecture for Office 365 on the Windows Azure Cloud Environment representing SaaS.





Exhibit 18. Reference Architecture for SharePoint in an Azure Cloud Environment Representing SaaS.

Exhibit 19 is an example of Google Apps reference architecture

8.6.11 Describe security procedures (background checks, foot printing logging, etc.) which are in place regarding Offoror's employees who have access to sensitive data.

Unisys Response:

Unisys and our partners mandate background verification for our associates to comply with laws of various countries. New associates are subject to a security check and verification of previous employment before they begin work, which includes a past employment check; an education check; Federal and state criminal checks; and, if required, drug testing for associates. For our cleared associates and associates working in the government space, additional screening extending to a Federal clearance is performed in accordance with the requirement descriptions.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS operates under a shared security responsibility model, where AWS is responsible for the security of the underlying cloud infrastructure and customers are responsible for securing the workloads they deploy in AWS. AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS conducts pre-employment criminal background checks, as permitted by law, for employees commensurate with their position and level of access. The AWS SOC reports provides additional details regarding the controls in place for background verification.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Pursuant to local laws, regulations, ethics and contractual constraints, Microsoft full-time employees based in the United States are required to successfully complete a standard background check as part of the hiring process. Background checks may include reviewing information on a candidate's education, employment, and criminal history.

Google Public and Community Cloud for SaaS:

Google personnel are required to conduct themselves in a manner consistent with the company's guidelines regarding confidentiality, business ethics, appropriate usage, and professional standards. Google conducts reasonably appropriate background checks to the extent legally permissible and in accordance with applicable local labor law and statutory regulations. Personnel are required to execute a confidentiality agreement and must acknowledge receipt of, and compliance with, Google's confidentiality and privacy policies. Personnel are provided with security training. Personnel handling Customer Data are required to complete additional requirements appropriate to their role (eg certifications). Google's personnel will not process Customer Data without authorization.

8.6.12 Describe the security measures and standards (i.e. NIST) which the Offeror has in place to secure the confidentiality of data at rest and in transit.

Unisys Response:

Unisys adequately secures database systems from security threats internally and externally. We established a workflow mechanism that allows specific queries to be run in the database. Only administrators have the control to specify these queries, thereby confirming that even trusted users have access to limited features of the database. Database systems are adequately hardened from threats such as SQL injections and cross-site scripting (XSS) attacks. Connections to databases are controlled using appropriate features in the database.

To achieve and maintain security across our networks, Unisys uses network controls to separate various client networks and the public network as appropriate through switch and firewall technologies. We established an access policy to govern associates who have access to networks and services. To encrypt network traffic, we use the latest cryptographic controls.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS offers you the ability to add a layer of security to your data at rest in the cloud, providing scalable and efficient encryption features. These include:

- Data encryption capabilities available in AWS storage and database services, such as EBS, S3, Glacier, Oracle RDS, SQL Server RDS, and Redshift
- Flexible key management options that allow you to choose whether to have AWS manage the encryption keys or maintain complete control over your keys
- Dedicated, hardware-based cryptographic key storage options for customers to help satisfy compliance requirements

In addition, AWS provides APIs for you to integrate encryption and data protection with any of the services you develop or deploy in an AWS environment.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft encrypts customer data at transit and at rest and provides a platform that complies with NIST standards. Industry and government regulations such as HIPAA and FedRAMP as well as international standards such as ISO 27001 lay out specific safeguards through processes and policies. Microsoft Azure and its customers share a responsibility to implement sufficient mechanisms to meet those obligations. Specifically, Microsoft provides a compliant platform for services, applications, and data; Azure customers must design and configure their cloud environment to maintain the confidentiality and integrity of their information assets.

Microsoft Azure contains many tools to safeguard data according to a customer's security and compliance needs. One of the keys to data protection in the cloud is accounting for the possible states in which data may occur, and what controls are available for that state. Specifically, this accounting is for the following states of customer data:

- At rest: This includes information storage objects, containers, and types that exist statically on physical media, be it magnetic or optical disk.
- In transit: When data is transferred between components, locations, or programs, such as over the network, across a service bus (from on-premises to cloud and vice versa, including hybrid connections such as ExpressRoute), or during an input/output process, it is thought of as being in motion. Data in transit does not always mean a communications process with a component outside a cloud service; data moves internally, also (e.g., between two virtual networks).
- In use: (or in process): Dynamic data use could be on a table kept in virtual memory, transactions in a message queue, or encryption keys in the CPU cache. Information acted on in some way by the host or guest during a process, such as real-time database queries running in active memory (instead of a page file sent out to disk), could be in different security states, depending on whether it is encrypted or decrypted, and the operator's security context.

Additionally, there are two fundamental types of data at rest:

- Data in production. Data is in a form of storage, (e.g., Azure SQL Database), and compute processes need to access that storage during production operations. For this data, encryption at rest aims to protect the data in that storage (whereas the compute aspect deals with data in use).
- Data not in production. Data is in a form of storage (e.g., a virtual hard disk [VHD], but that VHD is not in production use. For example, the data may be part of an upgrade operation, but the VHD is not yet loaded or mounted. Data encryption at rest is applicable here, but the compute aspect is not relevant to this scenario.

Google Public and Community Cloud for SaaS:

Google stores data in a multi-tenant environment on Google-owned servers. Data, the Services database and file system architecture are replicated between multiple geographically dispersed data centers. Google logically isolates data on a per End User basis at the application layer. Google logically isolates each Customer's data, and logically separates each End User's data from the data of other End Users, and data for an authenticated End User will not be displayed to another End User (unless the former End User or an Administrator allows the data to be shared). A central authentication system is used across all Services to increase uniform security of data. The Customer will be given control over specific data sharing policies. Those policies, in accordance with the functionality of the Services, will enable Customer to determine the product sharing settings applicable to End Users for specific purposes. Customer may choose to make use of certain logging capability that Google may make available via the Services, products and APIs. Customer agrees that its use of the APIs is subject to the API Terms of Use. Google agrees that changes to the APIs will not result in the degradation of the overall security of the Services.

Data centers are typically connected via high-speed private links to provide secure and fast data transfer between data centers. This is designed to prevent data from being read, copied, altered or removed without authorization during electronic transfer or transport or while being recorded onto data storage media. Google transfers data via Internet standard protocols.

8.6.13 Describe policies and procedures regarding notification to both the State and the Cardholders of a data breach, as defined in this RFP, and the mitigation of such a breach.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys maintains a well-defined incident management framework that confirms that security events are given the utmost priority. The framework's design allows it to receive manual and automated notification on various incidents, further classifying them into levels of criticality until they are resolved and a root cause analysis is done. We use technologies such as SIEM and ArcSight for incident management. A Unisys Security Operations Center (SOC) continuously monitors security events 24x7 and correlates those events with log management solutions. We notify subscribers to our SOC services of security events and breaches and their progress toward resolution.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS Customers retain the responsibility to monitor their own environment for privacy breaches.

AWS has implemented a formal, documented incident response policy and program (including instructions on how to report internal and external security incidents). The policy addresses purpose, scope, roles, responsibilities, and management commitment. Systems within AWS are extensively instrumented to monitor key operational and security metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key metrics. When a threshold is crossed, the AWS incident response process is initiated. The Amazon Incident Response team employs industry - standard diagnostic procedures to drive resolution during business - impacting events. Staff operates 24x7x365 coverage to detect incidents and manage the impact to resolution.

AWS utilizes a three-phased approach to manage incidents:

1. Activation and Notification Phase: Incidents for AWS begin with the detection of an event. This can come from several sources including:
 - Metrics and alarms - AWS maintains an exceptional situational awareness capability, most issues are rapidly detected from 24x7x365 monitoring and alarming of real time metrics and service dashboards. The majority of incidents are detected in this manner. AWS utilizes early indicator alarms to proactively identify issues that may ultimately impact Customers.
 - Trouble ticket entered by an AWS employee
 - Calls to the 24X7X365 technical support hotline. If the event meets incident criteria, then the relevant on -call support engineer will start an engagement utilizing AWS Event Management Tool system to start the engagement and page relevant program resolvers (e.g. Security team). The resolvers will perform an analysis of the incident to determine if additional resolvers should be engaged and to determine the approximate root cause.
2. Recovery Phase - the relevant resolvers will perform break fix to address the incident. Once troubleshooting, break fix and affected components are addressed, the call leader will assign next steps in terms of follow -up documentation and follow - up actions and end the call engagement.
3. Reconstitution Phase - Once the relevant fix activities are complete the call leader will declare that the recovery phase is complete. Post mortem and deep root cause analysis of the incident will be assigned to the relevant team. The results of the post mortem will be reviewed by relevant senior management and relevant actions such as design changes etc. will be captured in a Correction of Errors (COE) document and tracked to completion.

In addition to the internal communication mechanisms detailed above, AWS has also implemented various methods of external communication to support its customer base and community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "Service Health Dashboard" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact.

AWS incident management program reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft Identity Manager and intrusion detection system tools are implemented within the Azure environment. Azure uses an early warning system to support real-time analysis of security events within its operational environment. Monitoring agents and the alert and incident management system generate near real-time alerts about events that could potentially compromise the system.

The Azure Incident Response process follows five main phases:

Identification – System and security alerts may be harvested, correlated, and analyzed. Events are investigated by Microsoft operational and security organizations. If an event indicates a security issue, the incident is assigned a severity classification and appropriately escalated within Microsoft. This escalation will include product, security, and engineering specialists.

Containment – The escalation team evaluates the scope and impact of an incident. The immediate priority of the escalation team is to ensure the incident is contained and data is safe. The escalation team forms the response, performs appropriate testing, and implements changes. In the case where in-depth investigation is required, content is collected from the subject systems using best-of-breed forensic software and industry best practices.

Eradication – After the situation is contained, the escalation team moves toward eradicating any damage caused by the security breach, and identifies the root cause for why the security issue occurred. If a vulnerability is determined, the escalation team reports the issue to product engineering.

Recovery – During recovery, software or configuration updates are applied to the system and services are returned to a full working capacity.

Lessons Learned – Each security incident is analyzed to ensure the appropriate mitigations applied to protect against future reoccurrence.

Google Public and Community Cloud for SaaS:

If Google becomes aware of a Data Incident, Google will promptly notify the customer of the Data Incident, and take reasonable steps to minimize harm and secure Customer Data. Notification(s) of any Data Incident(s) will be delivered to the Notification Email Address provided by Customer in connection with the Agreement or, at Google's discretion, by direct communication (e.g., by phone call or an in-person meeting). Customer acknowledges that it is solely responsible for ensuring the contact information given for purposes of the Notification Email Address is current and valid, and for fulfilling any third party notification obligations. Customer agrees that "Data Incidents" do not include: (i) unsuccessful access attempts or similar events that do not compromise the security or privacy of Customer Data, including pings, port scans, denial of service attacks and other network attacks on firewalls or networked systems; or (ii) accidental loss or disclosure of Customer Data caused by Customer's use of the Services or Customer's loss of account authentication credentials. Google's obligation to report or respond to a Data Incident under this Section will not be construed as an acknowledgement by Google of any fault or liability with respect to the Data Incident.

8.7 (E) MIGRATION AND REDEPLOYMENT PLAN

8.7.1 Offeror must describe how it manages the end of life activities of closing down a service to a Purchasing Entity and safely deprovisioning it before the Offeror is no longer contractually obligated to maintain the service, include planned and unplanned activities. An Offeror's response should include detail on how an Offeror maintains security of the data during this phase of an SLA, if the Offeror provides for redundancy during migration, and how portable the data is during migration.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

The client is primarily responsible for creation, usage, maintenance and destruction of its content when they are subscribed to services offered within the Cloud Account that Unisys would have provisioned.. In case when the subscription has ended, Unisys will provide a moratorium time period until the customer can safely backup or migrate their existing data, application stack and workloads as required. Once customer has completed with his backup or migration activities, Unisys will deprovision the Cloud Account.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

You (the client) are solely responsible for the development, content, operation, maintenance, and use of Your Content. For example, you are solely responsible for the technical operation of Your Content, including ensuring that calls you make to any Service are compatible with then-current APIs for that

Service; compliance of Your Content with the Acceptable Use Policy, the other Policies, and the law; any claims relating to Your Content; and properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Content violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act.

During the 30 days following termination (of the service) we will not erase any of Your Content as a result of the termination. You may retrieve Your Content from the Services only if you have paid any charges for any post-termination use of the Service Offerings and all other amounts due. We will provide you with the same post-termination data retrieval assistance that we generally make available to all customers.

Any additional post-termination assistance from us is subject to mutual agreement by you and us.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

During a customer's subscription, the customer will be able to access and extract customer data stored in each online service. Except for free trials, Microsoft retains customer data stored in the online service in a limited function account for 90 days after expiration or termination of the customer's subscription so that the customer can extract the data. After the 90-day retention period ends, Microsoft disables the customer's account and deletes the customer data.

Google Public and Community Cloud for SaaS:

During the Term, Google will provide Customer with the ability to correct, block, export and delete Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. Once Customer deletes Customer Data via the Services such that the Customer Data cannot be recovered by Customer (the "Customer-Deleted Data"), Google will delete the Customer-Deleted Data within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so. On the expiry or termination of the Agreement (or, if applicable on expiry of any post-termination period during which Google may agree to continue providing access to the Services), after a recovery period of up to 30 days following such expiry or termination, Google will thereafter delete the Customer-Deleted Data within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

8.7.2 Offeror must describe how it intends to provide an orderly return of data back to the Purchasing Entity, include any description in your SLA that describes the return of data to a customer.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

Unisys provides client the ability to download or migrate their data and workloads as and when a service has been terminated. Unisys also provides Cloud Application Migration Services that starts with a complete and analytical view of your application portfolio and cloud/data center capabilities. Using a set of criteria refined during customer engagements, Unisys can help you narrow the application pool of various cloud portable candidates and then provide the detailed assessment of costs and other factors that go into a cloud migration decision. Once the decision has been made, Unisys can help you with the planning and implementation of the application migration process including on-boarding of applications to AWS and other clouds.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

Amazon offers a variety of ways to transfer data in and out of AWS.

Direct Connect:

You can use a dedicated physical connection with AWS Direct Connect to accelerate network transfers between your datacenters and ours.

Using Direct Connect, you can establish a dedicated network connection between your network and one of the Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be

partitioned into multiple virtual interfaces. This approach means you can use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon Elastic Compute Cloud (EC2) instances running within an Amazon Virtual Private Cloud (Amazon VPC) using private IP space. At the same time, you can maintain network separation between the public and private environments. You can reconfigure virtual interfaces at any time to meet your changing needs.

AWS Import/Export Snowball:

Snowball is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of AWS. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, and secure, and can cost as little as one-fifth of the expense of using high-speed Internet for the same transfer.

AWS Import/Export Disk:

AWS Import/Export Disk accelerates moving large amounts of data into and out of the AWS cloud by using portable storage devices for transport. Disk transfers your data directly onto and off of your storage devices using the Amazon high-speed internal network and bypassing the Internet. For gigabyte-scale data sets that would travel at less than 10 Mbps, Disk is often faster than Internet transfer and more cost-effective than upgrading your connectivity.

Storage Gateways:

The AWS Storage Gateway service simplifies on-premises adoption of AWS storage. Your existing applications use industry-standard storage protocols to connect to a software appliance which stores your data in Amazon S3 and Amazon Glacier. With Storage Gateway, Data is compressed and securely transferred to AWS. SAN configurations offer stored or cached devices with point-in-time backups as Amazon Elastic Block Store (EBS) snapshots. Virtual tape library (VTL) configuration works with your existing backup software for cost-effective backup in Amazon S3 and long-term archival in Amazon Glacier.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: A customer's use of an Import/Export Service is conditioned on its compliance with the instructions provided by Microsoft for the preparation, treatment, and shipment of physical media containing its data (storage media). The customer is solely responsible for verifying the storage media and data are provided in compliance with laws and regulations. Microsoft has no duty toward the storage media and no liability for lost, damaged, or destroyed storage media. Storage media shipped to Microsoft must be shipped DAP Microsoft DCS Data Center (INCOTERMS 2010). Storage media shipped to the customer is shipped DAP Customer Dock (INCOTERMS 2010).

Google Public and Community Cloud for SaaS:

During the Term, Google will make available to Customer the Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. To the extent Customer, in its use and administration of the Services during the Term, does not have the ability to amend or delete Customer Data (as required by applicable law), or migrate Customer Data to another system or service provider, Google will, at Customer's reasonable expense, comply with any reasonable requests by Customer to assist in facilitating such actions to the extent Google is legally permitted to do so and has reasonable access to the relevant Customer Data.

8.8 (E) SERVICE OR DATA RECOVERY

8.8.1 Describe how you would respond to the following situations; include any contingency plan or policy.

- a. Extended downtime.**
- b. Suffers an unrecoverable loss of data.**

- c. *Offeror experiences a system failure.*
- d. *Ability to recover and restore data within 4 business hours in the event of a severe system outage.*
- e. *Describe your Recovery Point Objective (RPO) and Recovery Time Objective (RTO).*
- 8.8.2 *Describe your methodologies for the following backup and restore services:*
 - a. *Method of data backups*
 - b. *Method of server image backups*
 - c. *Digital location of backup storage (secondary storage, tape, etc.)*
 - d. *Alternate data center strategies for primary data centers in the continental United States.*

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

The advanced options of the Unisys Cloud Provisioning System allow users to gain significantly higher levels of default capability for backup and recovery operation.

Automation includes automatic detection of server failures. In a server failure, the Unisys Cloud Provisioning System automatically launches a new server based on the machine image configured for the server's group. When the replacement server comes on line, the Unisys Cloud Provisioning System reattaches volumes that were orphaned during the failure and restarts installed services.

For the Unisys Cloud Provisioning System, backups of participant applications and data are automated at two different levels. First, the Unisys Cloud Provisioning System calls scripts to run backups and uploads the backups to a cloud storage provider such as Amazon S3. For example, the State's scripts might create a compressed backup of its MySQL or SQL Server database. Second, in the Amazon cloud, the Unisys Cloud Provisioning System creates a complete file system backup by taking snapshots of attached EBS volumes. Both types of backups are run at custom intervals. Snapshots of images and data are captured on a user-defined schedule, can be moved to supported third-party cloud storage for increased disaster avoidance, and can be encrypted for increased security.

The Unisys Cloud Provisioning System also handles the creation and configuration of firewalls for the layers of a deployment. To create firewall rules, the Unisys Cloud Provisioning System opens access to database servers from application servers on the private IP address of the application servers. If load balancing is used, the application servers are reachable only through the load balancer to their respective private IP addresses.

The Unisys Cloud Provisioning System automatically scales the number of servers in a server group up or down according to a client's parameters. The Unisys Cloud Provisioning System periodically polls the health of the deployed services and reports the status inside the Unisys Cloud Provisioning System console.

Any time there is a state change (loss of server, backup failure, etc.), the Unisys Cloud Provisioning System sends alerts at the console and by email.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

The AWS platform enables a lightweight approach to backup and recovery due in part to the following characteristics.

- Computers are now virtual abstract resources instantiated by code; they are not hardware based.
- Capacity is available at an incremental cost instead of an upfront cost.
- Resource provisioning takes place in minutes, lending itself to real-time configuration.
- Server images are available on demand, can be maintained by an organization, and can be activated immediately.

These characteristics offer customers opportunities to recover deleted or corrupted data with less infrastructure overhead.

Protecting Configurations Instead of Servers


The Amazon Elastic Compute Cloud (Amazon EC2) service enables the backup and recovery of a standard server, such as a web server or an application server, so that customers can focus on protecting their configuration and the state of data instead of the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, and temporary files. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an Amazon Machine Image (AMI) and can connect to existing storage volumes; for example, Amazon Elastic Block Store (Amazon EBS). Additionally, when launching a new instance, it is possible to pass user data to the instance that can be accessed internally as dynamic configuration parameters.


A sample workflow is as follows.

- Launch a new instance of a web server, passing to it the identity of the web server and the security credentials required for initial setup. The instance is based upon a prebuilt AMI that contains the operating system and relevant web server application (e.g., Apache or IIS).
- Upon startup, a boot script accesses a designated and secured Amazon Simple Storage Service (Amazon S3) bucket that contains the specified configuration files.
- The configuration file contains various instructions for setting up the server (e.g., web server parameters, locations of related servers, additional software to install, and patch updates).
- The server executes the specified configuration and is ready for service. An Open Source tool for performing this process called cloud-init is already installed on Amazon Linux AMIs and is available for several other Linux distributions.

Exhibit 20 depicts a traditional backup approach, and Exhibit 21 depicts an Amazon EC2 backup approach.

Exhibit 20. Traditional Backup Approach.

UT008

Exhibit 21 Amazon EC2 Backup Approach.

UT009

In the sample workflow, there is no need to back up the server itself. The relevant configuration is in the combination of the AMI and the configuration files. The only components that require backup and recovery are the AMI and configuration files.

Amazon Machine Image (AMI)

AMIs that customers register are automatically stored in their account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable. This means that the underlying storage mechanism for the AMIs is protected from multiple failure scenarios.

It is also possible to share AMIs between separate AWS accounts. Consequently, customers can create totally independent copies of the AMI by doing the following:

- Sharing the original AMI to another specified AWS account controlled by the customer
- Starting a new instance based upon the shared AMI
- Creating a new AMI from that running instance.

The new AMI is then stored in the second account and is an independent copy of the original AMI. Customers can also create multiple copies of the AMI in the same account.

Configuration Files

Customers use a variety of version management approaches for configuration files; they can follow the same regime for the files used to configure their Amazon EC2 instances. For example, a customer could store different versions of configuration files in designated locations and securely control them like other code. That customer could then use the appropriate backup cycle (daily, weekly, or monthly) to back up these code repositories and snapshots to protected locations. Additionally, customers can use Amazon S3 to store their configuration files, taking advantage of the service's durability, in addition to backing up the files regularly to an alternate location.

Database and File Servers

Backing up data for database and file servers differs from the web and application layers. Database and file servers contain larger amounts of business data (tens of gigabytes to multiple terabytes) that must be retained and always protected. For this data, customers can leverage efficient data movement techniques such as snapshots to create backups that are fast, reliable, and space efficient.

For databases that are built on RAID sets of Amazon EBS volumes (and have total storage of less than 1 TB), an alternate backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. Although the destination Amazon EBS volume will have slower performance, it is not being used for data access and can be easily snapshotted to Amazon S3 using the Amazon EBS snapshot capability.

Disaster Recovery

The AWS cloud supports many popular disaster recovery architectures from “pilot light” environments that are ready to scale up at a moment's notice to “hot standby” environments that enable rapid failover. With data centers in 12 regions around the world (4 in the United States), AWS provides a set of cloud-based disaster recovery services that enable rapid recovery of IT infrastructure and data.

General Disaster Recovery, COOP, and Backup Requirements and Issues

A traditional disaster recovery approach has some of the following needs and requirements:

- Facilities to house additional infrastructure, including power and cooling
- Security to maintain the physical protection of assets
- Suitable capacity to scale the environment
- Support for repairing, replacing, and refreshing the infrastructure

- Contractual agreements with an Internet Service Provider (ISP) to provide Internet connectivity that can sustain bandwidth utilization for the environment under a full load
- Network infrastructure such as firewalls, routers, switches, and load balancers
- Enough server capacity to run mission-critical services, including storage appliances for the supporting data, and servers to run applications and backend services such as user authentication, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), monitoring, and alerting.

AWS Capabilities for Disaster Recovery, COOP, and Backup Solutions

With AWS, customers can eliminate the need for additional physical infrastructure, offsite data replication, and upkeep of spare capacity. AWS uses distinct and geographically diverse Availability Zones (AZs) that are engineered to be isolated from failures in other AZs. This innovative and unique AWS feature enables customers to protect applications from the failure of a single location, resulting in significant cost savings and increased agility to change and optimize resources during a DR scenario.


AWS offers the following high-level disaster recovery capabilities:

- **Fast Performance:** Fast, disk-based storage and retrieval of files
- **No Tape:** Eliminated costs of transporting, storing, and retrieving tape media and associated tape backup software
- **Compliance:** Minimized downtime to avoid breaching Service Level Agreements (SLAs)
- **Elasticity:** The ability to add data quickly and delete expired data easily, without handling media
- **Security:** Secure and durable cloud disaster recovery platform with industry-recognized certifications and audits
- **Partners:** AWS solution providers and system integration partners to help with deployments.

Solution Use Cases

AWS can enable customers to cost-effectively operate several disaster recovery strategies. **Exhibit 22** shows a spectrum of scenarios—“backup & restore,” “pilot light,” “warm standby,” and “multi-site”—arranged by how quickly a system can be available to users after a disaster recovery event.

Exhibit 22. Spectrum of Disaster Recovery Options.



Each disaster recovery option is discussed in more detail as follows.

- **Backup and Restore:** In most traditional environments, data is backed up to tape and sent off site regularly. Recovery time will be the longest using this method; lack of automation leads to increased costs. Using Amazon Simple Storage Service (Amazon S3) is ideal for backup data because it is designed to provide 99.99999999 percent durability

Amazon S3 is designed to provide 99.99999999 percent durability of objects over a given year. This durability level corresponds to an average annual expected loss of 0.00000001 percent of objects.

of objects over a given year. Transferring data to and from Amazon S3 is typically done over the network; therefore, it is accessible from anywhere. Additionally, with AWS Storage Gateway, customers can automatically back up on-premises data to Amazon S3.

- **Pilot Light for Simple Recovery into AWS Warm Standby Solution:** The idea of the pilot light is an analogy that comes from the gas heater. In a gas heater, a small idle flame that is always on can quickly ignite the entire furnace to heat a house as needed. This scenario is analogous to a backup-and-restore scenario; however, customers must confirm that they have the most critical core elements of their system already configured and running in AWS (the pilot light). When the time comes for recovery, customers would rapidly provision a full-scale production environment around the critical core.
- **Warm Standby Solution in AWS:** The term “warm standby” is used to describe a disaster recovery scenario in which a scaled-down version of a fully functional environment is always running in the cloud. It further decreases recovery time because, in this scenario, some services are always running. By identifying business-critical systems, customers can fully duplicate these systems on AWS and have them always on.
- **Multi-Site Solution Deployed on AWS and On Site:** A multisite solution runs in AWS and on a customer’s existing on-premises infrastructure in an active/active configuration. During a disaster, an organization can simply send the traffic to AWS servers, which can scale to handle the full production load.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Azure Backup can now back up your on-premises application workloads, including Microsoft SQL Server, Hyper-V virtual machines, Microsoft SharePoint, and Microsoft Exchange. You can back up your applications to a local disk or to Azure, allowing you to eliminate local tape libraries and leverage the unlimited storage capability of Azure. You can also manage all your on-premises backups from a single user interface. Backup continues to support backups of your production IaaS virtual machines in Azure and to help protect your Windows client data and your shared files and folders.

System Center Data Protection Manager is an option for on-premises, Azure, or Cloud Only backup and recovery.

Azure Site Recovery: ASR’s enhanced VMware to Azure scenario is now Generally Available. This GA release, among other enhancements, is designed to help customers benefit from the following key functionality:

Elimination of IaaS-based replication and orchestration components/appliance

MSI-based unified setup of on-premises components, which significantly reduces the time and complexity to onboard to the scenario

Non-disruptive disaster recovery testing with Test Failover

ASR-integrated failback experience without vContinuum, with support for alternate location recovery, and original location recovery

Disk-based replication from source machines, and driver installation without needing a source reboot

Multi-VM Application and Crash-Consistent Replication for Windows and Linux

Migration of protected machines from the in-market – Legacy – VMware to Azure scenario to the Enhanced VMware to Azure scenario

Enterprise-grade enhancements such as support for FQDNs, custom ports, and installation paths

Support for CentOS & RHEL 6.7, vCenter Server 6.0

For Microsoft Government Community Cloud (GCC) Services as defined in Microsoft's service terms and conditions, customer content is stored at rest in the United States. For the GCC versions of Exchange Online, SharePoint Online, Skype for Business, and Dynamics CRM Online, customer content is stored in encrypted format. In the GCC version of Azure Core Services, customers can encrypt nonpublic customer content. For the non-GCC (public) versions of the equivalent services, as well as for Microsoft Intune Online Services, certain types of customer content are stored at rest in the United States, if users in the United States set them up. The terms and conditions governing where customer data is stored can be found in the Microsoft Online Services Terms. For the non-GCC version of Azure Core Services, customers are given the choice of Microsoft worldwide data centers in which they can store and process data. For purposes of GCC Services, customer content means the subset of customer data created by users. For Office 365 Services, customer content at least includes Exchange Online mailbox content (email body, calendar entries, and the content of email attachments), SharePoint Online site content and the files stored at that site, and Skype for Business Online archived conversations. For Microsoft Dynamics CRM Online Services, customer content consists of the entities of customer data managed by Microsoft Dynamics CRM Online Services.

Google Public and Community Cloud for SaaS: Google's data storage methodology ensures that data is replicated across multiple servers and data centers. If a single server begins to exhibit signs of poor performance the customer data will be seamlessly moved to a different health server in a manner that is seamless to the end user. There is no likely instance of an extended downtime.

The likelihood of an unrecoverable loss of data would be more likely due to an end user omission or bad act. Customers have the ability, via Google Apps Vault, to set retention policies on email data and have Data Recovery tools in the Admin Panel that will allow for end user deleted data to be restored up to 25 days after it was deleted.

Systems fail all the time but Google designed its infrastructure to shield customers from that with the methods that will move customer data away from hardware that having issues.

A system outage would not necessarily lead to data loss. Google's infrastructure would be able to serve up the data from any of the redundant locations that was used as part of the real time data replication.

Google has both objectives set to zero based on the inherent design of the infrastructure.

The following methodologies are used for backup and restore services:

Online tape backups

Standard server backup techniques

Encrypted tapes stored on site in locked vaults. Some tapes may be stored with an offsite data storage company.

Google leverages all the data centers all the time to ensure rapid and real time data replication.

8.9 (E) DATA PROTECTION

8.9.1 *Specify standard encryption technologies and options to protect sensitive data, depending on the particular service model that you intend to provide under this Master Agreement, while in transit or at rest.*

Unisys Response:

All of the offerings in this proposal that are provided by our partners use SSL/TLS encryption for browser-based transactions. For non-browser based access, Unisys provides Secure Shell, Secure FTP based access. Additionally, various security APIs with documentation are available for customers to transact with our environments securely.

8.9.2 *Describe whether or not it is willing to sign relevant and applicable Business Associate Agreement or any other agreement that may be necessary to protect data with a Purchasing Entity.*

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS

After an opportunity for review, Unisys is willing to sign a relevant and applicable Business Associate Agreement or any other agreement that may be necessary to protect data with a Purchasing Entity.

8.9.3 *Offeror must describe how it will only use data for purposes defined in the Master Agreement, participating addendum, or related service level agreement. Offeror shall not use the government data or government related data for any other purpose including but not limited to data mining. Offeror or its subcontractors shall not resell nor otherwise redistribute information gained from its access to the data received as a result of this RFP.*

Unisys Response: Unisys is contractually obliged to safeguard a client's data by signing a confidentiality agreement; we cascade this clause to our partners as well. When joining Unisys, associates must sign a Non-Disclosure Agreement. We reiterate the importance of maintaining the confidentiality of client data in annual security awareness training. Unisys understands that breaches of confidential agreements are liable to legal action.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Customers maintain ownership of their customer content and select which AWS services process, store and host their customer content. AWS does not access or use customer content for any purpose other than as legally required and for maintaining the AWS services and providing them to our customers and their end users. We never use customer content or derive information from it for marketing or advertising.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Customer data will be used only to provide a Purchasing Entity with Online Services, including purposes compatible with providing them. Microsoft does not use customer data or derive information from it for advertising or similar commercial purposes. As between the parties, the Purchasing Entity retains the right, title, and interest in customer data. Microsoft acquires no rights in customer data other than the rights that the customer grants to Microsoft to provide the customer with Online Services. This paragraph does not affect Microsoft's rights in software or Microsoft Online Services licenses to the Purchasing Entity. Microsoft uses data mining only to provide the cloud services, subject to the above-mentioned restrictions. Microsoft does not use data mining for unrelated commercial purposes, advertising or advertising-related purposes, or for purposes other than security or service delivery analysis that are not explicitly authorized.

Google Public and Community Cloud for SaaS: Google Apps customers own their data. Google does not sell customer data to third parties. Google offers a detailed Data Processing Amendment that describes our commitment to protecting customer data. Google will not process your data for any purpose other than to fulfill our contractual obligations. Further, Google commits to deleting data from our systems within 180 days of the customer deleting it in our services. Finally, we provide tools to make it easy for the customer to take their data if they choose to stop using our services altogether, without penalty or additional cost imposed by Google.

8.10 (E) SERVICE LEVEL AGREEMENTS

8.10.1 *Offeror must describe whether your sample Service Level Agreement is negotiable. If not describe how it benefits purchasing entity's not to negotiate your Service Level Agreement.*

8.10.2 *Offeror, as part of its proposal, must provide a sample of its Service Level Agreement, which should define the performance and other operating parameters in which the infrastructure must operate to meet IT System and Purchasing Entity's requirements.*

The sample SLA's cover a broad range of services. They are not negotiable. We are constantly evaluating and updating our SLA's to meet the needs of our customers. We are confident that our SLA's meet the overall requirements of the participating entity. The sample SLA's are available in the appendix sections itemized below.

AWS SLA is available in the Appendix 8.

Microsoft SLA is available in the Appendix 9.

Google Apps SLA is available in the Appendix 10.

8.11 (E) DATA DISPOSAL

Specify your data disposal procedures and policies and destruction confirmation process.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys data disposal and retention policies guide the requirement to maintain information for a predetermined time. Additionally, we develop procedures for archiving the information and guidelines for destroying it. We customize storage and disposal according to a client's requirements or Unisys defined policies, with a high priority on the client's requirements. The Unisys standard for media handling lists the types of disposal mechanisms to be followed for various media used for data storage. Disposal mechanisms include simple erasure, secure erasure, physical destruction, and the use of commercial data destruction services.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: It is important for customers to understand the following important basics of data ownership and management in the cloud shared responsibility model:

- Customers continue to own their data.
- Customers choose the geographic locations in which to store their data—it does not move unless the customer decides to move it.
- Customers can download or delete their data whenever they like.
- Customers should consider the sensitivity of their data, and decide if and how to encrypt the data while it is in transit and at rest.

AWS provides customers with the ability to delete their data. However, AWS customers retain control and ownership of their data; customers must manage their data. When a storage device reaches the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from exposure to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M, National Industrial Security Program Operating Manual or NIST 800-88, Guidelines for Media Sanitization to destroy data as part of the decommissioning process. Decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft uses best practice procedures and a wiping solution that is NIST 800-88 compliant. For hard drives that cannot be wiped, Microsoft uses a destruction process that destroys it (shredding) and renders the recovery of information impossible (e.g., disintegrating, shredding, pulverizing, or incinerating). The asset type determines the appropriate means of disposal. Microsoft retains records of the destruction. Microsoft Azure services use approved media storage and disposal management services. Paper documents are destroyed by approved means at the predetermined end-of-life cycle.

Data destruction techniques vary, depending on the type of data object being destroyed, whether it be subscriptions, storage, virtual machines, or databases. In Azure's multitenant environment, Microsoft pays careful attention to prevent one customer's data from leaking into another customer's data. When a customer deletes data, no other customer (usually including the customer that once owned the data) can gain access to that deleted data.

Azure follows NIST 800-88, Guidelines on Media Sanitization, which addresses the principal concern of preventing from unintentional release. These guidelines encompass electronic and physical sanitization.

Microsoft restricts access to administer Microsoft Azure Active Directory portal and services in accordance with assigned privileges and the associated subscription of the customer account. When approved for deletion, Microsoft follows the procedures outlined in DSI-07.1 to remove customer data.

Google Public and Community Cloud for SaaS: When retired from Google's systems, hard disks containing customer information are subjected to a data destruction process before leaving Google's premises. First, disks are logically wiped by authorized individuals using a process approved by the Google Security Team. Then, another authorized individual performs a second inspection to confirm that the disk has been successfully wiped. These erase results are logged by the drive's serial number for

tracking. Finally, the erased drive is released to inventory for reuse and redeployment. If the drive cannot be erased due to hardware failure, it is securely stored until it can be physically destroyed. Each facility is audited on a weekly basis to monitor compliance with the disk erase policy.

8.12 (E) PERFORMANCE MEASURES AND REPORTING

8.12.1 Describe your ability to guarantee reliability and uptime greater than 99.5%. Additional points will be awarded for 99.9% or greater availability.

Unisys Response:

Unisys and our partners in this proposal guarantee availability of 99.5 percent on some service and 99.9 percent availability on other services. The uptime availabilities are contained in the various SLA's provided with this proposal. We are able to provide this availability from our fully redundant, fault-tolerant, geographically dispersed data centers that are connected with fully redundant, fault-tolerant, high-speed connections. Additionally, our talented and skilled 24x365 staff and our best practices and processes can provide this reliability.

8.12.2 Provide your standard uptime service and related Service Level Agreement (SLA) criteria.

Unisys Response:

The SLAs in **Appendix 8, 9, and 10** of this proposal provide the SLA uptimes for our offerings.

8.12.3 Specify and provide the process to be used for the participating entity to call/contact you for support, who will be providing the support, and describe the basis of availability.

Unisys Response:

Unisys enables our clients to directly contact our provider's subsystems. Issues raised directly with the provider will be handled by the SLAs provided by the provider. Unisys also identified a contract manager. Our contract manager will act as the single point of contact (SPOC) for each Participating Entity or specific departments. Each client can directly reach the Unisys SPOC for responses to queries.

8.12.4 Describe the consequences/SLA remedies if the Respondent fails to meet incident response time and incident fix time.

Unisys Response:

Details are provided in the SLAs in **Appendix 8, 9, and 10** of this proposal.

8.12.5 Describe the firm's procedures and schedules for any planned downtime.

Unisys Response:

Unisys and our partners will communicate planned downtime to the user community in advance in email notifications and email reminders.

8.12.6 Describe the consequences/SLA remedies if disaster recovery metrics are not met.

Unisys Response:

Details are covered in the SLAs provided in **Appendix 8, 9 and 10** of this proposal.

8.12.7 Provide a sample of performance reports and specify if they are available over the Web and if they are real-time statistics or batch statistics.

Unisys Response:

For customized cloud solutions surrounding private and hybrid implementations we can customize the frequency and granularity of the performance reports as required. For the public and community cloud deployments, we have provided the sample reports in the exhibits below. Please note the exhibits below are only a sample. Our platform interface is robust, extensible and customizable. Furthermore, the platform offers several documented opportunities to integrate our platform with third party tools for data ingestion, manipulation and decision making.

Exhibit 23 AWS Sample Performance Report




Exhibit 24 Microsoft Sample Performance Report




Exhibit 25 Google Sample Performance Report

orangescape Tasks Reports Admin Masters Upgrade dinesh.varadharajan@orangescape.com									
Reports -> Transport Admin									
Dashboard	Process Step performance summary								Last 12 Months
KPI REPORTS	Process Step	No of Requests	Allocated Time	Avg Time	Max Time	Min Time	SLA Breached	Requests Queried	Requests Queried
Workflow Summary	Manager Approval	64	2 days	1 day 16 hours	26 days 4 hours	44 seconds	23	37	13
Workflow Performance	Finance Approval	30	2 days	1 day 6 hours	23 days 4 hours	8 minutes	2	4	1
Process Step Performance	Head of Finance Approval	74	2 days	1 day 19 hours	42 days 6 hours	45 seconds	14	4	2
Participants Performance	Group Manager Approval	76	2 days	3 days 10 hours	26 days 1 hour	2 minutes	29	5	3
ADHOC REPORTS	Head of Admin	66	2 days	20 hours	14 days 17 hours	2 minutes	11	6	3
You haven't created any adhoc reports. Create your first report by adding a shortcut from within the dashboard.	Payment Processing	75	2 days	14 hours	8 days 21 hours	28 seconds	8	1	1
	Payment Approval	70	2 days	3 days 6 hours	22 days 2 hours	14 minutes	39	1	1
	Payment Completion	19	2 days	7 hours	5 days 9 hours	1 minute	3	1	0
	Payment Acknowledgement	66	2 days	6 hours	4 days	23 seconds	3	0	0
	Process Audit and Comments	52	2 days	2 days	17 days 23 hours	2 minutes	13	19	10

8.12.8 Ability to print historical, statistical, and usage reports locally.**Unisys Response:**

Unisys provides each of the Participating Entities access to a Billing and Consumption portal. This portal will have the ability to provide a tabular view of a client's current as well as past billing and usage details of cloud resources on a real-time basis. It can be accessed by browser and provides clients the ability to print historical, statistical, and usage reports locally onto a client's machine in an excel or pdf format. Clients will have the ability to filter information of billing and usage and obtain granular information at a service, instance or a pre-defined group level.

8.12.9 Offeror must describe whether or not its on-demand deployment is supported 24x365.**Unisys Response:**

The offerings in this proposal provide on-demand deployment 24x365.

8.12.10 Offeror must describe its scale-up and scale-down, and whether it is available 24x365.**Unisys Response:**

The offerings in this proposal provide scaling 24x365.

8.13 (E) CLOUD SECURITY ALLIANCE

Describe your level of disclosure with CSA Star Registry for each Solution offered.

- Completion of a CSA STAR Self-Assessment, as described in Section 5.5.3.**
- Completion of Exhibits 1 and 2 to Attachment B.**
- Completion of a CSA STAR Attestation, Certification, or Assessment.**
- Completion CSA STAR Continuous Monitoring.**

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Appendix 3

See Appendix 3

According to the CSA definitions, AWS aligns with Level 2 in the following determinations in its third-party audits for SOC and ISO:

- Level 2 Attestation is based on SOC 2, which can be requested under the NDA. The SOC 2 report audit attests that a third-party auditor validated AWS to confirm that AWS control objectives are appropriately designed and operating effectively.
- Level 2 Certification is based on ISO 27001:2005. The AWS ISO 27001:2005.

The AWS self-assessed assertions in the CSA STAR Registry Self-Assessment are backed by independent third-party audits across multiple compliance programs. AWS continues to assert that it raises the bar on CSA's attestation and certification program.

Per the CSA website, CSA Level 3 Continuous Monitoring is still under development. AWS has implemented and documented a Continuous Monitoring Plan which defines AWS' approach to conducting continuous monitoring with its authorizing officials within the FedRAMP Security Assessment Framework. It is based on the continuous monitoring process described in NIST SP 800-137, Information Security Continuous Monitoring for Federal Information Systems and Organization, and has been reviewed and validated by a third-party assessor as part of our annual FedRAMP Assessment. It is made available to customers within the AWS FedRAMP Package which can be obtained under NDA through <https://aws.amazon.com/compliance/contact/>

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft is on the Board of Directors of the CSA. All the questionnaires are available in the appendix section as outlined below.

Appendix 4 – Microsoft O365 CCM

See Appendix 4

Appendix 5 – Microsoft Azure CCM

See Appendix 5

Appendix 6 – Microsoft Dynamics CCM

See Appendix 6

Google Public and Community Cloud for SaaS:

Appendix 7 – Google CAIQ

Google has completed the CAIQ and it is available in the Appendix

See Appendix 7

8.14 (E) SERVICE PROVISIONING

8.14.1 Describe in detail how your firm processes emergency or rush services implementation requests by a Purchasing Entity.

Unisys Response: The standard process and the lead times for onboarding a Participating Entity are as follows:

- Assuming the Participating Entity has a Participating Addendum (PA) in place - Upon the qualification and the understanding of the requirements from the Participating Entity Unisys will deliver a statement of work (SOW) within 5 business days.
- Once the Participating Entity sign's and returns the SOW, Unisys will provision the necessary account(s) in 10 business days and provide the required credentials to the Participating Entity so that they can start using the service.

As an incumbent, Unisys has the proven past performance in expediting "rush" service requests. We have delivered the above set of onboarding in less than half the time and we can do even better. It depends on the nature of the request and cooperation and participation from the participating entity.

Service provisioning once the customer has been on-boarded, is delivered on a self-service basis. Such provisioning services are available 24X365. Participating Entities can provision all services available via the Unisys portal or the AWS EC2 portal or the Microsoft Azure platform or the Google Cloud Platform within 5-10 minutes or less.

8.14.2 Describe in detail the standard lead-time for provisioning your Solutions.

Unisys Response: The standard process and the lead times for onboarding a Participating Entity are as follows:

- Assuming the Participating Entity has a Participating Addendum (PA) in place - Upon the qualification and the understanding of the requirements from the Participating Entity Unisys will deliver a statement of work (SOW) within 5 business days.
- Once the Participating Entity sign's and returns the SOW, Unisys will provision the necessary account(s) in 10 business days and provide the required credentials to the Participating Entity so that they can start using the service.

Once a participating entity has been on boarded, provisioning of services requires no lead time. It is based on a self-provisioning system. As an example all services available via the Unisys portal or the AWS EC2 portal or the Microsoft Azure platform or the Google platform can be self-provisioned within 5-10 minutes or less.

For Private and Hybrid cloud solutions that are customized Unisys will provide a detailed project plan based on the scope of services requested.

8.15 (E) BACK UP AND DISASTER PLAN

8.15.1 Ability to apply legal retention periods and disposition by agency per purchasing entity policy and/or legal requirements.

Unisys Response:

Customers can customize retention periods in accordance with purchasing entity policy, legal requirements, or both across all providers in this RFP.

8.15.2 Describe any known inherent disaster recovery risks and provide potential mitigation strategies.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: We feel very confident about our physically and logically redundant, geographically distributed data center architecture that is constantly replicated. Our provided SLA's speak to our performance and capability in providing the needed protection against such risks.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Amazon's Disaster Recovery Approach is explained in detail in section 8.8 (Service or Data Recovery)

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft's Disaster Recovery Approach is explained in detail in section 8.8 (Service or Data Recovery)

Google Public and Community Cloud for SaaS:

Google's Disaster Recovery Approach is explained in detail in section 8.8 (Service or Data Recovery)

8.15.3 Describe the infrastructure that supports multiple data centers within the United States, each of which supports redundancy, failover capability, and the ability to run large scale applications independently in case one data center is lost.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has set up primary and secondary data centers. If required by a contract, applications and devices are configured in high availability mode, monitored, and tested regularly. Unisys establishes data mirroring controls and performs backups

regularly. We define recovery parameters such as RTOs and RPOs. We perform risk assessments and business impact analyses regularly on our data centers. In case of a data center outage, the disaster recovery trigger would automatically kick in and will enable failover to other data centers.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Amazon's infrastructure has a high level of availability and provides customers with the features to deploy a resilient IT architecture. AWS designed its systems to tolerate system or hardware failures with minimal impact on customers. The Amazon Infrastructure Group directs Business Continuity Management for data centers at AWS.

Availability

Data centers are built in clusters in various global regions. Data centers are on line and serving customers; no data center is "cold." During a failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that during a data center failure, there is sufficient capacity to enable traffic to be load balanced to the remaining sites.

AWS provides customers with the flexibility to move instances and store data in multiple geographic regions as well as across multiple availability zones in each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated in a typical metropolitan region and located on lower risk flood plains (specific flood zone categorization varies by region). In addition to discreet uninterruptible power supply (UPS) and onsite backup generation facilities, they are each fed by different grids from independent utilities to further reduce single points of failure. Availability zones are redundantly connected to multiple Tier 1 transit providers.

Customers should architect their AWS use to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in most failure modes, including natural disasters or system failures.

Fault-tolerant Design

Amazon's infrastructure has a high level of availability and provides customers with the capability to deploy a resilient IT architecture. AWS designed its systems to tolerate system or hardware failures with minimal impact on customers.

Clients should architect their AWS use to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures. However, clients should be aware of location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between regions unless proactively done by the customer, thereby enabling customers with these types of data placement and privacy requirements to establish compliant environments. Note that communication between regions is across public Internet infrastructure; therefore, appropriate encryption methods should be used to protect sensitive data.

As of this proposal, there are 12 regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), EU (Ireland), EU (Frankfurt), Asia-Pacific (Singapore), Asia-Pacific (Tokyo), Asia-Pacific (Sydney), Asia-Pacific (Seoul), South America (São Paulo), and China (Beijing).

AWS GovCloud (US) is an isolated AWS region designed to allow Federal agencies and customers to move workloads to the cloud by helping them meet certain regulatory and compliance requirements. The AWS GovCloud (US) framework allows Federal agencies and their contractors to comply with U.S. International Traffic in Arms Regulations (ITAR) regulations and Federal Risk and Authorization Management Program (FedRAMP) requirements. AWS GovCloud (US) received an Agency Authorization to Operate (ATO) from the U.S. Department of Health and Human Services (HHS) to use a FedRAMP-accredited Third Party Assessment Organization (3PAO) for several AWS services.

The AWS GovCloud (US) Region provides the same fault-tolerant design as other regions, with two Availability Zones. Additionally, the AWS GovCloud (US) region is a mandatory AWS Virtual Private Cloud (VPC) service by default to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft built a fully physically and logically redundant computing platform. Microsoft also provides the flexibility to easily back up critical applications to georedundant cloud data centers and client on-premises or partner data centers. This adds yet another layer of disaster recovery capability. Microsoft delivers more than 200 cloud services that are hosted in more than 100 globally distributed data centers, edge computing nodes, and service operations centers. This infrastructure is supported by one of the world's largest built terabyte global networks, with an extensive dark fiber footprint, that connects them all.


Google Public and Community Cloud for SaaS: The table below lists the Google datacenters the in US.

Google US Data Centers

Product	North America
Compute Engine	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Preemptible VMs	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Container Engine	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Container Registry	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Cloud Dataflow	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Cloud Bigtable	Council Bluffs, IA - us-central1
Cloud Dataproc	Council Bluffs, IA - us-central1

The following diagram shows how these Google Cloud Platform components work together to build a scalable, resilient web application.

Exhibit 26 Google Cloud Platform



Each component in the example application architecture plays a role in ensuring the application is both scalable and resilient. This section briefly describes each of these services.

The HTTP load balancer exposes a single public IP address that customers use to access the application. This IP address can be associated with a DNS A record (e.g., example.com) or CNAME (e.g., www.example.com). Incoming requests are distributed across the instance groups in each zone according to each group's capacity. Within the zone, requests are spread evenly over the instances within the group. Although the HTTP load balancer can balance traffic across multiple regions, we're using it in a single region with multiple zones, as described in the next section.

A zone is an isolated location within a region. Zones have high-bandwidth, low-latency network connections to other zones in the same region. Google recommends deploying applications across multiple zones in a region.

An instance is a virtual machine hosted on Google's infrastructure. You can install and configure these instances just like physical servers. In this document, you can use startup scripts and Chef to configure instances with the application server and code for the web application.

An instance group is a collection of homogeneous instances that can be targeted by an HTTP load balancer. Instances are added to and removed from a group by an instance group manager. An instance group and corresponding manager are required for each zone you want to run in.

The Compute Engine Autoscaler adds or removes Google Compute Engine instances to an instance group by interfacing with the group's manager in response to traffic, CPU utilization, or other signals. In the example solution, the Autoscaler responds to the Request Per Second (RPS) metric of the HTTP load balancer. An Autoscaler is required for each instance group that you want scaled automatically.

Google Cloud SQL is a fully managed MySQL database. Replication, encryption, patches, and backups are managed by Google. A Cloud SQL instance is deployed to a single zone, and data is replicated to other zones automatically. The Redmine application used in this example is compatible with MySQL and works seamlessly with Cloud SQL.

Cloud Storage allows objects (usually files) to be stored and retrieved with a simple and scalable interface. In this solution, a Cloud Storage Bucket is used to distribute private SSL keys to the scalable Google Compute Engine instances, and is also used to store all files uploaded to the Redmine application, meaning no stateful information is stored on any instance's disks.

A dynamic decision is made throughout a user's log in session on which servers and which data centers will be used to store primary and backup copies of all data. If one server begins to fail, data is moved to another healthy server. If one data center is having broad localized issues, all data will be served up and stored at the other healthy data centers. No action is required by the customer to invoke this service and it is transparent to your end users.

8.16 (E) SOLUTION ADMINISTRATION

8.16.1 Ability of the Purchasing Entity to fully manage identity and user accounts.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System enables several authentication methods, including username/password, username/password + 2nd factor, identity federation (SAML 2.0), OpenID, and a proprietary Single Sign-On (SSO) API. User authentication credentials can be stored locally in the Unisys Cloud Provisioning System and authenticated there, or they can be leveraged from a customer-hosted LDAP directory (including AD). Note that the LDAP/AD model works only for the on-premises deployment of the Unisys Cloud Provisioning System and is not supported through our SaaS offering.

When username/password is the option for the local Unisys Cloud Provisioning System user store, passwords are stored securely using a versioned one-way hashing algorithm. This algorithm can be replaced if necessary.

The Unisys Cloud Provisioning System also enables users to have one or more API keys for talking to the Unisys Cloud Provisioning System API. These API keys enforce a user's access rights when the user is interacting with the Unisys Cloud Provisioning System API. A user can revoke an API key without affecting the others. These API keys are stored securely using AES-256 encryption based on an administrator-managed encryption key.

Each user can belong to zero or more groups. Users and groups can be maintained in the Unisys Cloud Provisioning System, or synchronized with an external LDAP or AD directory (regardless of whether authentication goes through LDAP or AD). In an account (such as an R&D Amazon Web Services account instead of a production account), each group is assigned a role. A group can therefore have a different role for the R&D account as opposed to the production account. In fact, no role can be assigned that results in no access rights for that role. Each user set up in the Unisys Cloud Provisioning System can share a common access point with Amazon Web Services—freeing system administrators from the requirement to create individual Amazon Web Services users.

Access rights are determined by the role. System administrators can create their own roles and map them to groups to suit business needs. System administrators can also define permissions for each role at a very granular level. An access right is defined in a RESOURCE, ACTION, QUALIFIER triad. The default for a new role is no access rights at all. System administrators can create an administrative role by setting the triad to ANY/ANY/ANY, or they can limit the ability to terminate virtual machines for a specific role by setting the right VM/Terminate/MINE. System administrators can define of permissions for roles.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Each participating agency can have its own AWS account. AWS provides several ways for customers to identify themselves and securely access their AWS account.

AWS Identity and Access Management (IAM) enables system administrators to securely control access to AWS services and resources for users. IAM enables system administrators to create and manage users in

AWS, and it enables them to grant access to AWS resources for users managed outside AWS in a corporate directory.

IAM also enables identity federation between corporate directory and AWS services. This enables system administrators to use existing corporate identities to grant secure and direct access to AWS resources, such as S3 buckets, without creating a new AWS identity for those users.

Exhibit 15 shows the AWS IAM tab on the AWS Management Console.




Exhibit 15. AWS IAM Tab.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Managing identity is just as important in the public cloud as it is on premises. To help with managing identity, Azure supports several different cloud identity technologies, which include the following.

Administrators can run Windows Server Active Directory (commonly called just AD) in the cloud using virtual machines created with Azure Virtual machines. This approach makes sense when administrators using Azure to extend an on-premises data center to the cloud.

Administrators can use Azure Active Directory to give users a single sign-on to Software as a Service (SaaS) applications. Microsoft Office 365 uses this technology, for example, and applications running on Azure or other cloud platforms can also use it.

Applications running in the cloud or on premises can use Azure Active Directory Access Control to let users log in using identities from Facebook, Google, Microsoft, and other identity providers.

There are several options for connecting the domain controllers in the cloud with those running on premises:

- Make them part of a single Active Directory domain.
- Create separate AD domains on-premises and in the cloud that are part of the same forest.
- Create separate AD forests in the cloud and on-premises, then connect the forests using cross-forest trusts or Windows Server Active Directory Federation Services (AD FS), which can also run in virtual machines on Azure.

As SaaS applications become more common, they raise an obvious question: What kind of directory service should these cloud-based applications use? Microsoft's answer to that question is Azure Active Directory.

There are two main options for using this directory service in the cloud:


- Individuals and organizations that use only SaaS applications can rely on Azure Active Directory as their sole directory service.
- Organizations that run Windows Server Active Directory can connect their on-premises directory to Azure Active Directory, then use it to give their users single sign-on (SSO) to SaaS applications.

Administrators can federate on-premises AD or other directory stores with Azure AD. Microsoft recommends that organizations synchronize their on-premises directory information with SSO to enable users removed from on-premises directories to be also removed from the Azure AD to maintain sufficient access controls. Deployment of a highly available Security Token Service (STS) (such as Active Directory Federation Services 2.0 or AD FS) on premises is required because this STS will become part of the authentication flow for every user login. Once federation is configured, Azure AD users whose identities are based on the federated domain can use their existing corporate login to authenticate to Azure AD services. Federation enables secure, token-based authentication and SSO across Azure applications.

Microsoft provides Multi-Factor Authentication for Azure administrators with a phone as the second factor; it also supports integration with third-party authentication solutions via on-premises STS integration.

Google Public and Community Cloud for SaaS: Each participating agency can have its own Google account. Google provides several ways for customers to identify themselves and securely access their AWS account. Google Identity and Access Management (IAM) enables system administrators to securely control access to Google services and resources for users. IAM enables system administrators to create and manage users in the Google Cloud. IAM also enables them to grant access to Google resources.

The Graphic below Google IAM Overview illustrates Google IAM model.



8.16.2 Ability to provide anti-virus protection, for data stores.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys Data centers and assets are protected by Symantec End Point Protection. Unisys infrastructure has policies defined that allows real-time background scanning in which a memory resident scanner runs continuously in background mode and automatically scans files and media upon access. All Unisys assets including their Data centers undergoes periodic full system scans that would detect for malicious and

unwanted programs, files and media. It is also the responsibility of the client to make sure to follow best practices while using Unisys infrastructure to refrain from accessing potential malicious sites and programs.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Moving IT infrastructure to AWS creates a shared responsibility model between the customer and AWS. This shared model can relieve customer operational burden because AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. Customers assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, and the configuration of the AWS-provided security group firewall. Customers should carefully consider the services they choose because their responsibilities vary, depending on the services used, the integration of those services into their IT environment, and applicable laws and regulations. To enhance security or meet more stringent compliance requirements (or do both), customers can use technology such as host-based firewalls, host-based intrusion detection and prevention, antivirus software, encryption, and key management. The nature of this shared responsibility also provides the flexibility and customer control that allows deployment of solutions that meet industry-specific certification requirements.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft Antimalware offers customers the ability to install an antimalware agent for PaaS roles and virtual machines. Based on System Center Endpoint Protection, this new feature brings proven on-premises security technology to the cloud.

We offer deep integration for Trend's Deep Security and SecureCloud products on the Azure platform. DeepSecurity is an antivirus solution; SecureCloud is an encryption solution. DeepSecurity will be deployed inside virtual machines using an extension model that Azure already announced. Using the portal UI and PowerShell, customers can choose to use DeepSecurity inside new virtual machines that are being spun up, or existing virtual machines that are already deployed.

Symantec End Point Protection (SEP) is also supported on Azure. Through portal integration, customers can specify that they intend to use SEP in a virtual machine. SEP can be installed on a brand new virtual machine at the Azure Portal or on an existing virtual machine using PowerShell.

Google Public and Community Cloud for SaaS:

Google's malware strategy begins with infection prevention by using manual and automated scanners to scour Google's search index for websites that may be vehicles for malware or phishing. Approximately one billion people use Google's Safe Browsing on a regular basis. Google's Safe Browsing technology examines billions of URLs per day looking for unsafe websites. Every day, we discover thousands of new unsafe sites, many of which are legitimate websites that have been compromised. When we detect unsafe sites, we show warnings on Google Search and in web browsers. In addition to our Safe Browsing solution, Google operates VirusTotal, a free online service that analyzes files and URLs enabling the identification of viruses, worms, trojans and other kinds of malicious content detected by antivirus engines and website scanners. VirusTotal's mission is to help in improving the antivirus and security industry and make the Internet a safer place through the development of free tools and services.

8.16.3 Ability to migrate all Purchasing Entity data, metadata, and usage data to a successor Cloud Hosting solution provider.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys provides client the ability to download or migrate their data and workloads as and when required by the client. Unisys provides Cloud Application Migration Services that starts with a complete and analytical view of your application portfolio and cloud/data center capabilities. Using a set of criteria refined during customer engagements, Unisys can help you narrow the application pool of various cloud

portable candidates and then provide the detailed assessment of costs and other factors that go into a cloud migration decision. Once the decision has been made, Unisys can help you with the planning and implementation of the application migration process including on-boarding of applications to AWS and other clouds.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Customers are free to export data. For large amounts of data, they can use the AWS Import/Export service to copy data to customer-supplied storage and deliver it to a customer-specified location. AWS Import/Export Disk accelerates moving large amounts of data into and out of the AWS cloud by using portable storage devices for transport. Disk transfers your data directly onto and off of your storage devices using the Amazon high-speed internal network and bypassing the Internet. For gigabyte-scale data sets that would transfer at less than 10 Mbps, Disk is often faster than Internet transfer and more cost-effective than upgrading your connectivity.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

A customer's use of an Import/Export Service is conditioned on its compliance with the instructions provided by Microsoft for the preparation, treatment, and shipment of physical media containing its data (storage media). The customer is solely responsible for verifying the storage media and data are provided in compliance with laws and regulations. Microsoft has no duty toward the storage media and no liability for lost, damaged, or destroyed storage media. Storage media shipped to Microsoft must be shipped DAP Microsoft DCS Data Center (INCOTERMS 2010). Storage media shipped to the customer is shipped DAP Customer Dock (INCOTERMS 2010).

Google Public and Community Cloud for SaaS: Customers are free to export data. For large amounts of data, they can use our documented API's to Import/Export /copy data to customer-supplied storage.

8.16.4 Ability to administer the solution in a distributed manner to different participating entities

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS: The Unisys Cloud Provisioning System can also be managed in a distributed way. Access controls enable managers to have a specific level of control in the system. For example, a person in the QA department in one city can manage the QA systems and budget controls, and a person in the Development department in another city can manage a different set of infrastructure and the corresponding budget.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Each Participating Entity can be provisioned with one or more separate accounts. The AWS environment is a virtualized, multitenant environment. AWS implemented security management processes, PCI controls, and other security controls designed to isolate each customer account from other customer accounts. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture was validated by an independent PCI Qualified Security Assessor (QSA) and found to comply with the requirements of PCI DSS version 2.0 published in October 2010.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Azure Role-Based Access Control (RBAC) enables fine-grained access management for Azure. Using RBAC, administrators can segregate duties on a DevOps team and grant users only the amount of access that they need to perform their jobs.

Each Azure subscription is associated with one Azure Active Directory. Only users, groups, and applications from that directory can be granted access to manage resources in the Azure subscription, using the Azure portal, Azure Command-Line tools and Azure Management APIs.

To grant access, administrators assign the appropriate RBAC role to users, groups, and applications at the right scope. To grant access to the entire subscription, administrators assign a role at the subscription scope. To grant access to a specific resource group in a subscription, administrators assign a role at the

resource group scope. They also can assign roles for specific resources, such as websites, virtual machines and subnets, to grant access only to a resource.

Google Public and Community Cloud for SaaS: Each Participating Entity can be provisioned with one or more separate accounts. The Google environment is a virtualized, multitenant environment. Google provides RBAC to the underlying infrastructure. Google systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software.

8.16.5 Ability to apply a participating entity's defined administration policies in managing a solution.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System supports major cloud providers with a unified management solution that includes access controls, financial tracking, and automation. The Unisys Cloud Provisioning System also has a REST-based API to allow for integration with internal and third-party systems. Each Participating Entity can set up its own policies for who has access, how budgets are handled, and more.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Each Participating Entity can be provisioned with one or more separate accounts. The AWS environment is a virtualized, multitenant environment. AWS implemented security management processes, PCI controls, and other security controls designed to isolate each customer account from other customer accounts. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture was validated by an independent PCI Qualified Security Assessor (QSA) and found to comply with the requirements of PCI DSS version 2.0 published in October 2010.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Each subscription in Azure belongs to only one directory, each resource group belongs to only one subscription, and each resource belongs to only one resource group. Access that administrators grant at parent scopes is inherited at child scopes. If administrators grant a reader role to an Azure AD group at the subscription scope, that group's members will be able to view every resource group and every resource in the subscription. If administrators grant the contributor role to an application at the resource group scope, the application will be able to manage resources in that resource group, but not other resource groups in the subscription.

The finer grained authorization model (Azure RBAC) is supported only for management operations of the Azure resources at the Azure portal and in Azure Resource Manager APIs. Not all data-level operations for Azure resources can be authorized by Azure RBAC. For example, the creation, reading, updating, and deletion of storage accounts can be controlled by RBAC, but the creation, reading, updating, and deletion of blobs or tables in the storage account cannot yet be controlled by RBAC. Similarly, the creation, reading, updating, and deletion of a SQL database can be controlled by RBAC, but the creation, reading, updating, and deletion of SQL tables in the database cannot yet be controlled by RBAC.

In Azure API Management, policies are a powerful capability of the system that allow a publisher to change the API's behavior through configuration. Policies are a collection of statements that are executed sequentially on the request or response of an API. Popular statements include format conversion from XML to JSON and call-rate limiting to restrict the amount of incoming calls from a developer. Many more policies are available out of the box.

Google Public and Community Cloud for SaaS:

Each Participating Entity can be provisioned with one or more separate accounts. The Google environment is a virtualized, multitenant environment. Google provides RBAC to the underlying infrastructure. Google systems are designed to prevent customers from accessing physical hosts or

instances not assigned to them by filtering through the virtualization software.

8.17 (E) HOSTING AND PROVISIONING

8.17.1 Documented cloud hosting provisioning processes, and your defined/standard cloud provisioning stack.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System includes not only basic provisioning (e.g., starting a server, creating a machine image, creating a new volume, attaching the volume to a server, creating a snapshot of a volume, creating a volume from a snapshot, reserving an IP address, creating a standard network or port, or creating or modifying a firewall), but also robust automation features such as creating and deploying multitier applications and keeping the application 100 percent available through autoscaling, autorecovery, and automatic backups. This enhancement is more significant than that in the base Amazon provisioning capability. Additionally, using the Unisys Cloud Provisioning System enables this capability across our supported cloud hosts and private cloud solutions.

AWS Public and Private Cloud for IaaS, PaaS, and SaaS:

There are several methods by which customers can provision a server or application services on AWS. Usually, customers use the web-based EC2 portal, which guides them through the provisioning process. The EC2 portal offers a wide array of preconfigured, fully automated deployable workloads to customize every step of a workload instantiation. Beyond EC2, customers can use external tools to deploy servers or application services to the EC2 portal. These external tools can be vendor products, Open Source solutions, and homegrown scripts and technologies. Exhibit # illustrates the various methods and tools of provisioning available from AWS.




Exhibit 26. AWS Provisioning Process, Methods, and Tools.**Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:**


Microsoft Azure Provisioning is an automated process that allows interaction from the Azure Portal or completely remotely using Azure APIs. For example, a virtual machine can be allowed to interact from the Azure portal in just a few minutes. The normal steps taken are planning for the deployment, selecting the virtual machine image, creating the virtual machine, and then managing the virtual machine. The virtual machine image can come from the images available in the gallery, including Windows and Linux. The processes support customized images from a customer's on-premises standard image template library. These processes can be repeated one at a time or automated with PowerShell to create hundreds of virtual machines at once.

This example is for automating a virtual machine, but the properties in Microsoft Azure can be automated using the same methods, at the portal or in PowerShell. Other examples for automating include cloud storage, SQL Servers, and virtual networks.

Microsoft offers the Azure Portal as the access gateway to create, modify, and decommission workloads on the Microsoft cloud platform. The Azure portal can be accessed over the Internet. **Appendix 19 – Getting Started with VMs on Windows Azure** outlines the provisioning process and the provisioning stack.


Google Public and Community Cloud for SaaS: The Cloud Platform provides options for computing and hosting. Customers can choose to work with a managed application platform, leverage container technologies to gain lots of flexibility, or build their own cloud-based infrastructure to have the most control and flexibility.

The exhibit below Google Hosting and Provisioning Services illustrates the spectrum.



The exhibit below illustrates how the Google Compute Engine works.

Google Compute Engine in pictures



8.17.2 Provide tool sets at minimum for:

1. **Deploying new servers (determining configuration for both stand alone or part of an existing server farm, etc.)**
2. **Creating and storing server images for future multiple deployments**
3. **Securing additional storage space**
4. **Monitoring tools for use by each jurisdiction's authorized personnel – and this should ideally cover components of a public (respondent hosted) or hybrid cloud (including Participating entity resources)**

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System enables the deployment of cloud servers in new and existing cloud server environments in the context of advanced governance and security to protect users from overspending, which violates ownership and governance rules.

The Unisys Cloud Provisioning System gives system administrators complete control over the creation, configuration, and deployment of not only server images, but also complete multitier application topologies that can be delivered across AWS regions or across cloud providers complete with firewall rules, load balancers, databases, and more.

The Unisys Cloud Provisioning System supports the functionality provided by other vendors such as Amazon and Microsoft. It also supports storage provisioning across many public and private clouds.

The Unisys Cloud Provisioning System supports the direct tracking of CPU usage, memory, network statistics, and disk I/O from the guest operating through our agent. The Unisys Cloud Provisioning System also tracks metrics provided by the cloud provider (e.g., AWS CloudWatch). These can be used for Unisys Cloud Provisioning System autoscaling. The Unisys Cloud Provisioning System graphs some of these metrics, but the metrics can be accessed at the API. The scripts for tracking loads in the guest using the Unisys agent can be extended to use a customer's own scripts or third-party monitoring tools.

Virtual machine load for autoscaling is based on these metrics or on custom metrics. System administrators can write their own scripts that execute on the guest and "vote" on autoscaling. The Unisys Cloud Provisioning System regularly tabulates votes and determines a course of action for scaling.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

The AWS Management Console is a single destination for managing AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. System administrators use the AWS Management Console to perform tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console also enables customers to manage the aspects of their AWS account, including accessing monthly spending by service, managing security credentials, or setting up new AWS Identity and Access Management (AWS IAM) users. The AWS Management Console supports AWS regions and allows customers to provision resources across multiple regions.

Command Line Interface

The AWS Command Line Interface (CLI) is a unified tool used to manage AWS cloud services. With just one tool to download and configure, customers can control multiple AWS resources from the command line and automate them through scripts. The AWS CLI introduces a new set of simple file commands for efficient transfer of files to and from Amazon Simple Storage Service (Amazon S3).

Use of Existing Management Tools

Many of the tools that organizations use to manage on-premises environments also can be integrated with AWS. Integrating an AWS environment can provide a simpler and quicker path for cloud adoption because a customer's operations team does not have to learn new tools or develop completely new processes. For example, AWS provides the following tools.

The AWS Management Portal for vCenter enables customers to manage their AWS resources using VMware vCenter. The portal installs as a vCenter plugin in the existing vCenter environment. Once installed, it enables customers to migrate VMware virtual machines to Amazon EC2 and manage AWS resources from within vCenter. The AWS resources that customers create at the portal can be located in their AWS account, even though those resources were created using vCenter. For experienced VMware administrators, AWS Management Portal for vCenter provides a familiar look and feel that can make it easy to start using AWS. AWS Management Portal for vCenter is available at no additional charge.

The Amazon EC2 VM Import Connector extends the capabilities of VMware vCenter to provide a familiar graphical user interface that customers can use to import their preexisting virtual machines to Amazon EC2. Using the connector, importing a virtual machine is as simple as selecting a virtual machine from the vSphere infrastructure and specifying the AWS region, Availability Zone, operating system, instance size, security group, and Amazon Virtual Private Cloud (Amazon VPC) details (if desired) to which the virtual machine should be imported. Once the virtual machine is imported, system administrators can launch it as an instance from the AWS Management Console and immediately take advantage of the features of Amazon EC2.

AWS Management Pack for Microsoft System Center enables system administrators to view and monitor their AWS resources directly in the Operations Manager console. This way, they can use a single, familiar console to monitor their resources, whether they are on the premises or in the AWS cloud. System administrators get a consolidated view of AWS resources across regions and Availability Zones. It also has built-in integration with Amazon CloudWatch so that the metrics and alarms defined in Amazon CloudWatch surface as performance counters and alerts at the Operations Manager console.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

The Azure Portal provides the flexibility to create one or thousands of virtual machines (VMs). A single VM can be created, and two or more VMs can be created and moved to an "availability set" so that the uptime SLA is applied.

The Azure Portal and the underlying Azure APIs are used to extend or scale systems easily. Sites (PaaS) options in Azure are scaled by simply configuring the Azure Portal to extend the web farm from 1 to 10 or more in accordance with customized needs. For example, the number of client connections, CPU

utilization, and other metrics could be configured once and the scaled-out systems would deploy automatically, without interaction from administrators.

The Azure Portal supports VMs to be deployed to a previously created VM farm. The VMs are created using automated processes and then configured either through the VM operating system processes or automated by using scripting such as PowerShell.

Additionally, custom virtual machines can be created using the same on-premises processes. They are then uploaded to the Azure Portal to include in the customer's VM library. Once the image is added to the library, the automation processes for single VM deployment or joining to a VM farm are exactly the same.

With a cloud service previously created, the addition of cloud storage can be automated using the Azure Portal or PowerShell. Azure Blob storage is a service that stores file data in the cloud. Blob storage can store text or binary data, such as a document, a media file, or an application installer. Blob storage is sometimes referred to as object storage.

Azure Blob storage can be completely automated using .NET, Node.js, Java, C++, PHP, Ruby, Python, IOS, and Xamarin.

Blob storage includes the following common uses:

- Serving images or documents directly to a browser
- Storing files for distributed access
- Streaming video and audio
- Performing secure backup and disaster recovery
- Storing data for analysis by an on-premises or Azure-hosted service.

Monitoring cloud services are included at the Azure portal and available from APIs at the Azure portal and on the individual VM available in the operating system. Administrators can monitor key performance metrics for their cloud services at the Azure classic portal. They can set the level of monitoring to minimal and verbose for each service role and can customize the monitoring displays. Verbose monitoring data is stored in a storage account, which administrators can access outside the portal.

Monitoring displays at the Azure classic portal are highly configurable. Administrators can choose the metrics they want to monitor on the metrics list on the Monitor page. They can choose which metrics to plot on metrics charts on the Monitor page and the dashboard.

By default, Azure provides minimal for a new cloud service using performance counters gathered from the host operating system for the roles' instances (virtual machines). The minimal metrics are limited to CPU Percentage, Data In, Data Out, Disk Read Throughput, and Disk Write Throughput. By configuring verbose monitoring, administrators can receive additional metrics based on performance data in the virtual machines (role instances). The verbose metrics enable closer analysis of issues that occur during application operations.

By default, performance counter data from role instances is sampled and transferred from the role instance at 3-minute intervals. When administrators enable verbose monitoring, the raw performance counter data is aggregated for each role instance and across role instances for each role at intervals of 5 minutes, 1 hour, and 12 hours. The aggregated data is purged after 10 days.

Google Public and Community Cloud for SaaS: The Google Cloud Platform consists of a set of physical assets, such as computers and hard disk drives, and virtual resources, such as virtual machines (VMs), that are contained in Google's data centers around the globe. Each data center location is in a global region. Regions include Central US, Western Europe, and East Asia. Each region is a collection of zones, which are isolated from each other within the region. Each zone is identified by a name that combines a letter identifier with the name of the region.

The Google Cloud Platform Console provides a web-based, graphical user interface that customers can use to manage their Cloud Platform projects and resources. Using the Cloud Platform Console, customers can create a new project, or choose an existing project, and use the resources that were create in the context of that project.

Google Cloud SDK provides the gcloud command-line tool, which gives customers another method of access. The gcloud tool can be used to manage both development workflow and the Cloud Platform resources.

Cloud Platform provides tools for logging and monitoring so you can keep track keep track of the performance and availability of your resources and applications.

Cloud Logging collects and stores logs from applications and services running on Cloud Platform. Customers can use Cloud Logging with App Engine or Compute Engine. The Logs Viewer in the Cloud Platform Console lets one see their logs. Logs can be exported to Cloud Storage, BigQuery, and Cloud Pub/Sub so cunstomers can process them more easily. The Cloud Logging Agent enables you to integrate third-party logs.

Cloud Monitoring provides dashboards and alerts for applications that run on Cloud Platform. Customers configure Cloud Monitoring using the Cloud Monitoring Console. Review performance metrics for cloud services, Compute Engine instances, and common open source servers such as MongoDB, Apache, Nginx, and Elasticsearch. Customers can use the Cloud Monitoring API to retrieve monitoring data and create custom metrics.

8.18 (E) TRIAL AND TESTING PERIODS (PRE- AND POST- PURCHASE)

8.18.1 Describe your testing and training periods that your offer for your service offerings.

Unisys Response:

We offer generous free trials in some cases for up to a full year. The product set offered during this free period is identical to the paid version of the product. Potential clients, existing clients, or both can use this benefit to take a test drive to learn the environment

8.18.2 Describe how you intend to provide a test and/or proof of concept environment for evaluation that verifies your ability to meet mandatory requirements.

Unisys Response:

In the public and the community spaces for IaaS, PaaS, and SaaS, our partners offer free trials that are generous and come with identical features and functionality of the paid version of the product. Potential clients, existing clients, or both can use this benefit to take a test drive to learn the environment. We also provide a significant amount of audiovisual material and other knowledge transfer media such as white papers and how-to documents. Potential clients, existing clients, or both can use this material for training. For the private and the hybrid spaces for IaaS, PaaS, and SaaS, Unisys customizes most of the solution offerings to the client's specifications. We provide access to user interfaces, onsite proofs of concept, and site visits to our facility, where clients can interact with our staff and test drive environments for feasibility and learning.

8.18.3 Offeror must describe what training and support it provides at no additional cost.

Unisys Response:

Unisys will provide Participating Entities with the following value-add across our cloud models and service models at no additional cost:

- **Provisioning:** Unisys will provision the access login ID for the Participating Entity.
- **Billing Portal:** Unisys will provide a customized, online web-based portal for the Participating Entities to monitor and report on their consumption, utilization, current balances, etc.

- **Billing:** Unisys will bill Participating Entities each month.
- **Consultation:** Unisys will guide Participating Entities on how to publicly available learning tools and content from our partners.

Additionally, the public-facing websites of Unisys and of our partners in this proposal are rich in training content, support material, and tools that our prospects, clients, or both can access at no additional cost.

8.19 (E) INTEGRATION AND CUSTOMIZATION

8.19.1 Describe how the Solutions you provide can be integrated to other complementary applications, and if you offer standard-based interface to enable additional integrations.

8.19.2 Describe the ways to customize and personalize the Solutions you provide to meet the needs of specific Purchasing Entities.

Unisys and its partners platforms on this RFP provide a completely integrated and built on the concept of pluggable, reusable components. The interfaces conform to a standard framework, protocols, and design tools. We provide formal framework documentation; how-to documents; toolsets; and prebuilt, reusable, and customizable components that clients can leverage for additional integration or customization. A significant amount of related content, training material, Software Development Kits (SDK) is freely available on our website. Our sales architects are available to provide free consultation on an as needed basis. Our professional services staff is available for contracted engagements. Typically, our customers will leverage the information available on our portals to integrate and customize the environment internally or customer can contract with us for this work.

8.20 (E) MARKETING PLAN

Describe your how you intend to market your Solutions to NASPO ValuePoint and Participating Entities.

Unisys Response:

Unisys designated sales and business development and our NASPO contract manager will perform this function. Additionally, we plan to take an active role in the NASPO community. We see ourselves bringing value to the NASPO community by sponsoring and attending NASPO events. We will raise awareness of NASPO in our own client base and work very closely with the State of Utah and the Participating Entities of this contract vehicle. Unisys has been a participant of NASCIO; we will continue to demonstrate our support through participation and sponsorships.

8.21 (E) RELATED VALUE-ADDED SERVICES TO CLOUD SOLUTIONS

Describe the valued-added services that you can provide as part of an awarded contract, e.g. consulting services pre- and post- implementation. Offerors may detail professional services in the RFP limited to assisting offering activities with initial setup, training and access to the services.

Value-add services at No Cost to the Participating Entities

If Unisys is awarded the contract, besides the onboarding consultation services, contract management, contract support organization, billing, and fee disbursements that are required by this contract, we will offer the following additional services at no cost to the Participating Entities:

- **Provisioning:** Unisys will provision the access login ID for the Participating Entity.
- **Billing Portal:** Unisys will provide a customized, online web-based portal for the Participating Entities to monitor and report on their consumption, utilization, current balances, etc.
- **Billing:** Unisys will bill Participating Entities each month.
- **Consultation:** Unisys will guide Participating Entities on how to locate publicly available learning tools and content from our partners.

Fee-based Public Sector Cloud based Solutions and Value-Add Professional Services:

Following is a summary of the public sector solutions and professional services offered by Unisys to the Participating Entities.

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud: Outlined below are the cloud based solution offerings from Unisys.

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
Unisys Justice Law Enforcement and Border Security Solutions for the Cloud		
Secure Image Management Solution	Unisys developed a highly secure solution to meet the forensic integrity standards and policy requirements of several of the largest police departments in the world. The secure storage of images and other multimedia content occurs in a tamperproof information archive – in effect, a digital evidence vault.	All
U-LEAF	U-LEAF provides a POLE-type data model (Person, Object, Location and Event) for the storage and recording of incidents and entities. The POLE model allows entities to be recorded in the system once. The recorded entities, however, can be linked to other entities and events as many times as necessary, to build the picture of an incident, or a network of associations.	All
Unisys Social Services Solutions for the Cloud		
Unisys 311	The Unisys “311 solution” is a Multi- Channel Citizen Service Delivery and Engagement solution that can be used by a Local Government or any Governmental Agency to deliver and manage non-emergency services, respond to inquiries and engage with their constituents. This offers multi-channel interaction including social media and mobile applications, supported by a knowledgebase and GIS systems and also interfaces with existing systems to manage work orders and service requests. The solution can be configured to meet specific needs and extend its functionality and flexibility through other apps, and integration to mobile apps, social media and other systems.	All
State/Local Government Enterprise Regulatory System (AMANDA)	The Unisys AMANDA Platform (from CSDC) is designed to provide a collection of back office functions such as Licensing, Permitting, Inspections, Land Use, Planning and a number of other functions in various form factors (desktop, mobile and tablet) and can also be deployed either on premise, hosted, or delivered via the Cloud. The solution includes the design, delivery, deployment and ongoing support and maintenance of the solution. The AMANDA platform can be configured specific to client needs in order to extend its functionality and interface with other systems as required.	All
Unisys Enterprise Content Management for the Cloud		
Infomage	Unisys Infomage is the Enterprise Content Management (ECM), Business Process Management (BPM) and Record Management (RM) solution for organizations looking to significantly improve business processes that depend heavily high-volume paper documents, documents generated from Internet transactions, Office documents, and other electronic documents that need to be accessed for automated and manual processes. Infomage can easily capture, manage, store, and access the content required for cases, inquiries, and process-centric work, regardless of	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	data structure or document origination, with a single intuitive user interface. Unisys InfoImage brings together ECM, imaging, workflow, document management, record management and integration technologies to form an integrated end-to-end solution.	
Unisys Horizontal Solutions for the Cloud		
Stealth	The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. We built a better way to deal with advanced threats for our clients by applying novel approaches to new threats. By substituting traditional hardware topology for software-based cryptography, our Stealth Microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users.	Private and Hybrid Cloud Only. For Public and Community Cloud, see AWS offering above. IaaS, PaaS and SaaS
VantagePoint	VantagePoint can also extend beyond the boundaries of IT with uses cases to support relevant business scenarios in, • Security: Strengthens the security posture of organizations through greater visibility into the Stealth enabled infrastructure. • Advanced Data Analytics: Cuts across data silos and sources relevant data to optimize the performance of Advanced Analytics solution. • Cloud and Infrastructure Services: Enables hybrid IT by seamlessly integrating with a variety of data sources and services and presenting them via a digital control panel, irrespective of technology underpinnings. • End User Services: Supports millennial users via personalized and secure access to relevant data and services, across endpoints. • Service Management: Cuts across ITSM platforms to improve service delivery through service automation, orchestration, and aggregation capabilities. • Facilities/Crisis Management: Helps organizations monitor their facilities and keep their Facilities Managers abreast of any security threats and risks. Also acts as a standard communication vehicle for users at times of crisis. (178) With VantagePoint, our clients gain a personalized, secure, and intuitive data and service aggregation platform that cuts across strategic and operational dimensions of business and accelerates digital transformation.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys Cloud Hosting	Unisys Hosted Private Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can provide, integrate and scale the cloud infrastructure to meet client's needs. Our management platform enables a single pane of glass to effectively manage client's workloads. Unisys Hybrid Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can source, integrate and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our management	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	platform enables a single pane of glass to effectively manage a variety of clouds within an organization, whether they be public or private.	
ServiceNow	ServiceNow offers enterprise service management software for human resources, law, facilities management, finance, marketing, and field operations in the cloud. ServiceNow has its STAR Self-Assessment available on the CSA's website. ServiceNow specializes in ITSM applications and provides forms-based workflow application development. ServiceNow has open integration options to variety platforms such as: Salesforce, SharePoint, and BMC Remedy Action Request System.	All
Unisys ClearPath Forward!	Unisys ClearPath Forward is an Intel based fabric computing platform designed to run mission critical applications that require predictable performance and low transport latency. The ClearPath Forward platform provides hardware partitioning technology similar to the Mainframes, but is designed to run Windows and Linux operating environments on commodity server hardware. The fabric interconnect included with the ClearPath Forward platform provides transport speeds at 56Gbps or higher. The ClearPath Forward platform fabric can be used in a hybrid mode with VMware.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys White Label Offerings	Unisys has strategic, long lasting relationships in the industry and is pleased to offer our partner products surrounding the cloud market space. This is an evolving list. Our current partner list includes: SalesForce, Oracle, Verizon, SHI, Decision Lens, Aptio, NetApp, VMware, EMC, NetApp, Birst, Okta, Box, WatchDox. We have created webpage on our public site that is updated with our list of current partners. The page is available here: http://www.unisys.com/ms/wsca-cloud-hosting/	All

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud.

Exhibit Y. Unisys Services Offerings: Outlined below are our professional services surrounding the cloud.

Unisys Professional Services	Description
Cloud Advisory Services	Unisys Cloud Advisory Services provide strategic and financial guidance on aligning IT with business objectives. This starts with a roadmap that outlines the vision of a hybrid IT based on a combination of existing data center, internal cloud, and external cloud resources to provide agility, flexibility, and control.
Data Center Planning, Design, and Implementation Services	Unisys Data Center Planning, Design, and Implementation Services offer a complete range of services that delivers cohesive, end-to-end optimization of data centers. With a wide range of services for discovery, analysis, optimization, virtualization, consolidation, and migration of data centers that can complement client efforts and fill gaps in skills and capacities and a combination of world-class people, processes, and technology with program and project management expertise, Unisys transforms clients' existing data centers to a business engine that provides agility at a lower cost.
CloudBuild Services	Unisys CloudBuild Services enable organizations to successfully build a cloud that is integrated with the overall business process, transforming their existing infrastructure to an agile IT-as-a-service model. The Unisys "8-Tracks" model, a 360-degree cloud view approach, covers eight critical data center domains. Together with ConOps, which includes industry best practices, Unisys enables a client's cloud infrastructure to meet the security, regulatory, and compliance requirements that enable us to deliver the most secure and reliable cloud in the marketplace.
Hybrid Cloud Strategy	Unisys Hybrid Cloud strategy helps organizations overcome key challenges when

Unisys Professional Services	Description
	planning to implement a cloud environment. Unisys provides governments and businesses with a comprehensive cloud architecture that gives clients cloud on their terms. We can source, integrate, and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our homegrown VantagePoint management platform enables a single pane of glass to effectively manage a variety of clouds in an organization, whether they be public or private.
Unisys Platform Services	
Platform Assessment Services	Unisys Platform Assessment Services enable governments and businesses to foster their platform adoption initiatives. Unisys analyzes organizations' business goals and objectives along with their technology landscape of their enterprise application portfolio and recommend a best-fit Platform Suitability Analysis to help them make the right decision in their move to PaaS.
Architecture Design Services	Unisys Architecture Design Services assist governments and businesses with designing an effective architecture suited for cloud. We provide several concepts and best practices that are essential to build highly scalable applications in the cloud – be it on-premises, public, or hybrid.
PaaS Enablement Services	Unisys PaaS Enablement Services analyze an organization's current set of applications and provide a strategy to use the best practices and engagement models of various platform providers such as Microsoft Azure, IBM Bluemix, SAP HCP, and SFDC AppCloud.
Application Migration Services	Unisys Application Migration Services assist organizations in their cloud migration activities by using tools and frameworks adhering to Unisys best practices that promise minimal downtime without affecting the day-to-day business processes.
Application Portability Services	Unisys Application Portability Services help organizations to develop cross-platform applications that can be scaled across multiple cloud platforms.
Application Development Services	Unisys Application Development Services help organizations to build rich, interactive applications focused on business logic and workflows using visual tools as well as cloud-based tools, architectures, and services that make their applications cloud ready.
IoT Development Services	Unisys IoT Development Services enable organizations to use IoT Services across various cloud platforms that enable apps to communicate and consume data collected by the connected devices, sensors, and gateways.
API Management Services	Unisys API Management Services provide a solution that addresses the aspects of the application programming interface (API) life cycle for on-premises and cloud environments and offer capabilities to create, run, manage, secure, and monetize APIs and microservices that deliver an integrated user experience and enable rapid deployment and simplified administration of APIs.
Horizontal Technology Integration Services	
Horizontal Technology Integration Services	Unisys is a world-recognized leader in integrating technology infrastructures across platforms and vendors. We particularly specialize in integrating solutions surrounding service desks, cloud provisioning, data center management, and the ITIL framework.
Unisys Service Management Services	
Maturity and Platform Technology Assessments	Unisys Service Management Consulting Services leverage subject matter experts with more than 15 years of experience in working and delivering services for ITIL and service management. We leverage that knowledge to provide best practices, lessons learned, and roadmap development for Service Management disciplines and benchmark them with industry and operational best practices.
Service Management Platform Implementation	Unisys brings experience in moving existing ITSM tool information, processes, and requirements to new cloud enterprise Service Management platforms. By leveraging industry best practices and structured methodologies, Unisys can establish transition and transformation plans that minimize risk to clients' ongoing operations and deliver speed to value with the new cloud-based Service Management platform.
Enterprise Cloud Service Management Platform Support	Unisys can provide ongoing day-to-day support in the ongoing management of a client's Service Management platform. This service allows the organization to focus on the business aspects instead of the routine activities necessary to support and maintain a platform. Unisys provides cost-effective solutions that leverage staff in many locations, driving the cost of servicing the environment to effective and efficient levels.

Exhibit Y. Unisys Services Offerings:

8.22 (E) SUPPORTING INFRASTRUCTURE

8.22.1 *Describe what infrastructure is required by the Purchasing Entity to support your Solutions or deployment models.*

8.22.2 *If required, who will be responsible for installation of new infrastructure and who will incur those costs?*

Unisys Response:

Our proposed cloud models and service delivery options across all providers and solution offerings require only a computer connected to the Internet. The Participating Entity will install this infrastructure and incur those costs.

8.23 (E) ALIGNMENT OF CLOUD COMPUTING REFERENCE ARCHITECTURE

Clarify how your architecture compares to the NIST Cloud Computing Reference Architecture, in particular, to describe how they align with the three domains e.g. Infrastructure as a Service (IaaS), Software as a Service (SaaS), and Platform as a Service (PaaS).

Unisys Response:

For information on the public and the community IaaS, PaaS, and SaaS domains, refer to the responses to RFP requirements 8.6.1, 8.6.2, and 8.6.6. We have provided our comprehensive list of compliance certifications. In the cases relevant to this proposal, the architecture provided by Amazon, Microsoft, and Google exceed the NIST Cloud Computing Reference Architecture requirements. For the private and hybrid cloud environments, Unisys builds and integrates the environment from the ground up, based on client's requirements, which typically exceed NIST requirements.

7.0 CONFIDENTIAL, PROTECTED OR PROPRIETARY INFORMATION

Section Title: Confidential, Protected or Proprietary Information. All confidential, protected or proprietary Information must be included in this section of proposal response. Do not incorporate protected information throughout the Proposal. Rather, provide a reference in the proposal response directing Lead State to the specific area of this protected Information section.

- *If there is no protected information, write "None" in this section.*

Failure to comply with this Section and Section 3.13 of the RFP releases the Lead State, NASPO ValuePoint, and Participating Entities from any obligation or liability arising from the inadvertent release of Offeror information.

Unisys Response:

None.

PRODUCTS & SERVICES

- Amazon EC2 >
- Product Details >
- Instances >
- Pricing >
- Purchasing Options >
- Developer Resources >
- FAQs >
- Getting Started >
- Amazon EC2 Run Command >

RELATED LINKS

- Amazon EC2 Dedicated Hosts
- Amazon EC2 Spot Instances
- Amazon EC2 Reserved Instances
- Amazon EC2 Dedicated Instances
- Windows Instances
- VM Import/Export
- AWS Management Portal for vCenter
- Management Console

Amazon EC2 Service Level Agreement

Last Updated June 1, 2013

This Amazon EC2 Service Level Agreement (“SLA”) is a policy governing the use of Amazon Elastic Compute Cloud (“Amazon EC2”) and Amazon Elastic Block Store (“Amazon EBS”) under the terms of the Amazon Web Services Customer Agreement (the “AWS Agreement”) between Amazon Web Services, Inc. and its affiliates (“AWS”, “us” or “we”) and users of AWS’ services (“you”). This SLA applies separately to each account using Amazon EC2 or Amazon EBS. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Service Commitment

AWS will use commercially reasonable efforts to make Amazon EC2 and Amazon EBS each available with a Monthly Uptime Percentage (defined below) of at least 99.95%, in each case during any monthly billing cycle (the “Service Commitment”). In the event Amazon EC2 or Amazon EBS does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- “Monthly Uptime Percentage” is calculated by subtracting from 100% the percentage of minutes during the month in which Amazon EC2 or Amazon EBS, as applicable, was in the state of “Region Unavailable.” Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon EC2 SLA Exclusion (defined below).
- “Region Unavailable” and “Region Unavailability” mean that more than one Availability Zone in which you are running an instance, within the same Region, is “Unavailable” to you.
- “Unavailable” and “Unavailability” mean:

 Menu

English ▼

My Account ▼

[Create an AWS Account](#)[Discussion Forum](#)[Amazon EC2 SLA](#)[Run Command](#)

- For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue.

- A “Service Credit” is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Commitments and Service Credits

Service Credits are calculated as a percentage of the total charges paid by you (excluding one-time payments such as upfront payments made for Reserved Instances) for either Amazon EC2 or Amazon EBS (whichever was Unavailable, or both if both were Unavailable) in the Region affected for the monthly billing cycle in which the Region Unavailability occurred in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0%	30%

We will apply any Service Credits only against future Amazon EC2 or Amazon EBS payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the Unavailability occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar (\$1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability, non-performance, or other failure by us to provide Amazon EC2 or Amazon EBS is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you must submit a claim by [opening a case in the AWS Support Center](#). To be eligible, the credit request must be received by us

1. the words “SLA Credit Request” in the subject line;

2. the dates and times of each Unavailability incident that you are claiming;
3. the affected EC2 instance IDs or the affected EBS volume IDs; and
4. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the Monthly Uptime Percentage of such request is confirmed by us and is less than the Service Commitment, then we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon EC2 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon EC2 or Amazon EBS, or any other Amazon EC2 or Amazon EBS performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon EC2 or Amazon EBS; (iii) that result from any actions or inactions of you or any third party, including failure to acknowledge a recovery volume; (iv) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); (v) that result from failures of individual instances or volumes not attributable to Region Unavailability; (vi) that result from any maintenance as provided for pursuant to the AWS Agreement; or (vii) arising from our suspension and termination of your right to use Amazon EC2 or Amazon EBS in accordance with the AWS Agreement (collectively, the “Amazon EC2 SLA Exclusions”). If availability is impacted by factors other than those used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

[Prior Version\(s\)](#)

**Menu****English** ▼**My Account** ▼**Create an AWS Account****Create a Free Account** [AWS on Twitter](#) [AWS on Facebook](#) [AWS on Google+](#) [AWS Blog](#) [What's New? RSS](#)**AWS & Cloud Computing**[What is Cloud Computing?](#)[Products & Services](#)[Customer Success](#)[Economics Center](#)[Architecture Center](#)[Security Center](#)[What's New](#)[Whitepapers](#)[AWS Blog](#)[Events](#)[Sustainable Energy](#)[Press Releases](#)[Analyst Reports](#)[Legal](#)**Solutions**[Websites & Website Hosting](#)[Business Applications](#)[Backup & Recovery](#)[Disaster Recovery](#)[Data Archive](#)[Big Data](#)[High Performance Computing](#)[Mobile Services](#)[Digital Marketing](#)[Game Development](#)[Digital Media](#)[Government & Education](#)**Resources & Training**[Developers](#)[Java on AWS](#)[JavaScript on AWS](#)[Mobile on AWS](#)[PHP on AWS](#)

 Menu

English ▼ My Account ▼

Create an AWS Account

[SDKs & Tools](#)
[AWS Marketplace](#)
[User Groups](#)
[Support Plans](#)
[Service Health Dashboard](#)
[Discussion Forums](#)
[FAQs](#)
[Documentation](#)
[Articles & Tutorials](#)
[Test Drives](#)

Manage Your Account
[Management Console](#)
[Billing & Cost Management](#)
[Personal Information](#)
[Payment Method](#)
[AWS Identity & Access Management](#)
[Security Credentials](#)
[Request Service Limit Increases](#)
[Contact Us](#)

Amazon Web Services is Hiring.

Amazon Web Services (AWS) is a dynamic, growing business unit within Amazon.com. We are currently hiring Software Development Engineers, Product Managers, Account Managers, Solutions Architects, Support Engineers, System Engineers, Designers and more. Visit our [Careers](#) page or our [Developer-specific Careers](#) page to learn more.

Amazon Web Services is an Equal Opportunity Employer.

Language

[Deutsch](#) | [English](#) | [Español](#) | [Français](#) | [Italiano](#) | [Português](#) | [Русский](#) | [日本語](#) | [한국어](#) | [中文 \(简体\)](#) | [中文 \(繁體\)](#)

SLA summary for Azure services

Last updated: September 2015

Azure Active Directory

We guarantee at least 99.9% availability of the Azure Active Directory Basic and Premium services. The services are considered available in the following scenarios:

- Users are able to login to the service, login to the Access Panel, access applications on the Access Panel and reset passwords.
- IT administrators are able to create, read, write and delete entries in the directory or provision or de-provision users to applications in the directory.

No SLA is provided for the Free tier of Azure Active Directory.

Download Active Directory SLA (<http://www.microsoftvolumelicensing.com/DocumentSearch.aspx?Mode=3&DocumentTypeId=37>)

View full details (</en-us/support/legal/sla/active-directory/>)

API Management

- We guarantee that API Management Service instances running in the Standard tier will respond to requests to perform operations at least 99.9% of the time.
- We guarantee that API Management Service instances running in the Premium tier deployed across two or more regions will respond to requests to perform operations at least 99.95% of the time.

No SLA is provided for the Developer tier of the API Management Service.

View full details (</en-us/support/legal/sla/api-management/>)

App Service

We guarantee that Web Apps running in a customer subscription will be available 99.95% of the time. No SLA is provided for Mobile Apps, Logic Apps, or API Apps while such services are still in Preview or for Apps under either the Free or Shared tiers.

View full details (</en-us/support/legal/sla/app-service/>)

Application Gateway

We guarantee that each Application Gateway Cloud Service having two or more medium or larger instances will be available at least 99.9% of the time.

[View full details \(/en-us/support/legal/sla/application-gateway/\)](/en-us/support/legal/sla/application-gateway/)

Automation

We guarantee that at least 99.9% of runbook jobs will start within 30 minutes of their planned start times.

We guarantee at least 99.9% availability of the Azure Automation DSC agent service.

No SLA is provided for the Free tier of Azure Automation.

[View full details \(/en-us/support/legal/sla/automation/\)](/en-us/support/legal/sla/automation/)

Backup

We guarantee at least 99.9% availability of the backup and restore functionality of the Azure Backup service.

[View full details \(/en-us/support/legal/sla/backup/\)](/en-us/support/legal/sla/backup/)

BizTalk Services

We guarantee that at least 99.9% of the time customers will have connectivity between their BizTalk Service Environments in the Basic, Standard and Premium tiers and our Internet gateway. We do not offer an SLA for the BizTalk Services Developer tier.

[View full details \(/en-us/support/legal/sla/biztalk-services/\)](/en-us/support/legal/sla/biztalk-services/)

Cache

We guarantee at least 99.9% of the time that customers will have connectivity between the Cache endpoints and our Internet gateway.

[View full details \(/en-us/support/legal/sla/cache/\)](/en-us/support/legal/sla/cache/)

CDN

We guarantee that at least 99.9% of the time CDN will respond to client requests and deliver the requested content without error. We will review and accept data from any commercially reasonable independent measurement system that you choose to monitor your content. You must select a set of agents from the measurement system's list of standard agents that are generally available and represent at least five geographically diverse locations in major worldwide metropolitan areas

(excluding PR of China).

[View full details \(/en-us/support/legal/sla/cdn/\)](/en-us/support/legal/sla/cdn/)

Cloud Services and Virtual Machines

- For Cloud Services, we guarantee that when you deploy two or more role instances in different fault and upgrade domains, your Internet facing roles will have external connectivity at least 99.95% of the time.
- For all Internet facing Virtual Machines that have two or more instances deployed in the same Availability Set, we guarantee you will have external connectivity at least 99.95% of the time.

[View full details \(/en-us/support/legal/sla/virtual-machines/\)](/en-us/support/legal/sla/virtual-machines/)

DocumentDB

We guarantee at least 99.95% of the time we will successfully process requests to perform operations against DocumentDB Resources.

[View full details \(/en-us/support/legal/sla/documentdb/\)](/en-us/support/legal/sla/documentdb/)

ExpressRoute

We guarantee a minimum of 99.9% ExpressRoute dedicated circuit availability.

[View full details \(/en-us/support/legal/sla/expressroute/\)](/en-us/support/legal/sla/expressroute/)

HDInsight

For HDInsight, we guarantee that any HDInsight Cluster that you deploy will have external connectivity at least 99.9% of the time over a monthly billing cycle.

[View full details \(/en-us/support/legal/sla/hdinsight/\)](/en-us/support/legal/sla/hdinsight/)

IoT Hub

- For IoT Hub, we promise that at least 99.9% of the time deployed IoT hubs will be able to send messages to and receive messages from registered devices and the Service will be able to perform create, read, update, and delete operations on IoT hubs.
- No SLA is provided for the Free Tier of IoT Hub.

[View full details \(/en-us/support/legal/sla/iot-hub/\)](/en-us/support/legal/sla/iot-hub/)

Key Vault

We guarantee that we will process Key Vault transactions within 5 seconds at least 99.9% of the time.

[View full details \(/en-us/support/legal/sla/key-vault/\)](/en-us/support/legal/sla/key-vault/)

Machine Learning

- For the Request Response Service (RRS), we guarantee 99.95% availability of API transactions.
- For the Batch Execution Service (BES) and management APIs, we guarantee 99.9% availability of API transactions.

No SLA is provided for the Free tier of Machine Learning.

[View full details \(/en-us/support/legal/sla/machine-learning/\)](/en-us/support/legal/sla/machine-learning/)

Media Services

- For Media Services Encoding, we guarantee 99.9% availability of REST API transactions.
- For Streaming, we will successfully service requests with a 99.9% availability guarantee for existing media content when at least one Streaming Unit is purchased.
- For Live Channels, we guarantee that running Channels will have external connectivity at least 99.9% of the time.
- For Content Protection, we guarantee that we will successfully fulfill key requests at least 99.9% of the time.
- For Indexer, we will successfully service Indexer Task requests processed with an Encoding Reserved Unit 99.9% of the time.

[View full details \(/en-us/support/legal/sla/media-services/\)](/en-us/support/legal/sla/media-services/)

Mobile Engagement

We guarantee at least 99.9% availability of REST API calls to the Azure Mobile Engagement Service.

No SLA is provided for the Free tier.

[View full details \(/en-us/support/legal/sla/mobile-engagement/\)](/en-us/support/legal/sla/mobile-engagement/)

Mobile Services

We guarantee 99.9% availability of REST API calls to all provisioned Azure Mobile Services running in Standard and Premium tiers in a customer subscription. No SLA is provided for the Free tier of Mobile Services.

[View full details \(/en-us/support/legal/sla/mobile-services/\)](/en-us/support/legal/sla/mobile-services/)

Multi-Factor Authentication

We guarantee 99.9% availability of Azure Multi-Factor Authentication. The service is considered unavailable when it is unable to receive or process authentication requests for the Multi-Factor authentication provider deployed in a customer subscription.

[View full details \(/en-us/support/legal/sla/multi-factor-authentication/\)](/en-us/support/legal/sla/multi-factor-authentication/)

Operational Insights

We guarantee that at least 99.9% of the time, log data will be indexed within six hours of the data being queued for indexing by the Operational Insights Service.

No SLA is provided for the Free tier of Azure Operational Insights.

[View full details \(/en-us/support/legal/sla/operational-insights/\)](/en-us/support/legal/sla/operational-insights/)

RemoteApp

We guarantee at least 99.9% of the time users will have connectivity to their applications through the RemoteApp service. No SLA is provided for the Free tier of RemoteApp.

[View full details \(/en-us/support/legal/sla/remoteapp/\)](/en-us/support/legal/sla/remoteapp/)

Scheduler

We guarantee that at least 99.9% of the time all scheduled jobs will initiate within 30 minutes of their planned execution times.

[View full details \(/en-us/support/legal/sla/scheduler/\)](/en-us/support/legal/sla/scheduler/)

Search

We guarantee at least 99.9% availability for index query requests when an Azure Search Service Instance is configured with two or more replicas, and index update requests when an Azure Search Service Instance is configured with three or more replicas. No SLA is provided for the Free tier.

[View full details \(/en-us/support/legal/sla/search/\)](/en-us/support/legal/sla/search/)

Service Bus

- For Service Bus Relays, we guarantee that at least 99.9% of the time, properly configured applications will be able to establish a connection to a deployed Relay.
- For Service Bus Queues and Topics, we guarantee that at least 99.9% of the time, properly configured applications will be able to send or receive messages or perform other operations on a deployed Queue or Topic.

- For Service Bus Basic and Standard Notification Hub tiers, we guarantee that at least 99.9% of the time, properly configured applications will be able to send notifications or perform registration management operations with respect to a Notification Hub.
- For Event Hubs Basic and Standard tiers, we guarantee that at least 99.9% of the time, properly configured applications will be able to send or receive messages or perform other operations on the Event Hub.

[View full details \(/en-us/support/legal/sla/service-bus/\)](/en-us/support/legal/sla/service-bus/)

Site Recovery

- For each Protected Instance configured for On-Premises-to-On-Premises Failover, we guarantee at least 99.9% availability of the Site Recovery service.
- For each Protected Instance configured for On-Premises-to-Azure planned and unplanned Failover, we guarantee a four-hour Recovery Time Objective for unencrypted Protected Instances, and a six-hour Recovery Time Objective for encrypted Protected Instance, depending on the size of the Protected Instance.

[View full details \(/en-us/support/legal/sla/site-recovery/\)](/en-us/support/legal/sla/site-recovery/)

SQL Database

Web and Business Tiers

We guarantee at least 99.9% of the time customers will have connectivity between their Web or Business Microsoft Azure SQL Database and our Internet gateway.

Basic, Standard, and Premium Tiers

We guarantee at least 99.99% of the time customers will have connectivity between their Basic, Standard, or Premium Microsoft Azure SQL Database and our Internet gateway.

[View full details \(/en-us/support/legal/sla/sql-database/\)](/en-us/support/legal/sla/sql-database/)

Storage

- We guarantee that at least 99.99% of the time, we will successfully process requests to read data from Read Access-Geo Redundant Storage (RA-GRS) Accounts, provided that failed attempts to read data from the primary region are retried on the secondary region.
- We guarantee that at least 99.9% of the time, we will successfully process requests to read data from Locally Redundant Storage (LRS), Zone Redundant Storage (ZRS), and Geo Redundant Storage (GRS) Accounts.
- We guarantee that at least 99.9% of the time, we will successfully process requests to write data to

Locally Redundant Storage (LRS), Zone Redundant Storage (ZRS), and Geo Redundant Storage (GRS) Accounts and Read Access-Geo Redundant Storage (RA-GRS) Accounts.

[View full details \(/en-us/support/legal/sla/storage/\)](/en-us/support/legal/sla/storage/)

StorSimple

We guarantee at least 99.9% availability of the backup, cloud tiering, and restore functionality of the Azure StorSimple service.

[View full details \(/en-us/support/legal/sla/storsimple/\)](/en-us/support/legal/sla/storsimple/)

Stream Analytics

- We guarantee at least 99.9% availability of the Stream Analytics API.
- We guarantee that 99.9% of the time, deployed Stream Analytics jobs will be either processing data or available to process data.

[View full details \(/en-us/support/legal/sla/stream-analytics/\)](/en-us/support/legal/sla/stream-analytics/)

Traffic Manager

We guarantee that DNS queries will receive a valid response from at least one of our Azure Traffic Manager name server clusters at least 99.99% of the time.

[View full details \(/en-us/support/legal/sla/traffic-manager/\)](/en-us/support/legal/sla/traffic-manager/)

Visual Studio Team Services

- We guarantee at least 99.9% availability of Visual Studio Team Services for paid Visual Studio Team Services users to access the associated Visual Studio Team Services account.
- We guarantee at least 99.9% availability to execute build operations using the paid Visual Studio Team Services Build Service.
- We guarantee at least 99.9% availability to execute load testing operations using the paid Visual Studio Team Services Load Testing Service.
- We guarantee at least 99.9% availability to execute build and deployment operations using the paid Visual Studio Team Services Build & Deployment Service.

[View full details \(/en-us/support/legal/sla/visual-studio-team-services/\)](/en-us/support/legal/sla/visual-studio-team-services/)

VPN Gateway

We guarantee 99.9% availability for each VPN Gateway.

[View full details \(/en-us/support/legal/sla/vpn-gateway/\)](/en-us/support/legal/sla/vpn-gateway/)

Microsoft will provide at least 90 days' notice for adverse material changes to any of the SLAs listed above.

Availability for all Azure services is calculated over a monthly billing cycle. Click here (<http://go.microsoft.com/fwlink/?linkid=517017&clcid=0x409>) to download SLA for most Microsoft Azure Services. The SLA for Active Directory can be found here (<http://www.microsoftvolumelicensing.com/DocumentSearch.aspx?Mode=3&DocumentTypeId=37>).



Amazon Web Services: Risk and Compliance

January 2016

(Consult <http://aws.amazon.com/compliance/aws-whitepapers/>

for the latest version of this paper)

This document is intended to provide information to assist AWS customers with integrating AWS into their existing control framework supporting their IT environment. This document includes a basic approach to evaluating AWS controls and provides information to assist customers with integrating control environments. This document also addresses AWS-specific information around general cloud computing compliance questions.

Table of Contents

Risk and Compliance Overview	3
<i>Shared Responsibility Environment</i>	<i>3</i>
<i>Strong Compliance Governance</i>	<i>4</i>
Evaluating and Integrating AWS Controls	4
<i>AWS IT Control Information</i>	<i>5</i>
<i>AWS Global Regions</i>	<i>5</i>
AWS Risk and Compliance Program	6
<i>Risk Management</i>	<i>6</i>
<i>Control Environment</i>	<i>6</i>
<i>Information Security</i>	<i>7</i>
AWS Certifications, Programs, Reports, and Third-Party Attestations	7
<i>CJIS</i>	<i>7</i>
<i>CSA</i>	<i>7</i>
<i>Cyber Essentials Plus</i>	<i>8</i>
<i>DoD SRG Levels 2 and 4</i>	<i>8</i>
<i>FedRAMP SM</i>	<i>8</i>
<i>FERPA</i>	<i>9</i>
<i>FIPS 140-2</i>	<i>9</i>
<i>FISMA and DIACAP</i>	<i>9</i>
<i>GxP</i>	<i>10</i>
<i>HIPAA</i>	<i>10</i>
<i>IRAP</i>	<i>11</i>
<i>ISO 9001</i>	<i>11</i>
<i>ISO 27001</i>	<i>12</i>
<i>ISO 27017</i>	<i>14</i>
<i>ISO 27018</i>	<i>14</i>
<i>ITAR</i>	<i>15</i>
<i>MPAA</i>	<i>16</i>
<i>MTCS Tier 3 Certification</i>	<i>16</i>

<i>NIST</i>	16
<i>PCI DSS Level 1</i>	17
<i>SOC 1/ISAE 3402</i>	17
<i>SOC 2</i>	19
<i>SOC 3</i>	19
<i>Key Compliance Questions and AWS</i>	20
AWS Contact	24
Appendix A: CSA Consensus Assessments Initiative Questionnaire v3.0.1	25
Appendix B: AWS alignment with the Australian Signals Directorate (ASD) Cloud Computing Security Considerations	62
Appendix C: Glossary of Terms	82

Risk and Compliance Overview

AWS and its customers share control over the IT environment, both parties have responsibility for managing the IT environment. AWS' part in this shared responsibility includes providing its services on a highly secure and controlled platform and providing a wide array of security features customers can use. The customers' responsibility includes configuring their IT environments in a secure and controlled manner for their purposes. While customers don't communicate their use and configurations to AWS, AWS does communicate its security and control environment relevant to customers. AWS does this by doing the following:

- Obtaining industry certifications and independent third-party attestations described in this document
- Publishing information about the AWS security and control practices in whitepapers and web site content
- Providing certificates, reports, and other documentation directly to AWS customers under NDA (as required)

For a more detailed description of AWS security please see:

[AWS Security Center](https://aws.amazon.com/security/): <https://aws.amazon.com/security/>

For a more detailed description of AWS Compliance please see

[AWS Compliance page](https://aws.amazon.com/compliance/): <https://aws.amazon.com/compliance/>

Additionally, The [AWS Overview of Security Processes Whitepaper](#) covers AWS' general security controls and service-specific security.

Shared Responsibility Environment

Moving IT infrastructure to AWS services creates a model of shared responsibility between the customer and AWS. This shared model can help relieve customer's operational burden as AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall. Customers should carefully consider the services they choose as their responsibilities vary depending on the services used, the integration of those



services into their IT environment, and applicable laws and regulations. It is possible for customers to enhance security and/or meet their more stringent compliance requirements by leveraging technology such as host based firewalls, host based intrusion detection/prevention, encryption and key management. The nature of this shared responsibility also provides the flexibility and customer control that permits the deployment of solutions that meet industry-specific certification requirements.

This customer/AWS shared responsibility model also extends to IT controls. Just as the responsibility to operate the IT environment is shared between AWS and its customers, so is the management, operation and verification of IT controls shared. AWS can help relieve customer burden of operating controls by managing those controls associated with the physical infrastructure deployed in the AWS environment that may previously have been managed by the customer. As every customer is deployed differently in AWS, customers can take advantage of shifting management of certain IT controls to AWS which results in a (new) distributed control environment. Customers can then use the AWS control and compliance documentation available to them (described in the [AWS Certifications and Third-party Attestations](#) section of this document) to perform their control evaluation and verification procedures as required.

The next section provides an approach on how AWS customers can evaluate and validate their distributed control environment effectively.

Strong Compliance Governance

As always, AWS customers are required to continue to maintain adequate governance over the entire IT control environment regardless of how IT is deployed. Leading practices include an understanding of required compliance objectives and requirements (from relevant sources), establishment of a control environment that meets those objectives and requirements, an understanding of the validation required based on the organization's risk tolerance, and verification of the operating effectiveness of their control environment. Deployment in the AWS cloud gives enterprises different options to apply various types of controls and various verification methods.

Strong customer compliance and governance might include the following basic approach:

1. Review information available from AWS together with other information to understand as much of the entire IT environment as possible, and then document all compliance requirements.
2. Design and implement control objectives to meet the enterprise compliance requirements.
3. Identify and document controls owned by outside parties.
4. Verify that all control objectives are met and all key controls are designed and operating effectively.

Approaching compliance governance in this manner will help companies gain a better understanding of their control environment and will help clearly delineate the verification activities to be performed.

Evaluating and Integrating AWS Controls

AWS provides a wide range of information regarding its IT control environment to customers through white papers, reports, certifications, and other third-party attestations. This documentation assists customers in understanding the controls in place relevant to the AWS services they use and how those controls have been validated. This information also assists customers in their efforts to account for and to validate that controls in their extended IT environment are operating effectively.

Traditionally, the design and operating effectiveness of control objectives and controls are validated by internal and/or external auditors via process walkthroughs and evidence evaluation. Direct observation/verification, by the customer or customer's external auditor, is generally performed to validate controls. In the case where



service providers, such as AWS, are used, companies request and evaluate third-party attestations and certifications in order to gain reasonable assurance of the design and operating effectiveness of control objective and controls. As a result, although customer's key controls may be managed by AWS, the control environment can still be a unified framework where all controls are accounted for and are verified as operating effectively. Third-party attestations and certifications of AWS can not only provide a higher level of validation of the control environment, but may relieve customers of the requirement to perform certain validation work themselves for their IT environment in the AWS cloud.

AWS IT Control Information

AWS provides IT control information to customers in the following two ways:

1. **Specific control definition.** AWS customers are able to identify key controls managed by AWS. Key controls are critical to the customer's control environment and require an external attestation of the operating effectiveness of these key controls in order to comply with compliance requirements—such as the annual financial audit. For this purpose, AWS publishes a wide range of specific IT controls in its Service Organization Controls 1 (SOC 1) Type II report. The SOC 1 report, formerly the Statement on Auditing Standards (SAS) No. 70, Service Organizations report, is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). The SOC 1 audit is an in-depth audit of both the design and operating effectiveness of AWS' defined control objectives and control activities (which include control objectives and control activities over the part of the infrastructure AWS manages). "Type II" refers to the fact that each of the controls described in the report are not only evaluated for adequacy of design, but are also tested for operating effectiveness by the external auditor. Because of the independence and competence of AWS' external auditor, controls identified in the report should provide customers with a high level of confidence in AWS' control environment. AWS' controls can be considered designed and operating effectively for many compliance purposes, including Sarbanes-Oxley (SOX) Section 404 financial statement audits. Leveraging SOC 1 Type II reports is also generally permitted by other external certifying bodies (e.g., ISO 27001 auditors may request a SOC 1 Type II report in order to complete their evaluations for customers).

Other specific control activities relate to AWS' Payment Card Industry (PCI) and Federal Information Security Management Act (FISMA) compliance. As discussed below, AWS is compliant with FISMA Moderate standards and with the PCI Data Security Standard. These PCI and FISMA standards are very prescriptive and require independent validation that AWS adheres to the published standard.

2. **General control standard compliance.** If an AWS customer requires a broad set of control objectives to be met, evaluation of AWS' industry certifications may be performed. With the AWS ISO 27001 certification, AWS complies with a broad, comprehensive security standard and follows best practices in maintaining a secure environment. With the PCI Data Security Standard (PCI DSS), AWS complies with a set of controls important to companies that handle credit card information. With AWS' compliance with the FISMA standards, AWS complies with a wide range of specific controls required by US government agencies. Compliance with these general standards provides customers with in-depth information on the comprehensive nature of the controls and security processes in place and can be considered when managing compliance.

AWS Global Regions

Data centers are built in clusters in various global regions. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Frankfurt), EU (Ireland), Asia Pacific (Seoul) Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing), and South America (Sao Paulo).



AWS Risk and Compliance Program

AWS provides information about its risk and compliance program to enable customers to incorporate AWS controls into their governance framework. This information can assist customers in documenting a complete control and governance framework with AWS included as an important part of that framework.

Risk Management

AWS management has developed a strategic business plan which includes risk identification and the implementation of controls to mitigate or manage risks. AWS management re-evaluates the strategic business plan at least biannually. This process requires management to identify risks within its areas of responsibility and to implement appropriate measures designed to address those risks.

In addition, the AWS control environment is subject to various internal and external risk assessments. AWS' Compliance and Security teams have established an information security framework and policies based on the Control Objectives for Information and related Technology (COBIT) framework and have effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.1, and the National Institute of Standards and Technology (NIST) Publication 800-53 Rev 3 (Recommended Security Controls for Federal Information Systems). AWS maintains the security policy, provides security training to employees, and performs application security reviews. These reviews assess the confidentiality, integrity, and availability of data, as well as conformance to the information security policy.

AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. These scans are done in a manner for the health and viability of the underlying AWS infrastructure and are not meant to replace the customer's own vulnerability scans required to meet their specific compliance requirements. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these types of scans can be initiated by submitting a request via the [AWS Vulnerability / Penetration Testing Request Form](#).

Control Environment

AWS manages a comprehensive control environment that includes policies, processes and control activities that leverage various aspects of Amazon's overall control environment. This control environment is in place for the secure delivery of AWS' service offerings. The collective control environment encompasses the people, processes, and technology necessary to establish and maintain an environment that supports the operating effectiveness of AWS' control framework. AWS has integrated applicable cloud-specific controls identified by leading cloud computing industry bodies into the AWS control framework. AWS continues to monitor these industry groups for ideas on which leading practices can be implemented to better assist customers with managing their control environment.

The control environment at Amazon begins at the highest level of the Company. Executive and senior leadership play important roles in establishing the Company's tone and core values. Every employee is provided with the Company's Code of Business Conduct and Ethics and completes periodic training. Compliance audits are performed so that employees understand and follow the established policies.



The AWS organizational structure provides a framework for planning, executing and controlling business operations. The organizational structure assigns roles and responsibilities to provide for adequate staffing, efficiency of operations, and the segregation of duties. Management has also established authority and appropriate lines of reporting for key personnel. Included as part of the Company's hiring verification processes are education, previous employment, and, in some cases, background checks as permitted by law and regulation for employees commensurate with the employee's position and level of access to AWS facilities. The Company follows a structured on-boarding process to familiarize new employees with Amazon tools, processes, systems, policies and procedures.

Information Security

AWS has implemented a formal information security program designed to protect the confidentiality, integrity, and availability of customers' systems and data. AWS publishes a security whitepaper that is available on the public website that addresses how AWS can help customers secure their data.

AWS Certifications, Programs, Reports, and Third-Party Attestations

AWS engages with external certifying bodies and independent auditors to provide customers with considerable information regarding the policies, processes, and controls established and operated by AWS.

CJIS

AWS complies with the FBI's Criminal Justice Information Services (CJIS) standard. We sign CJIS security agreements with our customers, including allowing or performing any required employee background checks according to the [CJIS Security Policy](#).

Law enforcement customers (and partners who manage CJI) are taking advantage of AWS services to improve the security and protection of CJI data, using the advanced security services and features of AWS, such as activity logging ([AWS CloudTrail](#)), encryption of data in motion and at rest (S3's Server-Side Encryption with the option to bring your own key), comprehensive key management and protection ([AWS Key Management Service](#) and [CloudHSM](#)), and integrated permission management (IAM federated identity management, multi-factor authentication).

AWS has created a Criminal Justice Information Services (CJIS) [Workbook](#) in a security plan template format aligned to the CJIS Policy Areas. Additionally, a CJIS Whitepaper has been developed to help guide customers in their journey to cloud adoption.

Visit the CJIS Hub Page: <https://aws.amazon.com/compliance/cjis/>

CSA

In 2011, the Cloud Security Alliance (CSA) launched [STAR](#), an initiative to encourage transparency of security practices within cloud providers. The [CSA Security, Trust & Assurance Registry](#) (STAR) is a free, publicly accessible registry that documents the security controls provided by various cloud computing offerings, thereby helping users assess the security of cloud providers they currently use or are considering contracting with. [AWS is a CSA STAR registrant](#) and has completed the Cloud Security Alliance (CSA) Consensus Assessments Initiative Questionnaire (CAIQ). This CAIQ published by the CSA provides a way to reference and



document what security controls exist in AWS' Infrastructure as a Service offerings. The CAIQ provides 298 questions a cloud consumer and cloud auditor may wish to ask of a cloud provider.

See: [Appendix A: CSA Consensus Assessments Initiative Questionnaire v3.0.1](#)

Cyber Essentials Plus

[Cyber Essentials Plus](#) is a UK Government-backed, industry-supported certification scheme introduced in the UK to help organizations demonstrate operational security against common cyber-attacks.

It demonstrates the baseline controls AWS implements to mitigate the risk from common Internet-based threats, within the context of the UK Government's "[10 Steps to Cyber Security](#)". It is backed by industry, including the Federation of Small Businesses, the Confederation of British Industry and a number of insurance organizations that offer incentives for businesses holding this certification.

Cyber Essentials sets out the necessary technical controls; the related assurance framework shows how the independent assurance process works for Cyber Essentials Plus certification through an annual external assessment conducted by an accredited assessor. Due to the regional nature of the certification, the certification scope is limited to EU (Ireland) region.

DoD SRG Levels 2 and 4

[The Department of Defense \(DoD\) Cloud Security Model \(SRG\)](#) provides a formalized assessment and authorization process for cloud service providers (CSPs) to gain a DoD Provisional Authorization, which can subsequently be leveraged by DoD customers. A Provisional Authorization under the SRG provides a reusable certification that attests to our compliance with DoD standards, reducing the time necessary for a DoD mission owner to assess and authorize one of their systems for operation on AWS. AWS currently holds provisional authorizations at Levels 2 and 4 of the SRG.

Additional information of the security control baselines defined for [Levels 2, 4, 5, and 6 can be found at: \[http://iase.disa.mil/cloud_security/Pages/index.aspx\]\(http://iase.disa.mil/cloud_security/Pages/index.aspx\)](#).

Visit the DoD Hub Page: <https://aws.amazon.com/compliance/dod/>

FedRAMP SM

AWS is a Federal Risk and Authorization Management Program (FedRAMPSM) Compliant Cloud Service Provider. AWS has completed the testing performed by a FedRAMPSM accredited Third-Party Assessment Organization (3PAO) and has been granted two Agency Authority to Operate (ATOs) by the US Department of Health and Human Services (HHS) after demonstrating compliance with FedRAMPSM requirements at the Moderate impact level. All U.S. government agencies can leverage the AWS Agency ATO packages stored in the FedRAMPSM repository to evaluate AWS for their applications and workloads, provide authorizations to use AWS, and transition workloads into the AWS environment. The two FedRAMPSM Agency ATOs encompass all U.S. regions (the AWS GovCloud (US) region and the AWS US East/West regions).

The following services are in the accreditation boundary for the regions stated above:



- [Amazon Redshift](#). Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools.
- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#). Amazon EC2 provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.
- [Amazon Simple Storage Service \(S3\)](#). Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.
- [Amazon Virtual Private Cloud \(VPC\)](#). Amazon VPC provides the ability for you to provision a logically isolated section of AWS where you can launch AWS resources in a virtual network that you define.
- [Amazon Elastic Block Store \(EBS\)](#). Amazon EBS provides highly available, highly reliable, predictable storage volumes that can be attached to a running Amazon EC2 instance and exposed as a device within the instance.
- [AWS Identity and Access Management \(IAM\)](#). IAM enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources.

For more information on AWS FedRAMPsm compliance please see the [AWS FedRAMPsm FAQs](#) at: <https://aws.amazon.com/compliance/fedramp/>

FERPA

[The Family Educational Rights and Privacy Act](#) (FERPA) (20 U.S.C. § 1232g; 34 CFR Part 99) is a Federal law that protects the privacy of student education records. The law applies to all schools that receive funds under an applicable program of the U.S. Department of Education. FERPA gives parents certain rights with respect to their children's education records. These rights transfer to the student when he or she reaches the age of 18, or attends a school beyond the high school level. Students to whom the rights have transferred are "eligible students."

AWS enables covered entities and their business associates subject to FERPA to leverage the secure AWS environment to process, maintain, and store protected education information.

AWS also offers a [FERPA-focused whitepaper](#) for customers interested in learning more about how they can leverage AWS for the processing and storage of educational data.

The "[FERPA Compliance on AWS Whitepaper](#)" outlines how companies can use AWS to process systems that facilitate FERPA compliance:
https://do.awsstatic.com/whitepapers/compliance/AWS_FERPA_Whitepaper.pdf

FIPS 140-2

[The Federal Information Processing Standard \(FIPS\) Publication 140-2](#) is a US government security standard that specifies the security requirements for cryptographic modules protecting sensitive information. To support customers with FIPS 140-2 requirements, SSL terminations in [AWS GovCloud \(US\)](#) operate using FIPS 140-2 validated hardware. AWS works with AWS GovCloud (US) customers to provide the information they need to help manage compliance when using the [AWS GovCloud \(US\) environment](#).

FISMA and DIACAP



AWS enables US government agencies to achieve and sustain compliance with the Federal Information Security Management Act ([FISMA](#)). The AWS infrastructure has been evaluated by independent assessors for a variety of government systems as part of their system owners' approval process. Numerous Federal Civilian and Department of Defense (DoD) organizations have successfully achieved security authorizations for systems hosted on AWS in accordance with the Risk Management Framework (RMF) process defined in NIST 800-37 and DoD Information Assurance Certification and Accreditation Process ([DIACAP](#)).

GxP

GxP is an acronym that refers to the regulations and guidelines applicable to life sciences organizations that make food and medical products such as drugs, medical devices, and medical software applications. The overall intent of GxP requirements is to ensure that food and medical products are safe for consumers and to ensure the integrity of data used to make product-related safety decisions.

AWS offers a [GxP whitepaper](#) which details a comprehensive approach for using AWS for GxP systems. This whitepaper provides guidance for using [AWS Products in the context of GxP](#) and the content has been developed in conjunction with AWS pharmaceutical and medical device customers, as well as software partners, who are currently using AWS Products in their validated GxP systems.

For more information on the GxP on AWS [please contact AWS Sales and Business Development](#).

For additional information please see our GxP Compliance FAQs:

<https://aws.amazon.com/compliance/gxp-part-11-annex-11/>

HIPAA

AWS enables covered entities and their business associates subject to the U.S. Health Insurance Portability and Accountability Act (HIPAA) to leverage the secure AWS environment to process, maintain, and store protected health information and AWS will be signing business associate agreements with such customers. AWS also offers a HIPAA-focused whitepaper for customers interested in learning more about how they can leverage AWS for the processing and storage of health information. The [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) whitepaper outlines how companies can use AWS to process systems that facilitate HIPAA and Health Information Technology for Economic and Clinical Health (HITECH) compliance.

Customers may use any AWS service in an account designated as a HIPAA account, but they should only process, store and transmit PHI in the HIPAA-eligible services defined in the BAA. There are nine HIPAA-eligible services today, including:

- [Amazon DynamoDB](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Cloud Compute \(EC2\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon Glacier](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(Amazon RDS\)](#) using only MySQL and Oracle engines
- [Amazon Simple Storage Service \(S3\)](#)



AWS follows a standards-based risk management program to ensure that the HIPAA-eligible services specifically support the security, control, and administrative processes required under HIPAA. Using these services to store and process PHI allows our customers and AWS to address the HIPAA requirements applicable to our utility-based operating model. AWS prioritizes and adds new eligible services based on customer demand.

For additional information please see our HIPAA Compliance FAQs:

<https://aws.amazon.com/compliance/hipaa-compliance/>

Architecting for HIPAA Security and Compliance on Amazon Web Services:

https://do.awsstatic.com/whitepapers/compliance/AWS_HIPAA_Compliance_Whitepaper.pdf

IRAP

The Information Security Registered Assessors Program (IRAP) enables Australian government customers to validate that appropriate controls are in place and determine the appropriate responsibility model for addressing the needs of the Australian Signals Directorate (ASD) Information Security Manual (ISM).

Amazon Web Services **has completed an independent assessment** that has determined all applicable ISM controls are in place relating to the processing, storage and transmission of Unclassified (DLM) for the AWS Sydney Region.

IRAP Compliance FAQs:

<https://aws.amazon.com/compliance/irap/>

For more information see: **Appendix B: AWS alignment with the Australian Signals Directorate (ASD) Cloud Computing Security Considerations**

ISO 9001

AWS has achieved ISO 9001 certification, AWS' ISO 9001 certification directly supports customers who develop, migrate and operate their quality-controlled IT systems in the AWS cloud. Customers can leverage AWS' compliance reports as evidence for their own ISO 9001 programs and industry-specific quality programs, such as GxP in life sciences, ISO 13485 in medical devices, AS9100 in aerospace, and ISO/TS 16949 in automotive. AWS customers who don't have quality system requirements will still benefit from the additional assurance and transparency that an ISO 9001 certification provides.

The ISO 9001 certification covers the quality management system over a specified scope of AWS services and Regions of operations (below) and services including:

- [AWS CloudFormation](#)
- [AWS Cloud Hardware Security Model \(HSM\)](#)
- [Amazon CloudFront](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [AWS Directory Service](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 VM Import/Export](#)
- [AWS Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Cloud Compute \(EC2\)](#)



- [Amazon EC2 Container Service \(ECS\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [AWS Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [AWS Storage Gateway](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [AWS WAF - Web Application Firewall](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- The underlying physical infrastructure and the AWS Management Environment

AWS' ISO 9001 accreditation covers AWS Regions including US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), South America (Sao Paulo), EU (Ireland), EU (Frankfurt) and Asia Pacific (Singapore), Asia Pacific (Sydney), and Asia Pacific (Tokyo).

ISO 9001:2008 is a global standard for managing the quality of products and services. The 9001 standard outlines a quality management system based on eight principles defined by the International Organization for Standardization (ISO) Technical Committee for Quality Management and Quality Assurance. They include:

- Customer focus
- Leadership
- Involvement of people
- Process approach
- System approach to management
- Continual Improvement
- Factual approach to decision-making
- Mutually beneficial supplier relationships

The AWS ISO 9001 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_9001_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 9001 certification at:

<https://aws.amazon.com/compliance/iso-9001-faqs/>

ISO 27001

AWS has achieved ISO 27001 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services including:



- [AWS CloudFormation](#)
- [Amazon CloudFront](#)
- [AWS Cloudtrail](#)
- [AWS Directory Service](#)
- [Amazon DynamoDB](#)
- [AWS Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Cloud Compute \(EC2\)](#)
- [Amazon EC2 Container Service \(ECS\)](#)
- [AWS Direct Connect](#)
- [Amazon EC2 VM Import/Export](#)
- [AWS Cloud Hardware Security Model \(HSM\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [AWS Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [AWS Storage Gateway](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [AWS WAF - Web Application Firewall](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- The underlying physical infrastructure (including GovCloud) and the AWS Management Environment

ISO 27001/27002 is a widely-adopted global security standard that sets out requirements and best practices for a systematic approach to managing company and customer information that's based on periodic risk assessments appropriate to ever-changing threat scenarios. In order to achieve the certification, a company must show it has a systematic and ongoing approach to managing information security risks that affect the confidentiality, integrity, and availability of company and customer information. This certification reinforces Amazon's commitment to providing significant information regarding our security controls and practices.

AWS' ISO 27001 accreditation covers AWS Regions including US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), South America (Sao Paulo), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Sydney), and Asia Pacific (Tokyo).

The AWS ISO 27001 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_27001_global_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 27001 certification at:

<https://aws.amazon.com/compliance/iso-27001-faqs/>



ISO 27017

ISO 27017 is the newest code of practice released by the International Organization for Standardization (ISO). It provides implementation guidance on information security controls that specifically relate to cloud services.

AWS has achieved ISO 27017 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services including:

- [Amazon CloudFront](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 Container Service \(ECS\)](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [Amazon SimpleDB](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [AWS Directory Service](#)
- [AWS Elastic Beanstalk](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Storage Gateway](#)
- [AWS WAF \(Web Application Firewall\)](#)
- [Elastic Load Balancing](#)
- [VM Import/Export](#)

The AWS ISO 27017 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_27017_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 27017 certification at:

<https://aws.amazon.com/compliance/iso-27017-faqs/>

ISO 27018



ISO 27018 is the first International code of practice that focuses on protection of personal data in the cloud. It is based on ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls applicable to public cloud Personally Identifiable Information (PII). It also provides a set of additional controls and associated guidance intended to address public cloud PII protection requirements not addressed by the existing ISO 27002 control set.

AWS has achieved ISO 27018 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services including:

- [Amazon CloudFront](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 Container Service \(ECS\)](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [Amazon SimpleDB](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [AWS Directory Service](#)
- [AWS Elastic Beanstalk](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Storage Gateway](#)
- [AWS WAF \(Web Application Firewall\)](#)
- [Elastic Load Balancing](#)
- [VM Import/Export](#)

The AWS ISO 27018 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_27018_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 27018 certification at:

<https://aws.amazon.com/compliance/iso-27018-faqs/>

ITAR

The [AWS GovCloud \(US\)](#) region supports US International Traffic in Arms Regulations ([ITAR](#)) compliance. As a part of managing a comprehensive ITAR compliance program, companies subject to ITAR export regulations must control unintended exports by restricting access to protected data to US Persons and restricting physical location of that data to the US. AWS GovCloud (US) provides an environment physically located in the US and where access by AWS Personnel is limited to US Persons, thereby allowing qualified companies to transmit, process, and store protected articles and data subject to ITAR restrictions. The AWS GovCloud (US) environment has been audited by an independent third-party to validate the proper controls are in place to support customer export compliance programs for this requirement.

MPAA

The Motion Picture Association of America (MPAA) has established a set of best practices for securely storing, processing and delivering protected media and content (<http://www.fightfilmtheft.org/facility-security-program.html>). Media companies use these best practices as a way to assess risk and security of their content and infrastructure. AWS has demonstrated alignment with the MPAA best practices and the AWS infrastructure is compliant with all applicable MPAA infrastructure controls. While the MPAA does not offer a “certification,” media industry customers can use the AWS MPAA documentation to augment their risk assessment and evaluation of MPAA-type content on AWS.

See the [AWS Compliance MPAA hub page](#) for additional details:
<https://aws.amazon.com/compliance/mpaa/>

MTCS Tier 3 Certification

The [Multi-Tier Cloud Security \(MTCS\)](#) is an operational Singapore security management Standard (SPRING SS 584:2013), based on ISO 27001/02 Information Security Management System (ISMS) standards. The certification assessment requires us to:

- Systematically evaluate our information security risks, taking into account the impact of company threats and vulnerabilities
- Design and implement a comprehensive suite of information security controls and other forms of risk management to address company and architecture security risks
- Adopt an overarching management process to ensure that the information security controls meet the our information security needs on an ongoing basis

View the MTCS Hub Page at:
<https://aws.amazon.com/compliance/aws-multitiered-cloud-security-standard-certification/>

NIST

In June 2015 The National Institute of Standards and Technology (NIST) released guidelines [800-171](#), "Final Guidelines for Protecting Sensitive Government Information Held by Contractors". This guidance is applicable to the protection of Controlled Unclassified Information (CUI) on nonfederal systems.

AWS is already compliant with these guidelines, and customers can effectively comply with NIST 800-171 immediately. NIST 800-171 outlines a subset of the NIST 800-53 requirements, a guideline under which AWS has already been audited under the FedRAMP program. The FedRAMP Moderate security control baseline is more rigorous than the recommended requirements established in Chapter 3 of 800-171, and includes a significant number of security controls above and beyond those required of FISMA Moderate systems that



protect CUI data. A detailed mapping is available in the [NIST Special Publication 800-171](#), starting on page D2 (which is page 37 in the PDF).

PCI DSS Level 1

AWS is Level 1 compliant under the Payment Card Industry (PCI) Data Security Standard (DSS). Customers can run applications on our PCI-compliant technology infrastructure for storing, processing, and transmitting credit card information in the cloud. In February 2013, the PCI Security Standards Council released PCI DSS Cloud Computing Guidelines. These guidelines provide customers who are managing a cardholder data environment with considerations for maintaining PCI DSS controls in the cloud. AWS has incorporated the PCI DSS Cloud Computing Guidelines into the AWS PCI Compliance Package for customers. The AWS PCI Compliance Package includes the AWS PCI Attestation of Compliance (AoC), which shows that AWS has been successfully validated against standards applicable to a Level 1 service provider under PCI DSS Version 3.1, and the AWS PCI Responsibility Summary, which explains how compliance responsibilities are shared between AWS and our customers in the cloud.

The following services are in scope for PCI DSS Level 1:

- [Auto Scaling](#)
- [AWS CloudFormation](#)
- [Amazon CloudFront](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [Amazon DynamoDB](#)
- [AWS Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon Glacier](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Workflow Service SWF](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- The underlying physical infrastructure (including GovCloud) and the AWS Management Environment

The latest scope of services and regions for the AWS PCI DSS Level 1 certification can be found at: <https://aws.amazon.com/compliance/pci-dss-level-1-faqs/>

SOC 1/ISAE 3402

Amazon Web Services publishes a Service Organization Controls 1 (SOC 1), Type II report. The audit for this report is conducted in accordance with American Institute of Certified Public Accountants (AICPA): AT 801



(formerly SSAE 16) and the International Standards for Assurance Engagements No. 3402 (ISAE 3402). This dual-standard report is intended to meet a broad range of financial auditing requirements for U.S. and international auditing bodies. The SOC 1 report audit attests that AWS' control objectives are appropriately designed and that the individual controls defined to safeguard customer data are operating effectively. This report is the replacement of the Statement on Auditing Standards No. 70 (SAS 70) Type II Audit report.

The AWS SOC 1 control objectives are provided here. The report itself identifies the control activities that support each of these objectives and the independent auditor's results of their testing procedures of each control.

Objective Area	Objective Description
Security Organization	Controls provide reasonable assurance that information security policies have been implemented and communicated throughout the organization.
Employee User Access	Controls provide reasonable assurance that procedures have been established so that Amazon employee user accounts are added, modified and deleted in a timely manner and reviewed on a periodic basis.
Logical Security	Controls provide reasonable assurance that policies and mechanisms are in place to appropriately restrict unauthorized internal and external access to data and customer data is appropriately segregated from other customers.
Secure Data Handling	Controls provide reasonable assurance that data handling between the customer's point of initiation to an AWS storage location is secured and mapped accurately.
Physical Security and Environmental Protection	Controls provide reasonable assurance that physical access to data centers is restricted to authorized personnel and that mechanisms are in place to minimize the effect of a malfunction or physical disaster to data center facilities.
Change Management	Controls provide reasonable assurance that changes (including emergency / non-routine and configuration) to existing IT resources are logged, authorized, tested, approved and documented.
Data Integrity, Availability and Redundancy	Controls provide reasonable assurance that data integrity is maintained through all phases including transmission, storage and processing.
Incident Handling	Controls provide reasonable assurance that system incidents are recorded, analyzed, and resolved.

The SOC 1 reports are designed to focus on controls at a service organization that are likely to be relevant to an audit of a user entity's financial statements. As AWS' customer base is broad, and the use of AWS services is equally as broad, the applicability of controls to customer financial statements varies by customer. Therefore, the AWS SOC 1 report is designed to cover specific key controls likely to be required during a financial audit, as well as covering a broad range of IT general controls to accommodate a wide range of usage and audit scenarios. This allows customers to leverage the AWS infrastructure to store and process critical data, including that which is integral to the financial reporting process. AWS periodically reassesses the selection of these controls to consider customer feedback and usage of this important audit report.

AWS' commitment to the SOC 1 report is ongoing, and AWS will continue the process of periodic audits. The SOC 1 report scope covers:

- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)



- [AWS Direct Connect](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 VM Import/Export](#)
- [Amazon Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon ElastiCache](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon Glacier](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow \(SWF\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [AWS Storage Gateway](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [Amazon Workspaces](#)

SOC 2

In addition to the SOC 1 report, AWS publishes a Service Organization Controls 2 (SOC 2), Type II report. Similar to the SOC 1 in the evaluation of controls, the SOC 2 report is an attestation report that expands the evaluation of controls to the criteria set forth by the American Institute of Certified Public Accountants (AICPA) Trust Services Principles. These principles define leading practice controls relevant to security, availability, processing integrity, confidentiality, and privacy applicable to service organizations such as AWS. The AWS SOC 2 is an evaluation of the design and operating effectiveness of controls that meet the criteria for the security and availability principles set forth in the AICPA's Trust Services Principles criteria. This report provides additional transparency into AWS security and availability based on a pre-defined industry standard of leading practices and further demonstrates AWS' commitment to protecting customer data. The SOC 2 report scope covers the same services covered in the SOC 1 report. See the SOC 1 description above for the in-scope services.

SOC 3

AWS publishes a Service Organization Controls 3 (SOC 3) report. The SOC 3 report is a publically-available summary of the AWS SOC 2 report. The report includes the external auditor's opinion of the operation of controls (based on the [AICPA's Security Trust Principles](#) included in the SOC 2 report), the assertion from AWS management regarding the effectiveness of controls, and an overview of AWS Infrastructure and Services. The AWS SOC 3 report includes all AWS data centers worldwide that support in-scope services. This is a great resource for customers to validate that AWS has obtained external auditor assurance without going through the process to request a SOC 2 report. The SOC 3 report scope covers the same services covered in the SOC 1 report. See the SOC 1 description above for the in-scope services. [View the AWS SOC 3 report here.](#)



Key Compliance Questions and AWS

This section addresses generic cloud computing compliance questions specifically for AWS. These common compliance questions listed may be of interest when evaluating and operating in a cloud computing environment and may assist in AWS customers' control management efforts.

Ref	Cloud Computing Question	AWS Information
1	Control ownership. Who owns which controls for cloud-deployed infrastructure?	For the portion deployed into AWS, AWS controls the physical components of that technology. The customer owns and controls everything else, including control over connection points and transmissions. To help customers better understand what controls we have in place and how effectively they are operating, we publish a SOC 1 Type II report with controls defined around EC2, S3 and VPC, as well as detailed physical security and environmental controls. These controls are defined at a high level of specificity that should meet most customer needs. AWS customers that have signed a non-disclosure agreement with AWS may request a copy of the SOC 1 Type II report.
2	Auditing IT. How can auditing of the cloud provider be accomplished?	Auditing for most layers and controls above the physical controls remains the responsibility of the customer. The definition of AWS-defined logical and physical controls is documented in the SOC 1 Type II report, and the report is available for review by audit and compliance teams. AWS ISO 27001 and other certifications are also available for auditors to review.
3	Sarbanes-Oxley compliance. How is SOX compliance achieved if in-scope systems are deployed in the cloud provider environment?	If a customer processes financial information in the AWS cloud, the customer's auditors may determine that some AWS systems come into scope for Sarbanes-Oxley (SOX) requirements. The customer's auditors must make their own determination regarding SOX applicability. Because most of the logical access controls are managed by customer, the customer is best positioned to determine if its control activities meet relevant standards. If the SOX auditors request specifics regarding AWS' physical controls, they can reference the AWS SOC 1 Type II report which details the controls that AWS provides.
4	HIPAA compliance. Is it possible to meet HIPAA compliance requirements while deployed in the cloud provider environment?	HIPAA requirements apply to and are controlled by the AWS customer. The AWS platform allows for the deployment of solutions that meet industry-specific certification requirements such as HIPAA. Customers can use AWS services to maintain a security level that is equivalent or greater than those required to protect electronic health records. Customers have built healthcare applications compliant with HIPAA's Security and Privacy Rules on AWS. AWS provides additional information about HIPAA compliance on its web site, including a whitepaper on this topic .
5	GLBA compliance. Is it possible to meet GLBA certification requirements while deployed in the cloud provider environment?	Most GLBA requirements are controlled by the AWS customer. AWS provides means for customers to protect data, manage permissions, and build GLBA-compliant applications on AWS infrastructure. If the customer requires specific assurance that physical security controls are operating effectively, they can reference the AWS SOC 1 Type II report as relevant.

Ref	Cloud Computing Question	AWS Information
6	Federal regulation compliance. Is it possible for a US Government agency to be compliant with security and privacy regulations while deployed in the cloud provider environment?	US Federal agencies can be compliant under a number of compliance standards, including the Federal Information Security Management Act (FISMA) of 2002, Federal Risk and Authorization Management Program (FedRAMP), the Federal Information Processing Standard (FIPS) Publication 140-2, and the International Traffic in Arms Regulations (ITAR). Compliance with other laws and statutes may also be accommodated depending on the requirements set forth in the applicable legislation.
7	Data location. Where does customer data reside?	AWS customers designate in which physical region their data and their servers will be located. Data replication for S3 data objects is done within the regional cluster in which the data is stored and is not replicated to other data center clusters in other regions. AWS customers designate in which physical region their data and their servers will be located. AWS will not move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Frankfurt), EU (Ireland), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing), and South America (Sao Paulo).
8	E-Discovery. Does the cloud provider meet the customer's needs to meet electronic discovery procedures and requirements?	AWS provides infrastructure, and customers manage everything else, including the operating system, the network configuration, and the installed applications. Customers are responsible for responding appropriately to legal procedures involving the identification, collection, processing, analysis, and production of electronic documents they store or process using AWS. Upon request, AWS may work with customers who require AWS' assistance in legal proceedings.
9	Data center tours. Are data center tours by customers allowed by the cloud provider?	No. Due to the fact that our data centers host multiple customers, AWS does not allow data center tours by customers, as this exposes a wide range of customers to physical access of a third party. To meet this customer need, an independent and competent auditor validates the presence and operation of controls as part of our SOC 1 Type II report. This broadly accepted third-party validation provides customers with the independent perspective of the effectiveness of controls in place. AWS customers that have signed a non-disclosure agreement with AWS may request a copy of the SOC 1 Type II report. Independent reviews of data center physical security is also a part of the ISO 27001 audit, the PCI assessment, ITAR audit, and the FedRAMP sm testing programs.
10	Third-party access. Are third parties allowed access to the cloud provider data centers?	AWS strictly controls access to data centers, even for internal employees. Third parties are not provided access to AWS data centers except when explicitly approved by the appropriate AWS data center manager per the AWS access policy. See the SOC 1 Type II report for specific controls related to physical access, data center access authorization, and other related controls.

Ref	Cloud Computing Question	AWS Information
11	Privileged actions. Are privileged actions monitored and controlled?	Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored. In addition, customer data is and server instances are logically isolated from other customers by default. Privileged user access control is reviewed by an independent auditor during the AWS SOC 1, ISO 27001, PCI, ITAR, and FedRAMP sm audits.
12	Insider access. Does the cloud provider address the threat of inappropriate insider access to customer data and applications?	AWS provides specific SOC 1 controls to address the threat of inappropriate insider access, and the public certification and compliance initiatives covered in this document address insider access. All certifications and third-party attestations evaluate logical access preventative and detective controls. In addition, periodic risk assessments focus on how insider access is controlled and monitored.
13	Multi-tenancy. Is customer segregation implemented securely?	The AWS environment is a virtualized, multi-tenant environment. AWS has implemented security management processes, PCI controls, and other security controls designed to isolate each customer from other customers. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture has been validated by an independent PCI Qualified Security Assessor (QSA) and was found to be in compliance with all requirements of PCI DSS version 3.1 published in April 2015. Note that AWS also has single-tenancy options. Dedicated Instances are Amazon EC2 instances launched within your Amazon Virtual Private Cloud (Amazon VPC) that run hardware dedicated to a single customer. Dedicated Instances let you take full advantage of the benefits of Amazon VPC and the AWS cloud while isolating your Amazon EC2 compute instances at the hardware level.
14	Hypervisor vulnerabilities. Has the cloud provider addressed known hypervisor vulnerabilities?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. See the AWS security whitepaper for more information on the Xen hypervisor and instance isolation.
15	Vulnerability management. Are systems patched appropriately?	AWS is responsible for patching systems supporting the delivery of service to customers, such as the hypervisor and networking services. This is done as required per AWS policy and in accordance with ISO 27001, NIST, and PCI requirements. Customers control their own guest operating systems, software and applications and are therefore responsible for patching their own systems.
16	Encryption. Do the provided services support encryption?	Yes. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB, and EC2. IPsec tunnels to VPC are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Customers may also use third-party encryption technologies. Refer to the AWS Security white paper for more information.

Ref	Cloud Computing Question	AWS Information
17	Data ownership. What are the cloud provider's rights over customer data?	AWS customers retain control and ownership of their data. AWS errs on the side of protecting customer privacy and is vigilant in determining which law enforcement requests we must comply with. AWS does not hesitate to challenge orders from law enforcement if we think the orders lack a solid basis.
18	Data isolation. Does the cloud provider adequately isolate customer data?	All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Amazon S3 provides advanced data access controls. Please see the AWS security whitepaper for more information about specific data services' security.
19	Composite services. Does the cloud provider layer its service with other providers' cloud services?	AWS does not leverage any third-party cloud providers to deliver AWS services to customers.
20	Physical and environmental controls. Are these controls operated by the cloud provider specified?	Yes. These are specifically outlined in the SOC 1 Type II report. In addition, other certifications AWS supports such as ISO 27001 and FedRAMP sm require best practice physical and environmental controls.
21	Client-side protection. Does the cloud provider allow customers to secure and manage access from clients, such as PC and mobile devices?	Yes. AWS allows customers to manage client and mobile applications to their own requirements.
22	Server security. Does the cloud provider allow customers to secure their virtual servers?	Yes. AWS allows customers to implement their own security architecture. See the AWS security whitepaper for more details on server and network security.
23	Identity and Access Management. Does the service include IAM capabilities?	AWS has a suite of identity and access management offerings, allowing customers to manage user identities, assign security credentials, organize users in groups, and manage user permissions in a centralized way. Please see the AWS web site for more information.
24	Scheduled maintenance outages. Does the provider specify when systems will be brought down for maintenance?	AWS does not require systems to be brought offline to perform regular maintenance and system patching. AWS' own maintenance and system patching generally do not impact customers. Maintenance of instances themselves is controlled by the customer.
25	Capability to scale. Does the provider allow customers to scale beyond the original agreement?	The AWS cloud is distributed, highly secure and resilient, giving customers massive scale potential. Customers may scale up or down, paying for only what they use.
26	Service availability. Does the provider commit to a high level of availability?	AWS does commit to high levels of availability in its service level agreements (SLA). For example, Amazon EC2 commits to annual uptime percentage of at least 99.95% during the service year. Amazon S3 commits to monthly uptime percentage of at least 99.9%. Service credits are provided in the case these availability metrics are not met.

Ref	Cloud Computing Question	AWS Information
27	Distributed Denial Of Service (DDoS) attacks. How does the provider protect their service against DDoS attacks?	The AWS network provides significant protection against traditional network security issues and the customer can implement further protection. See the AWS Security Whitepaper for more information on this topic, including a discussion of DDoS attacks.
28	Data portability. Can the data stored with a service provider be exported by customer request?	AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport.
29	Service provider business continuity. Does the service provider operate a business continuity program?	AWS does operate a business continuity program. Detailed information is provided in the AWS Security Whitepaper.
30	Customer business continuity. Does the service provider allow customers to implement a business continuity plan?	AWS provides customers with the capability to implement a robust continuity plan, including the utilization of frequent server instance back-ups, data redundancy replication, and multi-region/availability zone deployment architectures.
31	Data durability. Does the service specify data durability?	Amazon S3 provides a highly durable storage infrastructure. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 Region. Once stored, Amazon S3 maintains the durability of objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. Data stored in S3 is designed to provide 99.99999999% durability and 99.99% availability of objects over a given year.
32	Backups. Does the service provide backups to tapes?	AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 service is designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy. For information on data durability and redundancy, please refer to the AWS web site.
33	Price increases. Will the service provider raise prices unexpectedly?	AWS has a history of frequently reducing prices as the cost to provide these services reduces over time. AWS has reduced prices consistently over the past several years.
34	Sustainability. Does the service provider company have long term sustainability potential?	AWS is a leading cloud provider and is a long-term business strategy of Amazon.com. AWS has very high long term sustainability potential.

AWS Contact

Customers can request the reports and certifications produced by our third-party auditors or can request more information about AWS Compliance by contacting [AWS Sales and Business Development](#). The representative will route customers to the proper team depending on nature of the inquiry. For additional information on AWS Compliance, see the [AWS Compliance site](#) or send questions directly to awscompliance@amazon.com.



Appendix A: CSA Consensus Assessments Initiative Questionnaire v3.0.1

The Cloud Security Alliance (CSA) is a “not-for-profit organization with a mission to promote the use of best practices for providing security assurance within Cloud Computing, and to provide education on the uses of Cloud Computing to help secure all other forms of computing.” [Reference

<https://cloudsecurityalliance.org/about/>] A wide range of industry security practitioners, corporations, and associations participate in this organization to achieve its mission.

The CSA Consensus Assessments Initiative Questionnaire provides a set of questions the CSA anticipates a cloud consumer and/or a cloud auditor would ask of a cloud provider. It provides a series of security, control, and process questions which can then be used for a wide range of uses, including cloud provider selection and security evaluation. AWS has completed this questionnaire with the answers below.

Control Group	CID	Consensus Assessment Questions	AWS Response
Application & Interface Security <i>Application Security</i>	AIS-01.1	Do you use industry standards (Build Security in Maturity Model [BSIMM] benchmarks, Open Group ACS Trusted Technology Provider Framework, NIST, etc.) to build in security for your Systems/Software Development Lifecycle (SDLC)?	The AWS system development lifecycle incorporates industry best practices which include formal design reviews by the AWS Security Team, threat modeling and completion of a risk assessment. Refer to the AWS Overview of Security Processes for further details. AWS has in place procedures to manage new development of resources. Refer to ISO 27001 standard, Annex A, domain 14 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	AIS-01.2	Do you use an automated source code analysis tool to detect security defects in code prior to production?	
	AIS-01.3	Do you use manual source-code analysis to detect security defects in code prior to production?	
	AIS-01.4	Do you verify that all of your software suppliers adhere to industry standards for Systems/Software Development Lifecycle (SDLC) security?	
	AIS-01.5	(SaaS only) Do you review your applications for security vulnerabilities and address any issues prior to deployment to production?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Application & Interface Security <i>Customer Access Requirements</i>	AIS-02.1	Are all identified security, contractual and regulatory requirements for customer access contractually addressed and remediated prior to granting customers access to data, assets and information systems?	AWS Customers retain responsibility to ensure their usage of AWS is in compliance with applicable laws and regulations. AWS communicates its security and control environment to customers through industry certifications and third-party attestations, white papers (available at http://aws.amazon.com/compliance) and providing certifications, reports and other relevant documentation directly to AWS Customers.
	AIS-02.2	Are all requirements and trust levels for customers' access defined and documented?	
Application & Interface Security <i>Data Integrity</i>	AIS-03.1	Are data input and output integrity routines (i.e., reconciliation and edit checks) implemented for application interfaces and databases to prevent manual or systematic processing errors or corruption of data?	AWS data integrity controls as described in AWS SOC reports illustrates the data integrity controls maintained through all phases including transmission, storage and processing. In addition, refer to ISO 27001 standard, Annex A, domain 14 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Application & Interface Security <i>Data Security / Integrity</i>	AIS-04.1	Is your Data Security Architecture designed using an industry standard (e.g., CDSA, MULITSAFE, CSA Trusted Cloud Architectural Standard, FedRAMP, CAESARS)?	AWS Data Security Architecture was designed to incorporate industry leading practices. Refer to AWS Certifications, reports and whitepapers for additional details on the various leading practices that AWS adheres to (available at http://aws.amazon.com/compliance).
Audit Assurance & Compliance <i>Audit Planning</i>	AAC-01.1	Do you produce audit assertions using a structured, industry accepted format (e.g., CloudAudit/A6 URI Ontology, CloudTrust, SCAP/CYBEX, GRC XML, ISACA's Cloud Computing Management Audit/Assurance Program, etc.)?	AWS obtains certain industry certifications and independent third-party attestations and provides certain certifications, reports and other relevant documentation directly to AWS Customers.
Audit Assurance & Compliance <i>Independent Audits</i>	AAC-02.1	Do you allow tenants to view your SOC2/ISO 27001 or similar third-party audit or certification reports?	AWS provides third-party attestations, certifications, Service Organization Controls (SOC) reports and other relevant compliance reports directly to our customers under NDA. The AWS ISO 27001 certification can be downloaded here: http://do.awsstatic.com/certifications/iso_27001_global_certification.pdf. The AWS SOC 3 report can be downloaded here: https://do.awsstatic.com/whitepapers/compliance/soc3_amazon_web_services.pdf. AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat
	AAC-02.2	Do you conduct network penetration tests of your cloud service infrastructure regularly as prescribed by industry best practices and guidance?	
	AAC-02.3	Do you conduct application penetration tests of your cloud infrastructure regularly as prescribed by industry best practices and guidance?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	AAC - 02.4	Do you conduct internal audits regularly as prescribed by industry best practices and guidance?	assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership.
	AAC - 02.5	Do you conduct external audits regularly as prescribed by industry best practices and guidance?	In addition, the AWS control environment is subject to regular internal and external audits and risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment.
	AAC - 02.6	Are the results of the penetration tests available to tenants at their request?	
	AAC - 02.7	Are the results of internal and external audits available to tenants at their request?	
	AAC - 02.8	Do you have an internal audit program that allows for cross-functional audit of assessments?	
Audit Assurance & Compliance <i>Information System Regulatory Mapping</i>	AAC -03.1	Do you have the ability to logically segment or encrypt customer data such that data may be produced for a single tenant only, without inadvertently accessing another tenant's data?	All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers retain control and ownership of their data, thus it is their responsibility to choose to encrypt the data. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security
	AAC - 03.2	Do you have capability to recover data for a specific customer in the case of a failure or data loss?	
	AAC - 03.3	Do you have the capability to restrict the storage of customer data to specific countries or geographic locations?	AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 and Glacier services are designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy. For information on data durability and redundancy, please refer to the AWS website. AWS Customers designate in which physical region their content will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US)(Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region, and South America (Sao Paulo).

Control Group	CID	Consensus Assessment Questions	AWS Response
	AAC - 03.4	Do you have a program in place that includes the ability to monitor changes to the regulatory requirements in relevant jurisdictions, adjust your security program for changes to legal requirements, and ensure compliance with relevant regulatory requirements?	AWS monitors relevant legal and regulatory requirements. Refer to ISO 27001 standard Annex 18 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Business Continuity Management & Operational Resilience <i>Business Continuity Planning</i>	BCR -01.1	Do you provide tenants with geographically resilient hosting options?	Data centers are built in clusters in various global regions. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Refer to AWS Overview of Cloud Security whitepaper for additional details - available at http://aws.amazon.com/security .
	BCR -01.2	Do you provide tenants with infrastructure service failover capability to other providers?	
Business Continuity Management & Operational Resilience <i>Business Continuity Testing</i>	BCR -02.1	Are business continuity plans subject to test at planned intervals or upon significant organizational or environmental changes to ensure continuing effectiveness?	AWS Business Continuity Policies and Plans have been developed and tested in alignment with ISO 27001 standards. Refer to ISO 27001 standard, annex A domain 17 for further details on AWS and business continuity.
Business Continuity Management & Operational Resilience <i>Power / Telecommunications</i>	BCR -03.1	Do you provide tenants with documentation showing the transport route of their data between your systems?	AWS Customers designate in which physical region their data and servers will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. AWS SOC reports provides additional details. Customers can also choose their network path to AWS facilities, including over dedicated, private networks where the customer controls the traffic routing.
	BCR - 03.2	Can tenants define how their data is transported and through which legal jurisdictions?	
Business Continuity Management & Operational Resilience Documentation	BCR -04.1	Are information system documents (e.g., administrator and user guides, architecture diagrams, etc.) made available to authorized personnel to ensure configuration, installation and operation of the information system?	Information System Documentation is made available internally to AWS personnel through the use of Amazon's Intranet site. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security/ . Refer to ISO 27001 Appendix A Domain 12.
Business Continuity Management & Operational Resilience <i>Environmental Risks</i>	BCR -05.1	Is physical protection against damage (e.g., natural causes, natural disasters, deliberate attacks) anticipated and designed with countermeasures applied?	AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 11.

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Equipment Location</i>	BCR -06.1	Are any of your data centers located in places that have a high probability/occurrence of high-impact environmental risks (floods, tornadoes, earthquakes, hurricanes, etc.)?	AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 11.
Business Continuity Management & Operational Resilience <i>Equipment Maintenance</i>	BCR -07.1	If using virtual infrastructure, does your cloud solution include independent hardware restore and recovery capabilities?	EBS Snapshot functionality allows customers to capture and restore virtual machine images at any time. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	BCR -07.2	If using virtual infrastructure, do you provide tenants with a capability to restore a Virtual Machine to a previous state in time?	
	BCR -07.3	If using virtual infrastructure, do you allow virtual machine images to be downloaded and ported to a new cloud provider?	
	BCR -07.4	If using virtual infrastructure, are machine images made available to the customer in a way that would allow the customer to replicate those images in their own off-site storage location?	
	BCR -07.5	Does your cloud solution include software/provider independent restore and recovery capabilities?	
Business Continuity Management & Operational Resilience <i>Equipment Power Failures</i>	BCR -08.1	Are security mechanisms and redundancies implemented to protect equipment from utility service outages (e.g., power failures, network disruptions, etc.)?	AWS equipment is protected from utility service outages in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS SOC reports provides additional details on controls in place to minimize the effect of a malfunction or physical disaster to the computer and data center facilities. In addition, refer to the AWS Cloud Security Whitepaper - available at http://aws.amazon.com/security.

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Impact Analysis</i>	BCR-09.1	Do you provide tenants with ongoing visibility and reporting of your operational Service Level Agreement (SLA) performance?	AWS CloudWatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com .
	BCR-09.2	Do you make standards-based information security metrics (CSA, CAMM, etc.) available to your tenants?	
	BCR-09.3	Do you provide customers with ongoing visibility and reporting of your SLA performance?	
Business Continuity Management & Operational Resilience <i>Policy</i>	BCR-10.1	Are policies and procedures established and made available for all personnel to adequately support services operations' roles?	Policies and Procedures have been established through AWS Security framework based upon NIST 800-53, ISO 27001, ISO 27017, ISO 27018, ISO 9001 standard and the PCI DSS requirements. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/compliance.
Business Continuity Management & Operational Resilience <i>Retention Policy</i>	BCR-11.1	Do you have technical control capabilities to enforce tenant data retention policies?	AWS provide customers with the ability to delete their data. However, AWS Customers retain control and ownership of their data so it is the customer's responsibility to manage data retention to their own requirements. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security. AWS errs on the side of protecting customer privacy and is vigilant in determining which law enforcement requests we must comply with. AWS does not hesitate to challenge orders from law enforcement if we think the orders lack a solid basis. For additional information refer to https://aws.amazon.com/compliance/data-privacy-faq/. AWS backup and redundancy mechanisms have been developed and tested in alignment with ISO 27001 standards. Refer to ISO 27001 standard, annex A domain 12 and the AWS SOC 2 report for additional information on AWS backup and redundancy mechanisms.
	BCR-11.2	Do you have a documented procedure for responding to requests for tenant data from governments or third parties?	
	BCR-11.4	Have you implemented backup or redundancy mechanisms to ensure compliance with regulatory, statutory, contractual or business requirements?	
	BCR-11.5	Do you test your backup or redundancy mechanisms at least annually?	
Change Control & Configuration Management <i>New Development / Acquisition</i>	CCC-01.1	Are policies and procedures established for management authorization for development or acquisition of new applications, systems, databases, infrastructure, services, operations and facilities?	Policies and Procedures have been established through AWS Security framework based upon NIST 800-53, ISO 27001, ISO 27017, ISO 27018, ISO 9001 standard and the PCI DSS requirements. Whether a customer is new to AWS or an advanced user, useful information about the services, ranging from introductions to advanced features, can be found on the AWS Documentation section of our website at https://aws.amazon.com/documentation/.
	CCC-01.2	Is documentation available that describes the installation, configuration and use of products/services/features?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Change Control & Configuration Management <i>Outsourced Development</i>	CCC-02.1	Do you have controls in place to ensure that standards of quality are being met for all software development?	AWS does not generally outsource development of software. AWS incorporates standards of quality as part of the system development lifecycle (SDLC) processes.
	CCC-02.2	Do you have controls in place to detect source code security defects for any outsourced software development activities?	Refer to ISO 27001 standard, Annex A, domain 14 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Quality Testing</i>	CCC-03.1	Do you provide your tenants with documentation that describes your quality assurance process?	AWS maintains an ISO 9001 certification. This is an independent validation of AWS quality system and determined that AWS activities comply with ISO 9001 requirements. AWS Security Bulletins notify customers of security and privacy events. Customers can subscribe to the AWS Security Bulletin RSS feed on our website. Refer to aws.amazon.com/security/security-bulletins/ .
	CCC-03.2	Is documentation describing known issues with certain products/services available?	AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com .
	CCC-03.3	Are there policies and procedures in place to triage and remedy reported bugs and security vulnerabilities for product and service offerings?	The AWS system development lifecycle (SDLC) incorporates industry best practices which include formal design reviews by the AWS Security Team, threat modeling and completion of a risk assessment. Refer to the AWS Overview of Security Processes for further details.
	CCC-03.4	Are mechanisms in place to ensure that all debugging and test code elements are removed from released software versions?	In addition, refer to ISO 27001 standard, Annex A, domain 14 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Unauthorized Software Installations</i>	CCC-04.1	Do you have controls in place to restrict and monitor the installation of unauthorized software onto your systems?	AWS' program, processes and procedures for managing malicious software is in alignment with ISO 27001 standards. Refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Production Changes</i>	CCC-05.1	Do you provide tenants with documentation that describes your production change management procedures and their roles / rights / responsibilities within it?	AWS SOC reports provides an overview of the controls in place to manage change management in the AWS environment. In addition, refer to ISO 27001 standard, Annex A, domain 12 for further details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Data Security & Information Lifecycle Management <i>Classification</i>	DSI-01.1	Do you provide a capability to identify virtual machines via policy tags/metadata (e.g., tags can be used to limit guest operating systems from booting/instantiating/transporting data in the wrong country)?	Virtual Machines are assigned to customers as a part of the EC2 service. Customers retain control over what resources are being used and where resources reside. Refer to the AWS website for additional details - http://aws.amazon.com .
	DSI-01.2	Do you provide a capability to identify hardware via policy tags/metadata/hardware tags (e.g., TXT/TPM, VN-Tag, etc.)?	AWS provides the ability to tag EC2 resources. A form of metadata, EC2 tags can be used to create user-friendly names, enhance searchability, and improve coordination between multiple users. The AWS Management Console has also supports tagging.
	DSI-01.3	Do you have a capability to use system geographic location as an authentication factor?	AWS provides the capability of conditional user access based on IP address. Customers can add conditions to control how users can use AWS, such as time of day, their originating IP address, or whether they are using SSL.
	DSI-01.4	Can you provide the physical location/geography of storage of a tenant's data upon request?	AWS provides customers the flexibility to place instances and store data within multiple geographic Regions. AWS Customers designate in which physical region their data and their servers will be located. AWS will not move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region and South America (Sao Paulo).
	DSI-01.5	Can you provide the physical location/geography of storage of a tenant's data in advance?	
	DSI-01.6	Do you follow a structured data-labeling standard (e.g., ISO 15489, Oasis XML Catalog Specification, CSA data type guidance)?	AWS Customers retain control and ownership of their data and may implement a structured data-labeling standard to meet their requirements.
	DSI-01.7	Do you allow tenants to define acceptable geographical locations for data routing or resource instantiation?	AWS provides customers the flexibility to place instances and store data within multiple geographic regions. AWS Customers designate in which physical region their data and their servers will be located. AWS will not move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region and South America (Sao Paulo).
Data Security & Information Lifecycle Management <i>Data Inventory / Flows</i>	DSI-02.1	Do you inventory, document, and maintain data flows for data that is resident (permanent or temporary) within the services' applications and infrastructure network and systems?	AWS Customers designate in which physical region their content will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US)(Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region, and South America (Sao Paulo).

Control Group	CID	Consensus Assessment Questions	AWS Response
	DSI-02.2	Can you ensure that data does not migrate beyond a defined geographical residency?	
Data Security & Information Lifecycle Management <i>eCommerce Transactions</i>	DSI-03.1	Do you provide open encryption methodologies (3.4ES, AES, etc.) to tenants in order for them to protect their data if it is required to move through public networks (e.g., the Internet)?	All of the AWS APIs are available via SSH-protected endpoints which provide server authentication. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Customers may also use third-party encryption technologies.
	DSI-03.2	Do you utilize open encryption methodologies any time your infrastructure components need to communicate with each other via public networks (e.g., Internet-based replication of data from one environment to another)?	Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Data Security & Information Lifecycle Management <i>Handling / Labeling / Security Policy</i>	DSI-04.1	Are policies and procedures established for labeling, handling and the security of data and objects that contain data?	AWS Customers retain control and ownership of their data and may implement a labeling and handling policy and procedures to meet their requirements.
	DSI-04.2	Are mechanisms for label inheritance implemented for objects that act as aggregate containers for data?	
Data Security & Information Lifecycle Management <i>Nonproduction Data</i>	DSI-05.1	Do you have procedures in place to ensure production data shall not be replicated or used in non-production environments?	AWS Customers retain control and ownership of their own data. AWS provides customers the ability to maintain and develop production and non-production environments. It is the responsibility of the customer to ensure that their production data is not replicated to non-production environments.
Data Security & Information Lifecycle Management <i>Ownership / Stewardship</i>	DSI-06.1	Are the responsibilities regarding data stewardship defined, assigned, documented and communicated?	AWS Customers retain control and ownership of their own data. Refer to the AWS Customer Agreement for additional information.
Data Security & Information Lifecycle Management <i>Secure Disposal</i>	DSI-07.1	Do you support secure deletion (e.g., degaussing/cryptographic wiping) of archived and backed-up data as determined by the tenant?	When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be

Control Group	CID	Consensus Assessment Questions	AWS Response
	DSI-07.2	Can you provide a published procedure for exiting the service arrangement, including assurance to sanitize all computing resources of tenant data once a customer has exited your environment or has vacated a resource?	decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security. Amazon EBS volumes are presented to you as raw unformatted block devices that have been wiped prior to being made available for use. Wiping occurs immediately before reuse so that you can be assured that the wipe process completed. If you have procedures requiring that all data be wiped via a specific method, such as those detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization"), you have the ability to do so on Amazon EBS. You should conduct a specialized wipe procedure prior to deleting the volume for compliance with your established requirements. Encryption of sensitive data is generally a good security practice, and AWS provides the ability to encrypt EBS volumes and their snapshots with AES-256. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. In order to be able to do this efficiently and with low latency, the EBS encryption feature is only available on EC2's more powerful instance types (e.g., M3, C3, R3, G2).
Datacenter Security <i>Asset Management</i>	DCS-01.1	Do you maintain a complete inventory of all of your critical assets that includes ownership of the asset?	In alignment with ISO 27001 standards, AWS Hardware assets are assigned an owner, tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 8 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	DCS-01.2	Do you maintain a complete inventory of all of your critical supplier relationships?	
Datacenter Security <i>Controlled Access Points</i>	DCS-02.1	Are physical security perimeters (e.g., fences, walls, barriers, guards, gates, electronic surveillance, physical authentication mechanisms, reception desks and security patrols) implemented?	Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. The AWS SOC reports provides additional details on the specific control activities executed by AWS. Refer to ISO 27001 standards; Annex A, domain 11 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Datacenter Security <i>Equipment Identification</i>	DCS-03.1	Is automated equipment identification used as a method to validate connection authentication integrity based on known equipment location?	AWS manages equipment identification in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Datacenter Security <i>Offsite Authorization</i>	DCS -04.1	Do you provide tenants with documentation that describes scenarios in which data may be moved from one physical location to another? (e.g., offsite backups, business continuity failovers, replication)	AWS Customers can designate which physical region their data will be located. AWS will not move customers' content from the selected Regions without notifying the customer unless required to comply with the law or requests of governmental entities. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Datacenter Security <i>Offsite equipment</i>	DCS -05.1	Can you provide tenants with evidence documenting your policies and procedures governing asset management and repurposing of equipment?	In alignment with ISO 27001 standards, when a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to ISO 27001 standards; Annex A, domain 8 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Datacenter Security <i>Policy</i>	DCS -06.1	Can you provide evidence that policies, standards and procedures have been established for maintaining a safe and secure working environment in offices, rooms, facilities and secure areas?	AWS engages with external certifying bodies and independent auditors to review and validate our compliance with compliance frameworks. AWS SOC reports provides additional details on the specific physical security control activities executed by AWS. Refer to ISO 27001 standards; Annex A, domain 11 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	DCS -06.2	Can you provide evidence that your personnel and involved third parties have been trained regarding your documented policies, standards and procedures?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security . AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification. In addition AWS SOC 1 and SOC 2 reports provides further information.
Datacenter Security <i>Secure Area Authorization</i>	DCS -07.1	Do you allow tenants to specify which of your geographic locations their data is allowed to move into/out of (to address legal jurisdictional considerations based on where data is stored vs. accessed)?	AWS Customers designate which physical region their data will be located. AWS will not move customers' content from the selected Regions without notifying the customer unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US)(Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region, and South America (Sao Paulo).

Control Group	CID	Consensus Assessment Questions	AWS Response
Datacenter Security <i>Unauthorized Persons Entry</i>	DCS -08.1	Are ingress and egress points, such as service areas and other points where unauthorized personnel may enter the premises, monitored, controlled and isolated from data storage and process?	Physical access is strictly controlled both at the perimeter and at building ingress points and includes, but is not limited to, professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. Physical access points to server locations are recorded by closed circuit television camera (CCTV) as defined in the AWS Data Center Physical Security Policy.
Datacenter Security <i>User Access</i>	DCS -09.1	Do you restrict physical access to information assets and functions by users and support personnel?	AWS Physical Security Mechanisms are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Encryption & Key Management <i>Entitlement</i>	EKM -01.1	Do you have key management policies binding keys to identifiable owners?	AWS provides customers the ability to use their own encryption mechanism for nearly all services including S3, EBS and EC2. VPC sessions are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications. AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Encryption & Key Management <i>Key Generation</i>	EKM -02.1	Do you have a capability to allow creation of unique encryption keys per tenant?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/).
	EKM -02.2	Do you have a capability to manage encryption keys on behalf of tenants?	Refer to AWS SOC reports for more details on KMS.
	EKM -02.3	Do you maintain key management procedures?	In addition, refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	EKM -02.4	Do you have documented ownership for each stage of the lifecycle of encryption keys?	Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications.
	EKM -02.5	Do you utilize any third party/open source/proprietary frameworks to manage encryption keys?	AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Encryption & Key	EKM -03.1	Do you encrypt tenant data at rest (on disk/storage) within your environment?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key

Control Group	CID	Consensus Assessment Questions	AWS Response
Management Encryption	EKM - 03.2	Do you leverage encryption to protect data and virtual machine images during transport across and between networks and hypervisor instances?	Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS SOC reports for more details on KMS. In addition, refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
	EKM - 03.3	Do you support tenant-generated encryption keys or permit tenants to encrypt data to an identity without access to a public key certificate (e.g. identity-based encryption)?	
	EKM - 03.4	Do you have documentation establishing and defining your encryption management policies, procedures and guidelines?	
Encryption & Key Management Storage and Access	EKM - 04.1	Do you have platform and data appropriate encryption that uses open/validated formats and standard algorithms?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS SOC reports for more details on KMS.
	EKM - 04.2	Are your encryption keys maintained by the cloud consumer or a trusted key management provider?	AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications.
	EKM - 04.3	Do you store encryption keys in the cloud?	
	EKM - 04.4	Do you have separate key management and key usage duties?	AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Governance and Risk Management Baseline Requirements	GR M- 01.1	Do you have documented information security baselines for every component of your infrastructure (e.g., hypervisors, operating systems, routers, DNS servers, etc.)?	In alignment with ISO 27001 standards, AWS maintains system baselines for critical components. Refer to ISO 27001 standards, Annex A, domain 14 and 18 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. Customers can provide their own virtual machine image. VM Import enables customers to easily import virtual machine images from your existing environment to Amazon EC2 instances.
	GR M- 01.2	Do you have a capability to continuously monitor and report the compliance of your infrastructure against your information security baselines?	
	GR M- 01.3	Do you allow your clients to provide their own trusted virtual machine image to ensure conformance to their own internal standards?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Governance and Risk Management <i>Risk Assessments</i>	GR M-02.1	Do you provide security control health data in order to allow tenants to implement industry standard Continuous Monitoring (which allows continual tenant validation of your physical and logical control status)?	AWS does publish independent auditor reports and certifications to provide customers with considerable information regarding the policies, processes, and controls established and operated by AWS. The relevant certifications and reports can be provided to AWS Customers. Continuous Monitoring of logical controls can be executed by customers on their own systems.
	GR M-02.2	Do you conduct risk assessments associated with data governance requirements at least once a year?	In alignment with ISO 27001 standard, AWS maintains a Risk Management program to mitigate and manage risk. In addition AWS maintains an AWS ISO 27018 certification. Alignment with ISO 27018 demonstrates to customers that AWS has a system of controls in place that specifically address the privacy protection of their content. For more information refer to the AWS Compliance ISO 27018 FAQ http://aws.amazon.com/compliance/iso-27018-faqs/ .
Governance and Risk Management <i>Management Oversight</i>	GR M-03.1	Are your technical, business, and executive managers responsible for maintaining awareness of and compliance with security policies, procedures, and standards for both themselves and their employees as they pertain to the manager and employees' area of responsibility?	The Control environment at Amazon begins at the highest level of the Company. Executive and senior leadership play important roles in establishing the Company's tone and core values. Every employee is provided with the Company's Code of Business Conduct and Ethics and completes periodic training. Compliance audits are performed so that employees understand and follow the established policies. Refer to AWS Risk & Compliance whitepaper for additional details - available at http://aws.amazon.com/compliance .
Governance and Risk Management <i>Management Program</i>	GR M-04.1	Do you provide tenants with documentation describing your Information Security Management Program (ISMP)?	AWS provides our customers with our ISO 27001 certification. The ISO 27001 certification is specifically focused on the AWS ISMS and measures how AWS internal processes follow the ISO standard. Certification means a third party accredited independent auditor has performed an assessment of our processes and controls and confirms they are operating in alignment with the ISO 27001 certification standard. For additional information refer to the AWS Compliance ISO 27001 FAQ website: http://aws.amazon.com/compliance/iso-27001-faqs/ .
	GR M-04.2	Do you review your Information Security Management Program (ISMP) least once a year?	
Governance and Risk Management <i>Management Support / Involvement</i>	GR M-05.1	Do you ensure your providers adhere to your information security and privacy policies?	AWS has established information security framework and policies which have integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, PCI DSS v3.1 and National Institute of Standards and Technology (NIST) Publication 800-53 (Recommended Security Controls for Federal Information Systems).
Governance and Risk Management <i>Policy</i>	GR M-06.1	Do your information security and privacy policies align with industry standards (ISO-27001, ISO-22307, CoBIT, etc.)?	AWS manages third-party relationships in alignment with ISO 27001 standards. AWS Third Party requirements are reviewed by independent external

Control Group	CID	Consensus Assessment Questions	AWS Response
	GR M-06.2	Do you have agreements to ensure your providers adhere to your information security and privacy policies?	auditors during audits for our PCI DSS, ISO 27001 and FedRAMP compliance. Information about the AWS Compliance programs is published publicly on our website at http://aws.amazon.com/compliance/.
	GR M-06.3	Can you provide evidence of due diligence mapping of your controls, architecture and processes to regulations and/or standards?	
	GR M-06.4	Do you disclose which controls, standards, certifications and/or regulations you comply with?	
Governance and Risk Management <i>Policy Enforcement</i>	GR M-07.1	Is a formal disciplinary or sanction policy established for employees who have violated security policies and procedures?	AWS provides security policies and security training to employees to educate them as to their role and responsibilities concerning information security. Employees who violate Amazon standards or protocols are investigated and appropriate disciplinary action (e.g. warning, performance plan, suspension, and/or termination) is followed.
	GR M-07.2	Are employees made aware of what actions could be taken in the event of a violation via their policies and procedures?	Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security . Refer to ISO 27001 Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Governance and Risk Management <i>Business / Policy Change Impacts</i>	GR M-08.1	Do risk assessment results include updates to security policies, procedures, standards and controls to ensure they remain relevant and effective?	Updates to AWS security policies, procedures, standards and controls occur on an annual basis in alignment with the ISO 27001 standard. Refer to ISO 27001 for additional information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification.
Governance and Risk Management <i>Policy Reviews</i>	GR M-09.1	Do you notify your tenants when you make material changes to your information security and/or privacy policies?	Our AWS Cloud Security Whitepaper and Risk and Compliance whitepapers, available at http://aws.amazon.com/security and http://aws.amazon.com/compliance , are updated on a regular basis to reflect updates to the AWS policies.
	GR M-09.2	Do you perform, at minimum, annual reviews to your privacy and security policies?	The AWS SOC reports provide details related to privacy and security policy review.
Governance and Risk Management <i>Assessments</i>	GR M-10.1	Are formal risk assessments aligned with the enterprise-wide framework and performed at least annually, or at planned intervals, determining the likelihood and impact of all identified risks, using qualitative and quantitative methods?	In alignment with ISO 27001 AWS has developed a Risk Management program to mitigate and manage risk. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification. Refer to AWS Risk and Compliance Whitepaper (available at aws.amazon.com/security) for additional details on AWS Risk Management Framework.

Control Group	CID	Consensus Assessment Questions	AWS Response
	GR M-10.2	Is the likelihood and impact associated with inherent and residual risk determined independently, considering all risk categories (e.g., audit results, threat and vulnerability analysis, and regulatory compliance)?	
Governance and Risk Management Program	GR M-11.1	Do you have a documented, organization-wide program in place to manage risk?	In alignment with ISO 27001, AWS maintains a Risk Management program to mitigate and manage risk.
	GR M-11.2	Do you make available documentation of your organization-wide risk management program?	AWS management has a strategic business plan which includes risk identification and the implementation of controls to mitigate or manage risks. AWS management re-evaluates the strategic business plan at least biannually. This process requires management to identify risks within its areas of responsibility and to implement appropriate measures designed to address those risks. AWS Risk Management program is reviewed by independent external auditors during audits for our PCI DSS, ISO 27001 and FedRAMP compliance.
Human Resources Asset Returns	HRS -01.1	Are systems in place to monitor for privacy breaches and notify tenants expeditiously if a privacy event may have impacted their data?	AWS Customers retain the responsibility to monitor their own environment for privacy breaches. The AWS SOC reports provides an overview of the controls in place to monitor AWS managed environment.
	HRS -01.2	Is your Privacy Policy aligned with industry standards?	
Human Resources Background Screening	HRS -02.1	Pursuant to local laws, regulations, ethics and contractual constraints, are all employment candidates, contractors and involved third parties subject to background verification?	AWS conducts criminal background checks, as permitted by applicable law, as part of pre-employment screening practices for employees commensurate with the employee's position and level of access to AWS facilities. The AWS SOC reports provides additional details regarding the controls in place for background verification.
Human Resources Employment Agreements	HRS -03.1	Do you specifically train your employees regarding their specific role and the information security controls they must fulfill?	In alignment with ISO 27001 standard, all AWS employees complete periodic role based training that includes AWS Security training and requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. Refer to SOC reports for additional details.
	HRS -03.2	Do you document employee acknowledgment of training they have completed?	All personnel supporting AWS systems and devices must sign a non-disclosure agreement prior to being granted access. Additionally, upon hire, personnel are required to read and accept the Acceptable Use Policy and the Amazon Code of Business Conduct and Ethics (Code of Conduct) Policy.
	HRS -03.3	Are all personnel required to sign NDA or Confidentiality Agreements as a condition of employment to protect customer/tenant information?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	HRS - 03.4	Is successful and timed completion of the training program considered a prerequisite for acquiring and maintaining access to sensitive systems?	
	HRS - 03.5	Are personnel trained and provided with awareness programs at least once a year?	
Human Resources <i>Employment Termination</i>	HRS -04.1	Are documented policies, procedures and guidelines in place to govern change in employment and/or termination?	AWS Human Resources team defines internal management responsibilities to be followed for termination and role change of employees and vendors. AWS SOC reports provide additional details.
	HRS - 04.2	Do the above procedures and guidelines account for timely revocation of access and return of assets?	Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked. AWS SOC reports provide further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Human Resources <i>Portable / Mobile Devices</i>	HRS -05.1	Are policies and procedures established and measures implemented to strictly limit access to your sensitive data and tenant data from portable and mobile devices (e.g. laptops, cell phones and personal digital assistants (PDAs)), which are generally higher-risk than non-portable devices (e.g., desktop computers at the provider organization's facilities)?	Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
Human Resources <i>Nondisclosure Agreements</i>	HRS -06.1	Are requirements for non-disclosure or confidentiality agreements reflecting the organization's needs for the protection of data and operational details identified, documented and reviewed at planned intervals?	Amazon Legal Counsel manages and periodically revises the Amazon NDA to reflect AWS business needs.
Human Resources <i>Roles / Responsibilities</i>	HRS -07.1	Do you provide tenants with a role definition document clarifying your administrative responsibilities versus those of the tenant?	The AWS Cloud Security Whitepaper and the AWS Risk and Compliance Whitepaper provide details on the roles and responsibilities of AWS and those of our Customers. The whitepapers are available at: http://aws.amazon.com/security and http://aws.amazon.com/compliance .

Control Group	CID	Consensus Assessment Questions	AWS Response
Human Resources <i>Acceptable Use</i>	HRS -08.1	Do you provide documentation regarding how you may or access tenant data and metadata?	AWS has a formal access control policy that is reviewed and updated on an annual basis (or when any major change to the system occurs that impacts the policy). The policy addresses purpose, scope, roles, responsibilities and management commitment. AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content. Refer to the ISO 27001 standard and 27018 code of practice for additional information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 and ISO 27018.
	HRS -08.2	Do you collect or create metadata about tenant data usage through inspection technologies (search engines, etc.)?	
	HRS -08.3	Do you allow tenants to opt out of having their data/metadata accessed via inspection technologies?	
Human Resources <i>Training / Awareness</i>	HRS -09.1	Do you provide a formal, role-based, security awareness training program for cloud-related access and data management issues (e.g., multi-tenancy, nationality, cloud delivery model segregation of duties implications and conflicts of interest) for all persons with access to tenant data?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. AWS roles and responsibilities are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
	HRS -09.2	Are administrators and data stewards properly educated on their legal responsibilities with regard to security and data integrity?	
Human Resources <i>User Responsibility</i>	HRS -10.1	Are users made aware of their responsibilities for maintaining awareness and compliance with published security policies, procedures, standards and applicable regulatory requirements?	AWS has implemented various methods of internal communication at a global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employee as well as electronic mail messages and the posting of information via the Amazon intranet. Refer to ISO 27001 standard, Annex A, domain 7 and 8. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. In addition the AWS Cloud Security Whitepaper provides further details - available at http://aws.amazon.com/security.
	HRS -10.2	Are users made aware of their responsibilities for maintaining a safe and secure working environment?	
	HRS -10.3	Are users made aware of their responsibilities for leaving unattended equipment in a secure manner?	
Human Resources <i>Workspace</i>	HRS -11.1	Do your data management policies and procedures address tenant and service level conflicts of interests?	AWS data management policies are in alignment with ISO 27001 standard. Refer to ISO 27001 standard, Annex A, domain 8 and 9. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS SOC reports

Control Group	CID	Consensus Assessment Questions	AWS Response
	HRS-11.2	Do your data management policies and procedures include a tamper audit or software integrity function for unauthorized access to tenant data?	provides additional details on the specific control activities executed by AWS to prevent unauthorized access to AWS resources. AWS has identified auditable event categories across systems and devices within the AWS system. Service teams configure the auditing features to record continuously the security-related events in accordance with requirements. Audit records contain a set of data elements in order to support necessary analysis requirements. In addition, audit records are available for AWS Security team or other appropriate teams to perform inspection or analysis on demand, and in response to security-related or business-impacting events.
	HRS-11.3	Does the virtual machine management infrastructure include a tamper audit or software integrity function to detect changes to the build/configuration of the virtual machine?	
Identity & Access Management <i>Audit Tools Access</i>	IAM-01.1	Do you restrict, log and monitor access to your information security management systems? (E.g., hypervisors, firewalls, vulnerability scanners, network sniffers, APIs, etc.)	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
	IAM-01.2	Do you monitor and log privileged access (administrator level) to information security management systems?	AWS has identified auditable event categories across systems and devices within the AWS system. Service teams configure the auditing features to record continuously the security-related events in accordance with requirements. The log storage system is designed to provide a highly scalable, highly available service that automatically increases capacity as the ensuing need for log storage grows. Audit records contain a set of data elements in order to support necessary analysis requirements. In addition, audit records are available for AWS Security team or other appropriate teams to perform inspection or analysis on demand, and in response to security-related or business-impacting events. Designated personnel on AWS teams receive automated alerts in the event of an audit processing failure. Audit processing failures include, for example, software/hardware errors. When alerted, on-call personnel issue a trouble ticket and track the event until it is resolved. AWS logging and monitoring processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Identity & Access Management <i>User Access Policy</i>	IAM-02.1	Do you have controls in place ensuring timely removal of systems access that is no longer required for business purposes?	AWS SOC reports provides further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 9 for additional details. AWS has

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM - 02.2	Do you provide metrics to track the speed with which you are able to remove systems access that is no longer required for business purposes?	been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Identity & Access Management <i>Diagnostic / Configuration Ports Access</i>	IAM -03.1	Do you use dedicated secure networks to provide management access to your cloud service infrastructure?	Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored per the AWS access policy. In addition, customer data is and server instances are logically isolated from other customers by default. Privileged user access controls are reviewed by an independent auditor during the AWS SOC, ISO 27001, PCI, ITAR, and FedRAMP audits.
Identity & Access Management <i>Policies and Procedures</i>	IAM -04.1	Do you manage and store the identity of all personnel who have access to the IT infrastructure, including their level of access?	
	IAM - 04.2	Do you manage and store the user identity of all personnel who have network access, including their level of access?	
Identity & Access Management <i>Segregation of Duties</i>	IAM -05.1	Do you provide tenants with documentation on how you maintain segregation of duties within your cloud service offering?	Customers retain the ability to manage segregations of duties of their AWS resources. Internally, AWS aligns with ISO 27001 standards for managing segregation of duties. Refer to ISO 27001 standard, Annex A, domain 6 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Identity & Access Management <i>Source Code Access Restriction</i>	IAM -06.1	Are controls in place to prevent unauthorized access to your application, program or object source code, and assure it is restricted to authorized personnel only?	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IAM - 06.2	Are controls in place to prevent unauthorized access to tenant application, program or object source code, and assure it is restricted to authorized personnel only?	
Identity & Access Management <i>Third Party Access</i>	IAM -07.1	Do you provide multi-failure disaster recovery capability?	AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS SOC reports provides further details. ISO 27001 standard Annex A, domain 15 provides additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification.
	IAM - 07.2	Do you monitor service continuity with upstream providers in the event of provider failure?	
	IAM - 07.3	Do you have more than one provider for each service you depend on?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM - 07.4	Do you provide access to operational redundancy and continuity summaries, including the services you depend on?	
	IAM -07.5	Do you provide the tenant the ability to declare a disaster?	
	IAM - 07.6	Do you provided a tenant-triggered failover option?	
	IAM -07.7	Do you share your business continuity and redundancy plans with your tenants?	
Identity & Access Management <i>User Access Restriction / Authorization</i>	IAM -08.1	Do you document how you grant and approve access to tenant data?	AWS Customers retain control and ownership of their data. Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored. In addition, customer data and server instances are logically isolated from other customers by default. Privileged user access controls are reviewed by an independent auditor during the AWS SOC, ISO 27001, PCI, ITAR, and FedRAMP audits.
	IAM - 08.2	Do you have a method of aligning provider and tenant data classification methodologies for access control purposes?	
Identity & Access Management <i>User Access Authorization</i>	IAM -09.1	Does your management provision the authorization and restrictions for user access (e.g. employees, contractors, customers (tenants), business partners and/or suppliers) prior to their access to data and any owned or managed (physical and virtual) applications, infrastructure systems and network components?	Unique user identifiers are created as part of the onboarding workflow process in the AWS human resources management system. The device provisioning process helps ensure unique identifiers for devices. Both processes include manager approval to establish the user account or device. Initial authenticators are delivered to user's in-person and to devices as part of the provisioning process. Internal users can associate SSH public keys with their account. System account authenticators are provided to the requestor as part of the account creation process after the identity of the requestor is verified.
	IAM - 09.2	Do you provide upon request user access (e.g. employees, contractors, customers (tenants), business partners and/or suppliers) to data and any owned or managed (physical and virtual) applications, infrastructure systems and network components?	AWS has established controls to address the threat of inappropriate insider access. All certifications and third-party attestations evaluate logical access preventative and detective controls. In addition, periodic risk assessments focus on how insider access is controlled and monitored.
Identity & Access Management <i>User Access Reviews</i>	IAM -10.1	Do you require at least annual certification of entitlements for all system users and administrators (exclusive of users maintained by your tenants)?	In alignment with ISO 27001 standard, all access grants are reviewed on a periodic basis; explicit re-approval is required or access to the resource is automatically revoked. Controls specific to User Access reviews are outlined in the SOC reports. Exceptions in the User entitlement controls are documented in the SOC reports. Refer to ISO 27001 standards, Annex A, domain 9 for additional details.

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM-10.2	If users are found to have inappropriate entitlements, are all remediation and certification actions recorded?	AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IAM-10.3	Will you share user entitlement remediation and certification reports with your tenants, if inappropriate access may have been allowed to tenant data?	
Identity & Access Management <i>User Access Revocation</i>	IAM-11.1	Is timely deprovisioning, revocation or modification of user access to the organizations systems, information assets and data implemented upon any change in status of employees, contractors, customers, business partners or involved third parties?	Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked. AWS SOC reports provides further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 9 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IAM-11.2	Is any change in user access status intended to include termination of employment, contract or agreement, change of employment or transfer within the organization?	
Identity & Access Management <i>User ID Credentials</i>	IAM-12.1	Do you support use of, or integration with, existing customer-based Single Sign On (SSO) solutions to your service?	The AWS Identity and Access Management (IAM) service provides identity federation to the AWS Management Console. Multi-factor authentication is an optional feature that a customer can utilize. Refer to the AWS website for additional details - http://aws.amazon.com/mfa . AWS Identity and Access Management (IAM) supports identity federation for delegated access to the AWS Management Console or AWS APIs. With identity federation, external identities (federated users) are granted secure access to resources in your AWS account without having to create IAM users. These external identities can come from your corporate identity provider (such as Microsoft Active Directory or from the AWS Directory Service) or from a web identity provider, such as Amazon Cognito, Login with Amazon, Facebook, Google or any OpenID Connect (OIDC) compatible provider.
	IAM-12.2	Do you use open standards to delegate authentication capabilities to your tenants?	
	IAM-12.3	Do you support identity federation standards (SAML, SPML, WS-Federation, etc.) as a means of authenticating/authorizing users?	
	IAM-12.4	Do you have a Policy Enforcement Point capability (e.g., XACML) to enforce regional legal and policy constraints on user access?	
	IAM-12.5	Do you have an identity management system (enabling classification of data for a tenant) in place to enable both role-based and context-based entitlement to data?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM -12.6	Do you provide tenants with strong (multifactor) authentication options (digital certs, tokens, biometrics, etc.) for user access?	AWS Identity and Access Management (IAM) enables customers to securely control access to AWS services and resources for their users. Additional information about IAM can be found on website at https://aws.amazon.com/iam/ . AWS SOC reports provides details on the specific control activities executed by AWS.
	IAM -12.7	Do you allow tenants to use third-party identity assurance services?	
	IAM -12.8	Do you support password (minimum length, age, history, complexity) and account lockout (lockout threshold, lockout duration) policy enforcement?	
	IAM -12.9	Do you allow tenants/customers to define password and account lockout policies for their accounts?	
	IAM -12.10	Do you support the ability to force password changes upon first logon?	
	IAM -12.11	Do you have mechanisms in place for unlocking accounts that have been locked out (e.g., self-service via email, defined challenge questions, manual unlock)?	
Identity & Access Management <i>Utility Programs Access</i>	IAM -13.1	Are utilities that can significantly manage virtualized partitions (e.g., shutdown, clone, etc.) appropriately restricted and monitored?	In alignment with ISO 27001 standards, system utilities are appropriately restricted and monitored. AWS SOC reports provides details on the specific control activities executed by AWS. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IAM -13.2	Do you have a capability to detect attacks that target the virtual infrastructure directly (e.g., shimming, Blue Pill, Hyper jumping, etc.)?	
	IAM -13.3	Are attacks that target the virtual infrastructure prevented with technical controls?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>Audit Logging / Intrusion Detection</i>	IVS-01.1	Are file integrity (host) and network intrusion detection (IDS) tools implemented to help facilitate timely detection, investigation by root cause analysis and response to incidents?	AWS Incident response program (detection, investigation and response to incidents) has been developed in alignment with ISO 27001 standards, system utilities are appropriately restricted and monitored. AWS SOC reports provides additional details on controls in place to restrict system access. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IVS-01.2	Is physical and logical user access to audit logs restricted to authorized personnel?	
	IVS-01.3	Can you provide evidence that due diligence mapping of regulations and standards to your controls/architecture/processes has been done?	
	IVS-01.4	Are audit logs centrally stored and retained?	In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-01.5	Are audit logs reviewed on a regular basis for security events (e.g., with automated tools)?	AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Infrastructure & Virtualization Security <i>Change Detection</i>	IVS-02.1	Do you log and alert any changes made to virtual machine images regardless of their running state (e.g. dormant, off or running)?	Virtual Machines are assigned to customers as a part of the EC2 service. Customers retain control over what resources are being used and where resources reside. Refer to the AWS website for additional details - http://aws.amazon.com .
	IVS-02.2	Are changes made to virtual machines, or moving of an image and subsequent validation of the image's integrity, made immediately available to customers through electronic methods (e.g. portals or alerts)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>Clock Synchronization</i>	IVS-03.1	Do you use a synchronized time-service protocol (e.g., NTP) to ensure all systems have a common time reference?	In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-04.1	Do you provide documentation regarding what levels of system (network, storage, memory, I/O, etc.) oversubscription you maintain and under what circumstances/scenarios?	Details regarding AWS Service Limits and how to request an increase for specific services is available on the AWS website at http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html . AWS manages capacity and utilization data in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-04.2	Do you restrict use of the memory oversubscription capabilities present in the hypervisor?	
	IVS-04.3	Do your system capacity requirements take into account current, projected and anticipated capacity needs for all systems used to provide services to the tenants?	
Infrastructure & Virtualization Security <i>Capacity / Resource Planning</i>	IVS-04.4	Is system performance monitored and tuned in order to continuously meet regulatory, contractual and business requirements for all the systems used to provide services to the tenants?	
Infrastructure & Virtualization Security <i>Management - Vulnerability Management</i>	IVS-05.1	Do security vulnerability assessment tools or services accommodate the virtualization technologies being used (e.g. virtualization aware)?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. Regular internal and external vulnerability scans are performed on the host operating system, web application and databases in the AWS environment utilizing a variety of tools. Vulnerability scanning and remediation practices are regularly reviewed as a part of AWS continued compliance with PCI DSS and FedRAMP.
	IVS-06.1	For your IaaS offering, do you provide customers with guidance on how to create a layered security architecture equivalence using your virtualized solution?	AWS website provides guidance on creating a layered security architecture in a number of white papers available via the AWS public website - http://aws.amazon.com/documentation/ .
	IVS-06.2	Do you regularly update network architecture diagrams that include data flows between security domains/zones?	Boundary protection devices that employ rule sets, access control lists (ACL), and configurations enforce the flow of information between network fabrics. Several network fabrics exist at Amazon, each separated by devices that

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-06.3	Do you regularly review for appropriateness the allowed access/connectivity (e.g., firewall rules) between security domains/zones within the network?	control the flow of information between fabrics. The flow of information between fabrics is established by approved authorizations, which exist as access control lists (ACL) which reside on these devices. These devices control the flow of information between fabrics as mandated by these ACLs. ACLs are defined, approved by appropriate personnel, managed and deployed using AWS ACL-manage tool.
	IVS-06.4	Are all firewall access control lists documented with business justification?	Amazon's Information Security team approves these ACLs. Approved firewall rule sets and access control lists between network fabrics restrict the flow of information to specific information system services. Access control lists and rule sets are reviewed and approved, and are automatically pushed to boundary protection devices on a periodic basis (at least every 24 hours) to ensure rule-sets and access control lists are up-to-date.
Infrastructure & Virtualization Security <i>OS Hardening and Base Controls</i>	IVS-07.1	Are operating systems hardened to provide only the necessary ports, protocols and services to meet business needs using technical controls (i.e antivirus, file integrity monitoring and logging) as part of their baseline build standard or template?	AWS Network Management is regularly reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm. AWS implements least privilege throughout its infrastructure components. AWS prohibits all ports and protocols that do not have a specific business purpose. AWS follows a rigorous approach to minimal implementation of only those features and functions that are essential to use of the device. Network scanning is performed and any unnecessary ports or protocols in use are corrected. Regular internal and external vulnerability scans are performed on the host operating system, web application and databases in the AWS environment utilizing a variety of tools. Vulnerability scanning and remediation practices are regularly reviewed as a part of AWS continued compliance with PCI DSS and FedRAMP.
	IVS-08.1	For your SaaS or PaaS offering, do you provide tenants with separate environments for production and test processes?	AWS Customers retain the ability and the responsibility to create and maintain production and test environments. AWS website provides guidance on creating an environment utilizing the AWS services - http://aws.amazon.com/documentation/ .
Infrastructure & Virtualization Security <i>Production / Nonproduction Environments</i>	IVS-08.2	For your IaaS offering, do you provide tenants with guidance on how to create suitable production and test environments?	
	IVS-08.3	Do you logically and physically segregate production and non-production environments?	AWS Customers retain responsibility to manage their own network segmentation in adherence with their defined requirements. Internally, AWS network segmentation is aligned with ISO 27001 standards. Refer to ISO 27001 standard, Annex A. domain 13 for further detail. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Infrastructure & Virtualization Security <i>Segmentation</i>	IVS-09.1	Are system and network environments protected by a firewall or virtual firewall to ensure business and customer security requirements?	
	IVS-09.2	Are system and network environments protected by a firewall or virtual firewall to ensure compliance with legislative, regulatory and contractual requirements?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-09.3	Are system and network environments protected by a firewall or virtual firewall to ensure separation of production and non-production environments?	
	IVS-09.4	Are system and network environments protected by a firewall or virtual firewall to ensure protection and isolation of sensitive data?	
Infrastructure & Virtualization Security <i>VM Security - vMotion Data Protection</i>	IVS-10.1	Are secured and encrypted communication channels used when migrating physical servers, applications or data to virtual servers?	AWS provides customers the ability to use their own encryption mechanism for nearly all services including S3, EBS and EC2. VPC sessions are also encrypted.
	IVS-10.2	Do you use a network segregated from production-level networks when migrating physical servers, applications or data to virtual servers?	AWS Customers retain control and ownership of their own data. AWS provides customers the ability to maintain and develop production and non-production environments. It is the responsibility of the customer to ensure that their production data is not replicated to non-production environments.
Infrastructure & Virtualization Security <i>VMM Security - Hypervisor Hardening</i>	IVS-11.1	Do you restrict personnel access to all hypervisor management functions or administrative consoles for systems hosting virtualized systems based on the principle of least privilege and supported through technical controls (e.g. two-factor authentication, audit trails, IP address filtering, firewalls and TLS-encapsulated communications to the administrative consoles)?	AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. When user accounts are created, user accounts are created to have minimal access. Access above these least privileges requires appropriate authorization. Refer to AWS SOC reports for more information on Access Controls.
Infrastructure & Virtualization Security <i>Wireless Security</i>	IVS-12.1	Are policies and procedures established and mechanisms configured and implemented to protect the wireless network environment perimeter and to restrict unauthorized wireless traffic?	Policies, procedures and mechanisms to protect AWS network environment are in place. AWS security controls are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
	IVS-12.2	Are policies and procedures established and mechanisms implemented to ensure wireless security settings are enabled with strong encryption for authentication and transmission, replacing vendor default settings? (e.g., encryption keys, passwords, SNMP community strings)	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-12.3	Are policies and procedures established and mechanisms implemented to protect wireless network environments and detect the presence of unauthorized (rogue) network devices for a timely disconnect from the network?	
Infrastructure & Virtualization Security <i>Network Architecture</i>	IVS-13.1	Do your network architecture diagrams clearly identify high-risk environments and data flows that may have legal compliance impacts?	AWS Customers retain responsibility to manage their own network segmentation in adherence with their defined requirements. Internally, AWS network segmentation is aligned with the ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-13.2	Do you implement technical measures and apply defense-in-depth techniques (e.g., deep packet analysis, traffic throttling and black-holing) for detection and timely response to network-based attacks associated with anomalous ingress or egress traffic patterns (e.g., MAC spoofing and ARP poisoning attacks) and/or distributed denial-of-service (DDoS) attacks?	AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment. AWS security controls are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Interoperability & Portability APIs	IPY-01	Do you publish a list of all APIs available in the service and indicate which are standard and which are customized?	Details regarding AWS APIs can be found on the AWS website at https://aws.amazon.com/documentation/. In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources.
Interoperability & Portability <i>Data Request</i>	IPY-02	Is unstructured customer data available on request in an industry-standard format (e.g., .doc, .xls, or .pdf)?	Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
Interoperability & Portability <i>Policy & Legal</i>	IPY-03.1	Do you provide policies and procedures (i.e. service level agreements) governing the use of APIs for interoperability between your service and third-party applications?	
	IPY-03.2	Do you provide policies and procedures (i.e. service level agreements) governing the migration of application data to and from your service?	Customer retain control and ownership of their content. Customers can choose how they migrate applications and content both on and off the AWS platform at their discretion.

Control Group	CID	Consensus Assessment Questions	AWS Response
Interoperability & Portability <i>Standardized Network Protocols</i>	IPY-04.1	Can data import, data export and service management be conducted over secure (e.g., non-clear text and authenticated), industry accepted standardized network protocols?	AWS allows customers to move data as needed on and off AWS storage. Refer to http://aws.amazon.com/choosing-a-cloud-platform for more information on Storage options.
	IPY-04.2	Do you provide consumers (tenants) with documentation detailing the relevant interoperability and portability network protocol standards that are involved?	
Interoperability & Portability <i>Virtualization</i>	IPY-05.1	Do you use an industry-recognized virtualization platform and standard virtualization formats (e.g., OVF) to help ensure interoperability?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	IPY-05.2	Do you have documented custom changes made to any hypervisor in use, and all solution-specific virtualization hooks available for customer review?	
Mobile Security <i>Anti-Malware</i>	MOS-01	Do you provide anti-malware training specific to mobile devices as part of your information security awareness training?	AWS' program, processes and procedures to managing antivirus / malicious software is in alignment with ISO 27001 standards. Refer to ISO 27001 standard, Annex A, domain 12 for additional information.
Mobile Security <i>Application Stores</i>	MOS-02	Do you document and make available lists of approved application stores for mobile devices accessing or storing company data and/or company systems?	AWS has established an information security framework and policies and has effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.1 and the National Institute of Standards and Technology (NIST) Publication 800-53 (Recommended Security Controls for Federal Information Systems).
Mobile Security <i>Approved Applications</i>	MOS-03	Do you have a policy enforcement capability (e.g., XACML) to ensure that only approved applications and those from approved application stores be loaded onto a mobile device?	
Mobile Security <i>Approved Software for BYOD</i>	MOS-04	Does your BYOD policy and training clearly state which applications and applications stores are approved for use on BYOD devices?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Mobile Security <i>Awareness and Training</i>	MOS-05	Do you have a documented mobile device policy in your employee training that clearly defines mobile devices and the accepted usage and requirements for mobile devices?	
Mobile Security <i>Cloud Based Services</i>	MOS-06	Do you have a documented list of pre-approved cloud based services that are allowed to be used for use and storage of company business data via a mobile device?	
Mobile Security <i>Compatibility</i>	MOS-07	Do you have a documented application validation process for testing device, operating system and application compatibility issues?	
Mobile Security <i>Device Eligibility</i>	MOS-08	Do you have a BYOD policy that defines the device(s) and eligibility requirements allowed for BYOD usage?	
Mobile Security <i>Device Inventory</i>	MOS-09	Do you maintain an inventory of all mobile devices storing and accessing company data which includes device status (os system and patch levels, lost or decommissioned, device assignee)?	
Mobile Security <i>Device Management</i>	MOS-10	Do you have a centralized mobile device management solution deployed to all mobile devices that are permitted to store, transmit, or process company data?	
Mobile Security <i>Encryption</i>	MOS-11	Does your mobile device policy require the use of encryption for either the entire device or for data identified as sensitive enforceable through technology controls for all mobile devices?	
Mobile Security <i>Jailbreaking and Rooting</i>	MOS-12.1	Does your mobile device policy prohibit the circumvention of built-in security controls on mobile devices (e.g., jailbreaking or rooting)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	MOS -12.2	Do you have detective and preventative controls on the device or via a centralized device management system which prohibit the circumvention of built-in security controls?	
Mobile Security <i>Legal</i>	MOS -13.1	Does your BYOD policy clearly define the expectation of privacy, requirements for litigation, e-discovery and legal holds?	Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
	MOS -13.2	Do you have detective and preventative controls on the device or via a centralized device management system which prohibit the circumvention of built-in security controls?	
Mobile Security <i>Lockout Screen</i>	MOS -14	Do you require and enforce via technical controls an automatic lockout screen for BYOD and company owned devices?	
Mobile Security <i>Operating Systems</i>	MOS -15	Do you manage all changes to mobile device operating systems, patch levels and applications via your company's change management processes?	
Mobile Security <i>Passwords</i>	MOS -16.1	Do you have password policies for enterprise issued mobile devices and/or BYOD mobile devices?	
	MOS -16.2	Are your password policies enforced through technical controls (i.e. MDM)?	
	MOS -16.3	Do your password policies prohibit the changing of authentication requirements (i.e. password/PIN length) via a mobile device?	
Mobile Security <i>Policy</i>	MOS -17.1	Do you have a policy that requires BYOD users to perform backups of specified corporate data?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	MOS-17.2	Do you have a policy that requires BYOD users to prohibit the usage of unapproved application stores?	
	MOS-17.3	Do you have a policy that requires BYOD users to use anti-malware software (where supported)?	
Mobile Security <i>Remote Wipe</i>	MOS-18.1	Does your IT provide remote wipe or corporate data wipe for all company-accepted BYOD devices?	
	MOS-18.2	Does your IT provide remote wipe or corporate data wipe for all company-assigned mobile devices?	
Mobile Security <i>Security Patches</i>	MOS-19.1	Do your mobile devices have the latest available security-related patches installed upon general release by the device manufacturer or carrier?	
	MOS-19.2	Do your mobile devices allow for remote validation to download the latest security patches by company IT personnel?	
Mobile Security <i>Users</i>	MOS-20.1	Does your BYOD policy clarify the systems and servers allowed for use or access on the BYOD-enabled device?	AWS maintains contacts with industry bodies, risk and compliance organizations, local authorities and regulatory bodies as required by the ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. The AWS SOC reports provides details on the specific control activities executed by AWS. All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities.
	MOS-20.2	Does your BYOD policy specify the user roles that are allowed access via a BYOD-enabled device?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Contact / Authority Maintenance</i>	SEF-01.1	Do you maintain liaisons and points of contact with local authorities in accordance with contracts and appropriate regulations?	
	SEF-02.1	Do you have a documented security incident response plan?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Management</i>	SEF-02.2	Do you integrate customized tenant requirements into your security incident response plans?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	SEF-02.3	Do you publish a roles and responsibilities document specifying what you vs. your tenants are responsible for during security incidents?	The AWS Cloud Security Whitepaper (available at http://aws.amazon.com/security) provides additional details.
	SEF-02.4	Have you tested your security incident response plans in the last year?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Reporting</i>	SEF-03.1	Does your security information and event management (SIEM) system merge data sources (app logs, firewall logs, IDS logs, physical access logs, etc.) for granular analysis and alerting?	
	SEF-03.2	Does your logging and monitoring framework allow isolation of an incident to specific tenants?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Response Legal Preparation</i>	SEF-04.1	Does your incident response plan comply with industry standards for legally admissible chain-of-custody management processes and controls?	
	SEF-04.2	Does your incident response capability include the use of legally admissible forensic data collection and analysis techniques?	
	SEF-04.3	Are you capable of supporting litigation holds (freeze of data from a specific point in time) for a specific tenant without freezing other tenant data?	
	SEF-04.4	Do you enforce and attest to tenant data separation when producing data in response to legal subpoenas?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Response Metrics</i>	SEF-05.1	Do you monitor and quantify the types, volumes and impacts on all information security incidents?	
	SEF-05.2	Will you share statistical information for security incident data with your tenants upon request?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Supply Chain Management, Transparency and Accountability <i>Data Quality and Integrity</i>	STA-01.1	Do you inspect and account for data quality errors and associated risks, and work with your cloud supply-chain partners to correct them?	Customers retain control and ownership over the quality of their data and potential quality errors that may arise through their usage of AWS services. Refer to AWS SOC report for specific details in relation to Data Integrity and Access Management (including least privilege access)
	STA-01.2	Do you design and implement controls to mitigate and contain data security risks through proper separation of duties, role-based access, and least-privileged access for all personnel within your supply chain?	
Supply Chain Management, Transparency and Accountability <i>Incident Reporting</i>	STA-02.1	Do you make security incident information available to all affected customers and providers periodically through electronic methods (e.g. portals)?	AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC reports provides details on the specific control activities executed by AWS. The AWS Cloud Security Whitepaper (available at http://aws.amazon.com/security) provides additional details.
Supply Chain Management, Transparency and Accountability <i>Network / Infrastructure Services</i>	STA-03.1	Do you collect capacity and use data for all relevant components of your cloud service offering?	AWS manages capacity and utilization data in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	STA-03.2	Do you provide tenants with capacity planning and use reports?	
Supply Chain Management, Transparency and Accountability <i>Provider Internal Assessments</i>	STA-04.1	Do you perform annual internal assessments of conformance and effectiveness of your policies, procedures, and supporting measures and metrics?	AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 15 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Supply Chain Management, Transparency and Accountability <i>Third Party Agreements</i>	STA-05.1	Do you select and monitor outsourced providers in compliance with laws in the country where the data is processed, stored and transmitted?	Personnel security requirements for third-party providers supporting AWS systems and devices are established in a Mutual Non-Disclosure Agreement between AWS' parent organization, Amazon.com, and the respective third-party provider. The Amazon Legal Counsel and the AWS Procurement team define AWS third party provider personnel security requirements in contract agreements with the third party provider. All persons working with AWS information must at a minimum, meet the screening process for pre-employment background checks and sign a Non-Disclosure Agreement (NDA) prior to being granted access to AWS information. AWS does not generally outsource development of AWS services to subcontractors.
	STA-05.2	Do you select and monitor outsourced providers in compliance with laws in the country where the data originates?	
	STA-05.3	Does legal counsel review all third-party agreements?	
	STA-05.4	Do third-party agreements include provision for the security and protection of information and assets?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	STA-05.5	Do you provide the client with a list and copies of all sub processing agreements and keep this updated?	
Supply Chain Management, Transparency and Accountability <i>Supply Chain Governance Reviews</i>	STA-06.1	Do you review the risk management and governance processes of partners to account for risks inherited from other members of that partner's supply chain?	AWS maintains formal agreements with key third party suppliers and implements appropriate relationship management mechanisms in line with their relationship to the business. AWS' third party management processes are reviewed by independent auditors as part of AWS ongoing compliance with SOC and ISO 27001.
Supply Chain Management, Transparency and Accountability <i>Supply Chain Metrics</i>	STA-07.1	Are policies and procedures established, and supporting business processes and technical measures implemented, for maintaining complete, accurate and relevant agreements (e.g., SLAs) between providers and customers (tenants)?	
	STA-07.2	Do you have the ability to measure and address non-conformance of provisions and/or terms across the entire supply chain (upstream/downstream)?	
	STA-07.3	Can you manage service-level conflicts or inconsistencies resulting from disparate supplier relationships?	
	STA-07.4	Do you review all agreements, policies and processes at least annually?	
Supply Chain Management, Transparency and Accountability <i>Third Party Assessment</i>	STA-08.1	Do you assure reasonable information security across your information supply chain by performing an annual review?	
	STA-8.2	Does your annual review include all partners/third-party providers upon which your information supply chain depends?	
Supply Chain Management, Transparency	STA-09.1	Do you permit tenants to perform independent vulnerability assessments?	Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for

Control Group	CID	Consensus Assessment Questions	AWS Response
and Accountability <i>Third Party Audits</i>	STA-09.2	Do you have external third party services conduct vulnerability scans and periodic penetration tests on your applications and networks?	these types of scans can be initiated by submitting a request via the AWS Vulnerability / Penetration Testing Request Form . AWS Security regularly engages independent security firms to perform external vulnerability threat assessments. The AWS SOC reports provides additional details on the specific control activities executed by AWS.
Threat and Vulnerability Management <i>Antivirus / Malicious Software</i>	TVM-01.1	Do you have anti-malware programs that support or connect to your cloud service offerings installed on all of your systems?	AWS' program, processes and procedures to managing antivirus / malicious software is in alignment with ISO 27001 standards. Refer to AWS SOC reports provides further details. In addition, refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	TVM-01.2	Do you ensure that security threat detection systems using signatures, lists or behavioral patterns are updated across all infrastructure components within industry accepted time frames?	
Threat and Vulnerability Management <i>Vulnerability / Patch Management</i>	TVM-02.1	Do you conduct network-layer vulnerability scans regularly as prescribed by industry best practices?	Customers retain control of their own guest operating systems, software and applications and are responsible for performing vulnerability scans and patching of their own systems. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. AWS Security regularly scans all Internet-facing service endpoint IP addresses for vulnerabilities. AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. AWS' own maintenance and system patching generally do not impact customers. Refer to AWS Cloud Security Whitepaper for further information - available at http://aws.amazon.com/security . Refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	TVM-02.2	Do you conduct application-layer vulnerability scans regularly as prescribed by industry best practices?	
	TVM-02.3	Do you conduct local operating system-layer vulnerability scans regularly as prescribed by industry best practices?	
	TVM-02.4	Will you make the results of vulnerability scans available to tenants at their request?	
	TVM-02.5	Do you have a capability to rapidly patch vulnerabilities across all of your computing devices, applications and systems?	
	TVM-02.6	Will you provide your risk-based systems patching time frames to your tenants upon request?	
Threat and Vulnerability Management <i>Mobile Code</i>	TVM-03.1	Is mobile code authorized before its installation and use, and the code configuration checked, to ensure that the authorized mobile code operates according to a clearly defined security policy?	AWS allows customers to manage client and mobile applications to their own requirements.

Control Group	CID	Consensus Assessment Questions	AWS Response
	TVM - 03.2	Is all unauthorized mobile code prevented from executing?	

Appendix B: AWS alignment with the Australian Signals Directorate (ASD) Cloud Computing Security Considerations

The Cloud Computing Security Considerations was created to assist agencies in performing a risk assessment of services offered by Cloud Service Providers. The following provides AWS alignment to the Security Considerations, published on September 2012. For additional details refer to:

http://www.asd.gov.au/publications/csocprotect/Cloud_Computing_Security_Considerations.pdf

Key Area	Questions	AWS RESPONSE
Maintaining Availability and Business Functionality	a. Business criticality of data or functionality. Am I moving business critical data or functionality to the cloud?	AWS customers retain control and ownership of their content. Customers are responsible for the classification and use of their content.
	b. Vendor's business continuity and disaster recovery plan. Can I thoroughly review a copy of the vendor's business continuity and disaster recovery plan that covers the availability and restoration of both my data and the vendor's services that I use? How much time does it take for my data and the services that I use to be recovered after a disaster, and do the vendor's other customers that are larger and pay more money than me get prioritization?	AWS customers retain control and ownership of their data. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS SOC 1 Type 2 report provides further details. ISO 27001 standard Annex A, domain 11 provides additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification. Customers utilize AWS to enable faster disaster recovery of their critical IT systems without incurring the infrastructure expense of a second physical site. The AWS cloud supports many popular disaster recovery (DR) architectures from "pilot light" environments that are ready to scale up at a moment's notice to "hot standby" environments that enable rapid failover. To learn more about Disaster Recovery on AWS visit https://aws.amazon.com/disaster-recovery/. AWS provides customers with the capability to implement a robust continuity plan, including the utilization of frequent server instance back-ups, data redundancy replication, and multi-region/availability zone deployment architectures. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 9 and the AWS SOC 1 Type II report for additional information.

Key Area	Questions	AWS RESPONSE
	c. My data backup plan. Will I spend additional money to maintain an up to date backup copy of my data located either at my agency's premises, or stored with a second vendor that has no common points of failure with the first vendor?	AWS customers retain control and ownership of their content and it is the customer's responsibility to manage their data backup plans. AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 service is designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy. For information on data durability and redundancy, please refer to the AWS website. AWS offers a range of cloud computing services to support Disaster Recovery. To learn more about Disaster Recovery on AWS visit https://aws.amazon.com/disaster-recovery/.
	d. My business continuity and disaster recovery plan. Will I spend additional money to replicate my data or business functionality with a second vendor that uses a different data center and ideally has no common points of failure with the first vendor? This replication should preferably be configured to automatically "failover", so that if one vendor's services become unavailable, control is automatically and smoothly transitioned to the other vendor.	Customers retain control and ownership of their data. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security. AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. AWS data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. This means that Availability Zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by Region). In addition to discrete uninterruptable power supply (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability Zones are all redundantly connected to multiple tier-1 transit providers. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Distributing applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes including natural disasters or system failures. AWS SOC 1 Type 2 report provides further details. ISO 27001 standard Annex A, domain 11 provides additional details. AWS has been validated

Key Area	Questions	AWS RESPONSE
		and certified by an independent auditor to confirm alignment with ISO 27001 certification.
	e. My network connectivity to the cloud. Is the network connectivity between my agency's users and the vendor's network adequate in terms of availability, traffic throughput (bandwidth), delays (latency) and packet loss?	Customers can also choose their network path to AWS facilities, including multiple VPN endpoints in each AWS Region. In addition, AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security.

Key Area	Questions	AWS RESPONSE
	f. Vendor's guarantee of availability. Does the Service Level Agreement (SLA guarantee that the vendor will provide adequate system availability and quality of service, using their robust system architecture and business processes?	AWS does commit to high levels of availability in its service level agreements (SLAs). For example, Amazon EC2 commits to annual uptime percentage of at least 99.95% during the service year. Amazon S3 commits to monthly uptime percentage of at least 99.99% Service credits are provided in the case these availability metrics are not met. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Distributing applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes including natural disasters or system failures. AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel. AWS Network Management is regularly reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm.
	g. Impact of outages. Can I tolerate the maximum possible downtime of the SLA? Are the scheduled outage windows acceptable both in duration and time of day, or will scheduled outages interfere with my critical business processes?	AWS does not require systems to be brought offline to perform regular maintenance and system patching. AWS' own maintenance and system patching generally do not impact customers. Maintenance of instances themselves is controlled by the customer.
	h. SLA inclusion of scheduled outages. Does the SLA guaranteed availability percentage include scheduled outages?	AWS does not operate an environment with scheduled outage as AWS provides customers the ability to architect their environment to take advantage of multiple Availability Zones and regions.
	i. SLA compensation. Does the SLA adequately reflect the actual damage caused by a breach of the SLA such as unscheduled downtime or data loss?	AWS provides customer remuneration for losses they may incur due to outages in alignment with AWS' Service Level Agreement.

Key Area	Questions	AWS RESPONSE
	j. Data integrity and availability. How does the vendor implement mechanisms such as redundancy and offsite backups to prevent corruption or loss of my data, and guarantee both the integrity and the availability of my data?	AWS data integrity controls as described in AWS SOC 1 Type II report provides reasonable assurance that data integrity is maintained through all phases including transmission, storage and processing. In addition, refer to ISO 27001 standard, Annex A, domain 12 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. Data centers are built in clusters in various global regions. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. You choose where to store your data by specifying a region (for Amazon S3) or an availability zone within a region (for EBS). Data stored in Amazon Elastic Block Store (Amazon EBS) is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones. Amazon S3 provides a highly durable storage infrastructure. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 Region. Once stored, Amazon S3 maintains the durability of objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. Data stored in S3 is designed to provide 99.99999999% durability and 99.99% availability of objects over a given year. Refer to AWS Overview of Security Processes whitepaper for additional details - available at http://aws.amazon.com/security
	k. Data restoration. If I accidentally delete a file, email or other data, how much time does it take for my data to be partially or fully restored from backup, and is the maximum acceptable time captured in the SLA?	AWS customers retain control and ownership of their data. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region.
	l. Scalability. How much available spare computing resources does the vendor provide to enable my usage of the vendor's services to scale at short notice?	The AWS cloud is distributed, highly secure and resilient, giving customers large scaling potential. Customers may scale up or down, paying for only what they use.

Key Area	Questions	AWS RESPONSE
	m. Changing vendor. If I want to move my data to my agency or to a different vendor, or if the vendor suddenly becomes bankrupt or otherwise quits the cloud business, how do I get access to my data in a vendor-neutral format to avoid vendor lock-in? How cooperative will the vendor be? How do I ensure that my data is permanently deleted from the vendor's storage media? For Platform as a Service, which standards does the vendor use that facilitate portability and interoperability to easily move my application to a different vendor or to my agency?	Customers retain control and ownership of their data. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security. AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS.
Protecting Data from Unauthorized Access by a Third Party	a. Choice of cloud deployment model. Am I considering using a potentially less secure public cloud, a potentially more secure hybrid cloud or community cloud, or a potentially most secure private cloud?	AWS' Compliance and Security teams have established an information security framework and policies based on the Control Objectives for Information and related Technology (COBIT) framework. The AWS security framework integrates the ISO 27002 best practices and the PCI Data Security Standard. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security. AWS provides third-party attestations, certifications, Service Organization Controls 1 (SOC 1) Type II report and other relevant compliance reports directly to our customers under NDA. Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can easily customize the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your web servers that has access to the Internet, and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet. Additionally, you can create a Hardware Virtual Private Network (VPN) connection between your corporate data center and your VPC and leverage the AWS cloud as an extension of your corporate data center

Key Area	Questions	AWS RESPONSE
	b. Sensitivity of my data. Is my data to be stored or processed in the cloud classified, sensitive, private, or data that is publicly available such as information from my public web site? Does the aggregation of my data make it more sensitive than any individual piece of data? For example, the sensitivity may increase if storing a significant amount of data, or storing a variety of data that if compromised would facilitate identity theft. If there is a data compromise, could I demonstrate my due diligence to senior management, government officials and the public?	AWS customers retain control and ownership of their data and may implement a structured data-classification program to meet their requirements.
	c. Legislative obligations. What obligations do I have to protect and manage my data under various legislation, for example the Privacy Act, the Archives Act, as well as other legislation specific to the type of data? Will the vendor contractually accept adhering to these obligations to help me ensure that the obligations are met to the satisfaction of the Australian Government?	AWS customers retain responsibility to ensure their usage of AWS is within compliance of applicable laws and regulations. AWS communicates its security and control environment to customers through industry certifications and third-party attestations, white papers (available at http://aws.amazon.com/security) and providing certifications, reports and other relevant documentation directly to AWS customers. AWS has published a whitepaper on using AWS in the context of Australian privacy considerations, available here.

Key Area	Questions	AWS RESPONSE
	d. Countries with access to my data. In which countries is my data stored, backed up and processed? Which foreign countries does my data transit? In which countries is the failover or redundant data centers? Will the vendor notify me if the answers to these questions change?	AWS customers choose the AWS Region or regions in which their content and servers will be located. This allows customers with geographic specific requirements to establish environments in a location of their choice. AWS customers in Australia can choose to deploy their AWS services exclusively in the Asia Pacific (Sydney) region and store their content onshore in Australia. If the customer makes this choice, their content will be located in Australia unless the customer chooses to move the data. Customers can replicate and back up content in more than one region, but AWS does not move or replicate customer content outside of the customer's chosen region or regions. AWS is vigilant about customers' security and does not disclose or move data in response to a request from the Australian, U.S. or other government unless legally required to do so in order to comply with a legally valid and binding order, such as a subpoena or a court order, or as is otherwise required by applicable law. Non-U.S. governmental or regulatory bodies typically must use recognized international processes, such as Mutual Legal Assistance Treaties with the U.S. government, to obtain valid and binding orders. Additionally, our practice is to notify customers where practicable before disclosing their content so they can seek protection from disclosure, unless we are legally prevented from doing so.

Key Area	Questions	AWS RESPONSE
	e. Data encryption technologies. Are hash algorithms, encryption algorithms and key lengths deemed appropriate by the DSD ISM used to protect my data when it is in transit over a network, and stored on both the vendor's computers and on backup media? The ability to encrypt data while it is being processed by the vendor's computers is still an emerging technology and is an area of current research by industry and academia. Is the encryption deemed strong enough to protect my data for the duration of time that my data is sensitive?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. VPC sessions are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Customers may also use third-party encryption technologies. Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications. AWS cryptographic processes are reviewed by independent third-party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm. The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance. The AWS CloudHSM service works with Amazon Virtual Private Cloud (VPC). CloudHSMs are provisioned inside your VPC with an IP address that you specify, providing simple and private network connectivity to your Amazon Elastic Compute Cloud (EC2) instances. Placing CloudHSMs near your EC2 instances decreases network latency, which can improve application performance. AWS provides dedicated and exclusive access to CloudHSMs, isolated from other AWS customers. Available in multiple Regions and Availability Zones (AZs), AWS CloudHSM allows you to add secure and durable key storage to your Amazon EC2 applications.
	f. Media sanitization. What processes are used to sanitize the storage media storing my data at its end of life, and are the processes deemed appropriate by the DSD ISM?	When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security.

Key Area	Questions	AWS RESPONSE
	g. Vendor's remote monitoring and management. Does the vendor monitor, administer or manage the computers that store or process my data? If yes, is this performed remotely from foreign countries or from Australia? Can the vendor provide patch compliance reports and other details about the security of workstations used to perform this work, and what controls prevent the vendor's employees from using untrustworthy personally owned laptops?	Moving IT infrastructure to AWS services creates a model of shared responsibility between the customer and AWS. This shared model can help relieve customer's operational burden as AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall.
	h. My monitoring and management. Can I use my existing tools for integrity checking, compliance checking, security monitoring and network management, to obtain visibility of all my systems regardless of whether these systems are located locally or in the cloud? Do I have to learn to use additional tools provided by the vendor? Does the vendor even provide such a mechanism for me to perform monitoring?	AWS Cloudwatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com. The AWS Trusted Advisor inspects your AWS environment and makes recommendations when opportunities exist to save money, improve system performance and reliability, or help close security gap.
	i. Data ownership. Do I retain legal ownership of my data, or does it belong to the vendor and may be considered an asset for sale by liquidators if the vendor declares bankruptcy?	AWS customers retain ownership and control of their data. AWS only uses each customer's content to provide the AWS services selected by each customer to that customer and does not use customer content for any secondary purposes. AWS treats all customer content the same and has no insight as to what type of content the customer chooses to store in AWS. AWS simply makes available the compute, storage, database and networking services selected by customer – AWS does not require access to customer content to provide its services.

Key Area	Questions	AWS RESPONSE
	j. Gateway technologies. What technologies does the vendor use to create a secure gateway environment? Examples include firewalls, traffic flow filters, content filters, and antivirus software and data diodes where appropriate.	The AWS network provides significant protection against traditional network security issues and customers can implement further protection. Refer to the AWS Overview of Security whitepaper (available at http://aws.amazon.com/security) for additional details. Amazon assets (e.g. laptops) are configured with anti-virus software that includes e-mail filtering and malware detection. AWS Network Firewall management and Amazon's anti-virus program are reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm.
	k. Gateway certification. Is the vendor's gateway environment certified against government security standards and regulations?	AWS obtains certain industry certifications and independent third-party attestations which include the AWS Gateway environment.
	l. Email content filtering. For email Software as a Service, does the vendor provide customizable email content filtering that can enforce my agency's email content policy?	A Customer can utilize a system to host e-mail capabilities, however in that case it is the Customer's responsibility to employ the appropriate levels of spam and malware protection at e-mail entry and exit points and update spam and malware definitions when new releases are made available.

Key Area	Questions	AWS RESPONSE
	m. Policies and processes supporting the vendor's IT security posture. Can I have details of how the vendor's computer and network security posture is supported by policies and processes including threat and risk assessments, ongoing vulnerability management, a change management process that incorporates security, penetration testing, logging and regular log analysis, use of security products endorsed by the Australian Government, and compliance with Australian government security standards and regulations?	Policies and procedures have been established by AWS Information Security based upon the COBIT framework, ISO 27001 standards and the PCI DSS requirements. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. In addition AWS publishes a SOC 1 Type II report. Refer to the SOC 1 report for further details. The AWS Risk and Compliance whitepaper for additional details - available at http://aws.amazon.com/security. AWS customers are able to identify key controls managed by AWS. Key controls are critical to the customer's control environment and require an external attestation of the operating effectiveness of these key controls in order to comply with compliance requirements—such as the annual financial audit. For this purpose, AWS publishes a wide range of specific IT controls in its Service Organization Controls 1 (SOC 1) Type II report. The SOC 1 report, formerly the Statement on Auditing Standards (SAS) No. 70, Service Organizations report and formerly referred to as the Statement on Standards for Attestation Engagements No. 16 (SSAE 16) report, is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). The SOC 1 audit is an in-depth audit of both the design and operating effectiveness of AWS' defined control objectives and control activities (which include control objectives and control activities over the part of the infrastructure AWS manages). "Type II" refers to the fact that each of the controls described in the report are not only evaluated for adequacy of design, but are also tested for operating effectiveness by the external auditor. Because of the independence and competence of AWS' external auditor, controls identified in the report should provide customers with a high level of confidence in AWS' control environment.
	n. Technologies supporting the vendor's IT security posture. Can I have details of how the vendor's computer and network security posture is supported by direct technical controls including timely application of security patches, regularly updated antivirus software, defense in depth mechanisms to protect against unknown vulnerabilities, hardened operating systems and software applications configured with the strongest possible security settings, intrusion detection and prevention systems,	AWS provides third-party attestations, certifications, Service Organization Controls 1 (SOC 1) Type II report and other relevant compliance reports directly to our customers under NDA. AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment.

Key Area	Questions	AWS RESPONSE
	and data loss prevention mechanisms?	
	o. Auditing the vendor's IT security posture. Can I audit the vendor' implementation of security measures, including performing scans and other penetration testing of the environment provided to me? If there is justifiable reason why auditing is not possible, which reputable third party has performed audits and other vulnerability assessments? What sort of internal audits does the vendor perform, and which compliance standards and other recommended practices from organization's such as the Cloud Security Alliance are used for these assessments? Can I thoroughly review a copy of recent resulting reports?	AWS provides third-party attestations, certifications, Service Organization Controls 1 (SOC 1) Type II report and other relevant compliance reports directly to our customers under NDA. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these types of scans can be initiated by submitting a request via the AWS Vulnerability / Penetration Testing Request Form. AWS Security regularly engages independent security firms to perform external vulnerability threat assessments. The AWS SOC 1 Type 2 report provides additional details on the specific control activities executed by AWS.

Key Area	Questions	AWS RESPONSE
	p. User authentication. What identity and access management systems does the vendor support for users to log in to use Software as a Service?	AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources. AWS supports identity federation that makes it easier to manage users by maintaining their identities in a single place. AWS IAM includes support for the Security Assertion Markup Language (SAML) 2.0, an open standard used by many identity providers. This new feature enables federated single sign-on, or SSO, empowering users to log into the AWS Management Console or make programmatic calls to AWS APIs, by using assertions from a SAML-compliant identity provider, such as Shibboleth and Windows Active Directory Federation Services.
	q. Centralized control of data. What user training, policies and technical controls prevent my agency's users from using unapproved or insecure computing devices without a trusted operating environment to store or process sensitive data accessed using Software as a Service?	N/A

Key Area	Questions	AWS RESPONSE
	r. Vendor's physical security posture. Does the vendor use physical security products and devices that are endorsed by the Australian Government? How is the vendor's physical data center designed to prevent the tampering or theft of servers, infrastructure and the data stored thereon? Is the vendor's physical data center accredited by an authoritative third party?	The definition of AWS-defined logical and physical controls is documented in the SOC 1 Type II report, and the report is available for review by audit and compliance teams. AWS ISO 27001 and other certifications are also available for auditors to review. Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. Physical access is strictly controlled both at the perimeter and at building ingress points and includes, but is not limited to, professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. Physical access points to server locations are recorded by closed circuit television camera (CCTV) as defined in the AWS Data Center Physical Security Policy. Images are retained for 90 days, unless limited to 30 days by legal or contractual obligations AWS provides data center physical access and information to approved employees and contractors who have a legitimate business need for such privileges. All visitors are required to present identification and are signed in and escorted by authorized staff. See the SOC 1 Type II report for specific controls related to physical access, data center access authorization, and other related controls. Refer to ISO 27001 standard, Annex A, domain 9 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	s. Software and hardware procurement. What procurement process is used to ensure that cloud infrastructure software and hardware has been supplied by a legitimate source and has not been maliciously modified in transit?	In alignment with ISO 27001 standards, AWS Hardware assets are assigned an owner, tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standard, Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Key Area	Questions	AWS RESPONSE
Protecting Data from Unauthorized Access by the Vendor's Customers	a. Customer segregation. What assurance do I have that the virtualization and "multi-tenancy" mechanisms guarantee adequate logical and network segregation between multiple tenants, so that a malicious customer using the same physical computer as me cannot access my data?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers retain control and ownership of their data, thus it is their responsibility to choose to encrypt the data. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. VPC sessions are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security.
	b. Weakening my security posture. How would using the vendor's cloud infrastructure weaken my agency's existing network security posture? Would the vendor advertise me as one of their customers without my explicit consent, thereby assisting an adversary that is specifically targeting me?	AWS customers are considered confidential and would not advertise customer details without explicit consent. Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.
	c. Dedicated servers. Do I have some control over which physical computer runs my virtual machines? Can I pay extra to ensure that no other customer can use the same physical computer as me e.g. dedicated servers or virtual private cloud?	VPC allows customers to launch Amazon EC2 instances that are physically isolated at the host hardware level; they will run on single tenant hardware. A VPC can be created with 'dedicated' tenancy, in which case all instances launched into the VPC will utilize this feature. Alternatively, a VPC may be created with 'default' tenancy, but customers may specify 'dedicated' tenancy for particular instances launched into the VPC.
	d. Media sanitization. When I delete portions of my data, what processes are used to sanitize the storage media before it is made available to another customer, and are the processes deemed appropriate by the DSD ISM?	Customers retain ownership and control of their content and provide customers with the ability to delete their data. When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security.

Key Area	Questions	AWS RESPONSE
Protecting Data from Unauthorized Access by Rogue Vendor Employees	a. Data encryption key management. Does the vendor know the password or key used to decrypt my data, or do I encrypt and decrypt the data on my computer so the vendor only ever has encrypted data?	AWS Customers manage their own encryption unless they are utilizing AWS server side encryption service. In this case, AWS does create a unique encryption key per tenant. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security .
	b. Vetting of vendor's employees. What personnel employment checks and vetting processes does the vendor perform to ensure that employees are trustworthy?	AWS conducts criminal background checks, as permitted by applicable law, as part of pre-employment screening practices for employees commensurate with the employee's position and level of access to AWS facilities.
	c. Auditing vendor's employees. What robust identity and access management system do the vendor's employees use? What auditing process is used to log and review the actions performed by the vendor's employees?	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC 1 Type 2 report outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Overview of Security Processes whitepaper for additional details - available at http://aws.amazon.com/security .
	d. Visitors to data center. Are visitors to data centers escorted at all times, and is the name and other personal details of every visitor verified and recorded?	All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff. AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to data centers by AWS employees is routinely logged and audited.
	e. Physical tampering by vendor's employees. Is network cabling professionally installed to Australian standards or internationally acceptable standards, to help avoid the vendor's employees from accidentally connecting cables to the wrong computers, and to help readily highlight any deliberate attempts by the vendor's employees to tamper with the cabling?	Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. This includes appropriate protection for network cables. The AWS SOC 1 Type 2 report provides additional details on the specific control activities executed by AWS. Refer to ISO 27001 standard, Annex A, domain 9 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Key Area	Questions	AWS RESPONSE
	f. Vendor's subcontractors. Do the answers to these questions apply equally to all of the vendor's subcontractors?	Provisioning contractor / vendor access is managed the same for both employees and contractors, with responsibility shared across Human Resources (HR), Corporate Operations and Service Owners. Vendors are subject to the same access requirements as employees.
Handling Security Incidents	a. Timely vendor support. Is the vendor readily contactable and responsive to requests for support, and is the maximum acceptable response time captured in the SLA or simply a marketing claim that the vendor will try their best? Is the support provided locally, or from a foreign country, or from several foreign countries using an approach that follows the sun? What mechanism does the vendor use to obtain a real-time understanding of the security posture of my use of the vendor's services so that the vendor can provide support?	AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced and technical support engineers. The service helps customers of all sizes and technical abilities to successfully utilize the products and features provided by Amazon Web Services. All AWS Support tiers offer customers of AWS Infrastructure Services an unlimited number of support cases with pay-by-the-month pricing and no long-term contracts. The four tiers provide developers and businesses the flexibility to choose the support tiers that meet their specific needs.
	b. Vendor's incident response plan. Does the vendor have a security incident response plan that specifies how to detect and respond to security incidents, in a way that is similar to incident handling procedures detailed in the DSD ISM? Can I thoroughly review a copy?	The Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide 24 x 7 x 365 coverage to detect incidents and to manage the impact and resolution. AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC 1 Type 2 report provides details on the specific control activities executed by AWS. The AWS Overview of Security Processes whitepaper (available at http://aws.amazon.com/security) provides additional details.
	c. Training of vendor's employees. What qualifications, certifications and regular information security awareness training do the vendor's employees require, to know how to use the vendor's systems in a secure manner and to identify potential security incidents?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security .

Key Area	Questions	AWS RESPONSE
	d. Notification of security incidents. Will the vendor notify me via secure communications of security incidents that are more serious than an agreed threshold, especially in cases where the vendor might be liable? Will the vendor automatically notify law enforcement or other authorities, who may confiscate computing equipment used to store or process my data?	Notification of security incidents are handled on a case-by-case basis and as required by applicable law. Any notification is performed via secure communications.
	e. Extent of vendor support. How much assistance will the vendor provide me with investigations if there is a security breach such as an unauthorized disclosure of my data, or if there is a need to perform legal electronic discovery of evidence?	AWS provides infrastructure and customers manage everything else, including the operating system, the network configuration and the installed applications. Customers are responsible for responding appropriately to legal procedures involving the identification, collection, processing, analysis and production of electronic documents they store or process using AWS. Upon request, AWS may work with customers who require AWS' assistance in legal proceedings.
	f. My access to logs. How do I obtain access to time synchronized audit logs and other logs to perform a forensic investigation, and how are the logs created and stored to be suitable evidence for a court of law?	Customers retain control of their own guest operating systems, software and applications and are responsible for developing logical monitoring of the conditions of these systems. In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS CloudTrail provides a simple solution to log user activity that helps alleviate the burden of running a complex logging system. Refer to aws.amazon.com/cloudtrail for additional details. AWS Cloudwatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com.
	g. Security incident compensation. How will the vendor adequately compensate me if the vendor's actions, faulty software or hardware contributed to a security breach?	AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC 1 Type 2 report provides details on the specific control activities executed by AWS. The AWS Overview of Security Processes whitepaper (available at http://aws.amazon.com/security) provides additional details.

Key Area	Questions	AWS RESPONSE
	h. Data spills. If data that I consider is too sensitive to be stored in the cloud is accidentally placed into the cloud, referred to as a data spill, how can the spilled data be deleted using forensic sanitization techniques? Is the relevant portion of physical storage media zeroed whenever data is deleted? If not, how long does it take for deleted data to be overwritten by customers as part of normal operation, noting that clouds typically have significant spare unused storage capacity? Can the spilled data be forensically deleted from the vendor's backup media? Where else is the spilled data stored, and can it be forensically deleted?	Customers retain ownership and control of their content. All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPSec tunnels to VPC are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security.

Appendix C: Glossary of Terms

Authentication: Authentication is the process of determining whether someone or something is, in fact, who or what it is declared to be.

Availability Zone: Amazon EC2 locations are composed of regions and Availability Zones. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same region.

DSS: The Payment Card Industry Data Security Standard (DSS) is a worldwide information security standard assembled and managed by the Payment Card Industry Security Standards Council.

EBS: Amazon Elastic Block Store (EBS) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance.

FedRAMPsm: The Federal Risk and Authorization Management Program (FedRAMPsm) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMPsm is mandatory for Federal Agency cloud deployments and service models at the low and moderate risk impact levels.

FISMA: The Federal Information Security Management Act of 2002. The act requires each federal agency to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including those provided or managed by another agency, contractor, or other source.

FIPS 140-2: The Federal Information Processing Standard (FIPS) Publication 140-2 is a US government security standard that specifies the security requirements for cryptographic modules protecting sensitive information.

GLBA: The Gramm–Leach–Bliley Act (GLB or GLBA), also known as the Financial Services Modernization Act of 1999, sets forth requirements for financial institutions with regard to, among other things, the disclosure of nonpublic customer information and the protection of threats in security and data integrity.

HIPAA: The Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires the establishment of national standards for electronic health care transactions and national identifiers for providers, health insurance plans, and employers. The Administration Simplification provisions also address the security and privacy of health data. The standards are meant to improve the efficiency and effectiveness of the nation's health care system by encouraging the widespread use of electronic data interchange in the U.S. health care system.

Hypervisor: A hypervisor, also called Virtual Machine Monitor (VMM), is software/hardware platform virtualization software that allows multiple operating systems to run on a host computer concurrently.

IAM: AWS Identity and Access Management (IAM) enables a customer to create multiple Users and manage the permissions for each of these Users within their AWS Account.

ITAR: International Traffic in Arms Regulations (ITAR) is a set of United States government regulations that control the export and import of defense-related articles and services on the United States Munitions List (USML). Government agencies and contractors must comply with ITAR and restrict access to protected data.



ISAE 3402: The International Standards for Assurance Engagements No. 3402 (ISAE 3402) is the international standard on assurance engagements. It was put forth by the International Auditing and Assurance Standards Board (IAASB), a standard-setting board within the International Federation of Accountants (IFAC). ISAE 3402 is now the new globally recognized standard for assurance reporting on service organizations.

ISO 9001: AWS' ISO 9001 certification directly supports customers who develop, migrate and operate their quality-controlled IT systems in the AWS cloud. Customers can leverage AWS' compliance reports as evidence for their own ISO 9001 programs and industry-specific quality programs, such as GxP in life sciences, ISO 13485 in medical devices, AS9100 in aerospace, and ISO/TS 16949 in automotive. AWS customers who don't have quality system requirements will still benefit from the additional assurance and transparency that an ISO 9001 certification provides.

ISO 27001: ISO/IEC 27001 is an Information Security Management System (ISMS) standard published by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). ISO 27001 formally specifies a management system that is intended to bring information security under explicit management control. Being a formal specification means that it mandates specific requirements. Organizations that claim to have adopted ISO/IEC 27001 can therefore be audited and certified compliant with the standard.

NIST: National Institute of Standards and Technology. This agency sets detailed security standards as needed by industry or government programs. Compliance with FISMA requires agencies to adhere to NIST standards.

Object: The fundamental entities stored in Amazon S3. Objects consist of object data and metadata. The data portion is opaque to Amazon S3. The metadata is a set of name-value pairs that describe the object. These include some default metadata such as the date last modified and standard HTTP metadata such as Content-Type. The developer can also specify custom metadata at the time the Object is stored.

PCI: Refers to the Payment Card Industry Security Standards Council, an independent council originally formed by American Express, Discover Financial Services, JCB, MasterCard Worldwide and Visa International, with the goal of managing the ongoing evolution of the Payment Card Industry Data Security Standard.

QSA: The Payment Card Industry (PCI) Qualified Security Assessor (QSA) designation is conferred by the PCI Security Standards Council to those individuals that meet specific qualification requirements and are authorized to perform PCI compliance assessments.

SAS 70: Statement on Auditing Standards No. 70: Service Organizations is an auditing statement issued by the Auditing Standards Board of the American Institute of Certified Public Accountants (AICPA). SAS 70 provides guidance to service auditors when assessing the internal controls of a service organization (such as AWS) and issuing a service auditor's report. SAS 70 also provides guidance to auditors of financial statements of an entity that uses one or more service organizations. The SAS 70 report has been replaced by the Service Organization Controls 1 report.

Service: Software or computing ability provided across a network (e.g., EC2, S3, VPC, etc.).

Service Level Agreement (SLA): A service level agreement is a part of a service contract where the level of service is formally defined. The SLA is used to refer to the contracted delivery time (of the service) or performance.

SOC 1: Service Organization Controls 1 (SOC 1) Type II report, formerly the Statement on Auditing Standards (SAS) No. 70, Service Organizations report (formerly referred to as the SSAE 16 report), is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). The



international standard is referenced as the International Standards for Assurance Engagements No. 3402 (ISAE 3402).

SSAE 16 [deprecated]: The Statement on Standards for Attestation Engagements No. 16 (SSAE 16) is an attestation standard published by the Auditing Standards Board (ASB) of the American Institute of Certified Public Accountants (AICPA). The standard addresses engagements undertaken by a service auditor for reporting on controls at organizations that provide services to user entities, for which a service organization's controls are likely to be relevant to a user entities internal control over financial reporting (ICFR). SSAE 16 effectively replaces Statement on Auditing Standards No. 70 (SAS 70) for service auditor's reporting periods ending on or after June 15, 2011.

SOC 2: Service Organization Controls 2 (SOC 2) reports are intended to meet the needs of a broad range of users that need to understand internal control at a service organization as it relates to security, availability, processing integrity, confidentiality and privacy. These reports are performed using the AICPA Guide: Reporting on Controls at a Service Organizations Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy and are intended for use by stakeholders (e.g., customers, regulators, business partners, suppliers, directors) of the service organization that have a thorough understanding of the service organization and its internal controls.

SOC 3: Service Organization Controls 3 (SOC 3) reports are designed to meet the needs of users who want assurance on the controls at a service organization related to security, availability, processing integrity, confidentiality, or privacy but do not have the need for or the knowledge necessary to make effective use of a SOC 2 Report. These reports are prepared using the AICPA/Canadian Institute of Chartered Accountants (CICA) Trust Services Principles, Criteria, and Illustrations for Security, Availability, Processing Integrity, Confidentiality, and Privacy. Because they are general use reports, SOC 3 Reports can be freely distributed or posted on a website as a seal.

Virtual Instance: Once an AMI has been launched, the resulting running system is referred to as an instance. All instances based on the same AMI start out identical and any information on them is lost when the instances are terminated or fail.

Version History

January 2016

- Added GxP Compliance Program
- Twelfth region added (Asia Pacific - Seoul)

December 2015

- Updates to certifications and third-party attestations summaries
- Added ISO 27017 certification
- Added ISO 27018 certification
- Eleventh region added (China - Beijing)

November 2015

- Update to CSA v3.0.1

August 2015

- Updates to in-scope services for PCI 3.1
- Updates to regions in-scope for PCI 3.1

May 2015

- Tenth region added (EU - Frankfurt)
- Updates to in-scope services for SOC 3
- SSAE 16 language deprecated

Apr 2015

- Updates to in-scope services for: FedRAMPsm, HIPAA, SOC 1, ISO 27001, ISO 9001

Feb 2015

- Updates to FIPS 140-2 VPN endpoints and SSL-terminating load balancers
- Updates to PCI DSS verbiage

Dec 2014

- Updates to certifications and third-party attestations summaries

Nov 2013 version

- Edits to IPsec tunnel encryption verbiage

Jun 2013 version

- Updates to certifications and third-party attestations summaries
- Updates to Appendix C: Glossary of Terms
- Minor changes to formatting

Jan 2013 version

- Edits to certifications and third-party attestations summaries

Nov 2012 version

- Edits to content and updated certification scope
- Added reference to the SOC 2 and MPAA

Jul 2012 version

- Edits to content and updated certification scope
- Addition of the CSA Consensus Assessments Initiative Questionnaire (Appendix A)

Jan 2012 version

- Minor edits to content based on updated certification scope



- Minor grammatical edits

Dec 2011 version

- Change to Certifications and Third-party Attestation section to reflect SOC 1/SSAE 16, FISMA Moderate, International Traffic in Arms Regulations, and FIPS 140-2
- Addition of S3 Server Side Encryption
- Added additional cloud computing issue topics

May 2011 version

- Initial release

Notices

© 2010-2016 Amazon.com, Inc., or its affiliates. This document is provided for informational purposes only. It represents AWS' current product offerings as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS' products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.



Amazon Web Services: Overview of Security Processes

August 2015

(Please consult <http://aws.amazon.com/security/> for the latest version of this paper)



Table of Contents

Introduction	5
Shared Security Responsibility Model	5
AWS Security Responsibilities	6
Customer Security Responsibilities	6
AWS Global Infrastructure Security	7
AWS Compliance Program	7
Physical and Environmental Security	8
Fire Detection and Suppression	8
Power	8
Climate and Temperature	8
Management	8
Storage Device Decommissioning	8
Business Continuity Management	9
Availability	9
Incident Response	9
Company-Wide Executive Review	9
Communication	9
Network Security	10
Secure Network Architecture	10
Secure Access Points	10
Transmission Protection	10
Amazon Corporate Segregation	10
Fault-Tolerant Design	11
Network Monitoring and Protection	12
AWS Access	14
Account Review and Audit	14
Background Checks	14
Credentials Policy	14
Secure Design Principles	14
Change Management	15
Software	15
Infrastructure	15
AWS Account Security Features	16
AWS Credentials	16

Passwords	17
AWS Multi-Factor Authentication (AWS MFA)	17
Access Keys	18
Key Pairs.....	18
X.509 Certificates	18
Individual User Accounts.....	19
Secure HTTPS Access Points.....	19
Security Logs	19
AWS Trusted Advisor Security Checks	20
AWS Service-Specific Security	20
Compute Services.....	20
Amazon Elastic Compute Cloud (Amazon EC2) Security.....	20
Auto Scaling Security	24
Networking Services.....	25
Amazon Elastic Load Balancing Security	25
Amazon Virtual Private Cloud (Amazon VPC) Security	26
Amazon Route 53 Security.....	31
Amazon CloudFront Security	32
AWS Direct Connect Security.....	34
Storage Services	35
Amazon Simple Storage Service (Amazon S3) Security	35
AWS Glacier Security.....	37
AWS Storage Gateway Security	38
AWS Import/Export Security.....	39
Database Services	41
Amazon DynamoDB Security	41
Amazon Relational Database Service (Amazon RDS) Security.....	42
Amazon Redshift Security	46
Amazon ElastiCache Security	48
Application Services	50
Amazon CloudSearch Security	50
Amazon Simple Queue Service (Amazon SQS) Security.....	51
Amazon Simple Notification Service (Amazon SNS) Security	51
Amazon Simple Workflow Service (Amazon SWF) Security.....	52
Amazon Simple Email Service (Amazon SES) Security	52
Amazon Elastic Transcoder Service Security.....	53
Amazon AppStream Security	54
Analytics Services	55

Amazon Elastic MapReduce (Amazon EMR) Security	55
Amazon Kinesis Security	56
AWS Data Pipeline Security	56
Deployment and Management Services	57
AWS Identity and Access Management (AWS IAM)	57
Amazon CloudWatch Security	58
AWS CloudHSM Security	59
AWS CloudTrail Security	60
Mobile Services	60
Amazon Cognito	60
Amazon Mobile Analytics	62
Applications	62
Amazon WorkSpaces	62
Amazon WorkDocs	63
Appendix – Glossary of Terms	65

Introduction

Amazon Web Services (AWS) delivers a scalable cloud computing platform with high availability and dependability, providing the tools that enable customers to run a wide range of applications. Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence. This document is intended to answer questions such as, "How does AWS help me protect my data?" Specifically, AWS physical and operational security processes are described for the network and server infrastructure under AWS's management, as well as service-specific security implementations.

Shared Security Responsibility Model

Before we go into the details of how AWS secures its resources, we should talk about how security in the cloud is slightly different than security in your on-premises data centers. When you move computer systems and data to the cloud, security responsibilities become shared between you and your cloud service provider. In this case, AWS is responsible for securing the underlying infrastructure that supports the cloud, and you're responsible for anything you put on the cloud or connect to the cloud. This shared security responsibility model can reduce your operational burden in many ways, and in some cases may even improve your default security posture without additional action on your part.



Figure 1: AWS Shared Security Responsibility Model

The amount of security configuration work you have to do varies depending on which services you select and how sensitive your data is. However, there are certain security features—such as individual user accounts and credentials, SSL/TLS for data transmissions, and user activity logging—that you should configure no matter which AWS service you use. For more information about these security features, see the "AWS Account Security Features" section below.

AWS Security Responsibilities

Amazon Web Services is responsible for protecting the global infrastructure that runs all of the services offered in the AWS cloud. This infrastructure is comprised of the hardware, software, networking, and facilities that run AWS services. Protecting this infrastructure is AWS's number one priority, and while you can't visit our data centers or offices to see this protection firsthand, we provide several reports from third-party auditors who have verified our compliance with a variety of computer security standards and regulations (for more information, visit aws.amazon.com/compliance).

Note that in addition to protecting this global infrastructure, AWS is responsible for the security configuration of its products that are considered managed services. Examples of these types of services include Amazon DynamoDB, Amazon RDS, Amazon Redshift, Amazon Elastic MapReduce, Amazon WorkSpaces, and several other services. These services provide the scalability and flexibility of cloud-based resources with the additional benefit of being managed. For these services, AWS will handle basic security tasks like guest operating system (OS) and database patching, firewall configuration, and disaster recovery. For most of these managed services, all you have to do is configure logical access controls for the resources and protect your account credentials. A few of them may require additional tasks, such as setting up database user accounts, but overall the security configuration work is performed by the service.

Customer Security Responsibilities

With the AWS cloud, you can provision virtual servers, storage, databases, and desktops in minutes instead of weeks. You can also use cloud-based analytics and workflow tools to process your data as you need it, and then store it in your own data centers or in the cloud. Which AWS services you use will determine how much configuration work you have to perform as part of your security responsibilities.

AWS products that fall into the well-understood category of Infrastructure as a Service (IaaS)—such as Amazon EC2, Amazon VPC, and Amazon S3—are completely under your control and require you to perform all of the necessary security configuration and management tasks. For example, for EC2 instances, you're responsible for management of the guest OS (including updates and security patches), any application software or utilities you install on the instances, and the configuration of the AWS-provided firewall (called a security group) on each instance. These are basically the same security tasks that you're used to performing no matter where your servers are located.

AWS managed services like Amazon RDS or Amazon Redshift provide all of the resources you need in order to perform a specific task—but without the configuration work that can come with them. With managed services, you don't have to worry about launching and maintaining instances, patching the guest OS or database, or replicating databases—AWS handles that for you. But as with all services, you should protect your AWS Account credentials and set up individual user accounts with Amazon Identity and Access Management (IAM) so that each of your users has their own credentials and you can implement segregation of duties. We also recommend using multi-factor authentication (MFA) with each account, requiring the use of SSL/TLS to communicate with your AWS resources, and setting up API/user activity logging with AWS CloudTrail. For more information about additional measures you can take, refer to the [AWS Security Best Practices whitepaper](#) and recommended reading on the [AWS Security Resources](#) webpage.

AWS Global Infrastructure Security

AWS operates the global cloud infrastructure that you use to provision a variety of basic computing resources such as processing and storage. The AWS global infrastructure includes the facilities, network, hardware, and operational software (e.g., host OS, virtualization software, etc.) that support the provisioning and use of these resources. The AWS global infrastructure is designed and managed according to security best practices as well as a variety of security compliance standards. As an AWS customer, you can be assured that you're building web architectures on top of some of the most secure computing infrastructure in the world.

AWS Compliance Program

Amazon Web Services Compliance enables customers to understand the robust controls in place at AWS to maintain security and data protection in the cloud. As systems are built on top of [AWS cloud infrastructure](#), compliance responsibilities will be [shared](#). By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, AWS [Compliance enablers](#) build on traditional programs; helping customers to establish and operate in an AWS security control environment. The IT infrastructure that AWS provides to its customers is designed and managed in alignment with security best practices and a variety of IT security standards, including:

- SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70)
- SOC 2
- SOC 3
- FISMA, DIACAP, and FedRAMP
- DOD CSM Levels 1-5
- PCI DSS Level 1
- ISO 9001 / ISO 27001
- ITAR
- FIPS 140-2
- MTCS Level 3

In addition, the flexibility and control that the AWS platform provides allows customers to deploy solutions that meet several industry-specific standards, including:

- Criminal Justice Information Services (CJIS)
- Cloud Security Alliance (CSA)
- Family Educational Rights and Privacy Act (FERPA)
- Health Insurance Portability and Accountability Act (HIPAA)
- Motion Picture Association of America (MPAA)

AWS provides a wide range of information regarding its IT control environment to customers through white papers, reports, certifications, accreditations, and other third-party attestations. More information is available in the Risk and Compliance whitepaper available on the website: <http://aws.amazon.com/compliance/>.



Physical and Environmental Security

AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches. Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to data centers by AWS employees is logged and audited routinely.

Fire Detection and Suppression

Automatic fire detection and suppression equipment has been installed to reduce risk. The fire detection system utilizes smoke detection sensors in all data center environments, mechanical and electrical infrastructure spaces, chiller rooms and generator equipment rooms. These areas are protected by either wet-pipe, double-interlocked pre-action, or gaseous sprinkler systems.

Power

The data center electrical power systems are designed to be fully redundant and maintainable without impact to operations, 24 hours a day, and seven days a week. Uninterruptible Power Supply (UPS) units provide back-up power in the event of an electrical failure for critical and essential loads in the facility. Data centers use generators to provide back-up power for the entire facility.

Climate and Temperature

Climate control is required to maintain a constant operating temperature for servers and other hardware, which prevents overheating and reduces the possibility of service outages. Data centers are conditioned to maintain atmospheric conditions at optimal levels. Personnel and systems monitor and control temperature and humidity at appropriate levels.

Management

AWS monitors electrical, mechanical, and life support systems and equipment so that any issues are immediately identified. Preventative maintenance is performed to maintain the continued operability of equipment.

Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Business Continuity Management

Amazon's infrastructure has a high level of availability and provides customers the features to deploy a resilient IT architecture. AWS has designed its systems to tolerate system or hardware failures with minimal customer impact. Data center Business Continuity Management at AWS is under the direction of the Amazon Infrastructure Group.

Availability

Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.

AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by Region). In addition to discrete uninterruptable power supply (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability zones are all redundantly connected to multiple tier-1 transit providers.

You should architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure modes, including natural disasters or system failures.

Incident Response

The Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide 24x7x365 coverage to detect incidents and to manage the impact and resolution.

Company-Wide Executive Review

Amazon's Internal Audit group has recently reviewed the AWS services resiliency plans, which are also periodically reviewed by members of the Senior Executive management team and the Audit Committee of the Board of Directors.

Communication

AWS has implemented various methods of internal communication at a global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employees; regular management meetings for updates on business performance and other matters; and electronics means such as video conferencing, electronic mail messages, and the posting of information via the Amazon intranet.

AWS has also implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "[Service Health Dashboard](#)" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact. The "AWS [Security Center](#)" is available to provide you with security and compliance details about AWS. You can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.



Network Security

The AWS network has been architected to permit you to select the level of security and resiliency appropriate for your workload. To enable you to build geographically dispersed, fault-tolerant web architectures with cloud resources, AWS has implemented a world-class network infrastructure that is carefully monitored and managed.

Secure Network Architecture

Network devices, including firewall and other boundary devices, are in place to monitor and control communications at the external boundary of the network and at key internal boundaries within the network. These boundary devices employ rule sets, access control lists (ACL), and configurations to enforce the flow of information to specific information system services.

ACLs, or traffic flow policies, are established on each managed interface, which manage and enforce the flow of traffic. ACL policies are approved by Amazon Information Security. These policies are automatically pushed using AWS's ACL-Manage tool, to help ensure these managed interfaces enforce the most up-to-date ACLs.

Secure Access Points

AWS has strategically placed a limited number of access points to the cloud to allow for a more comprehensive monitoring of inbound and outbound communications and network traffic. These customer access points are called API endpoints, and they allow secure HTTP access (HTTPS), which allows you to establish a secure communication session with your storage or compute instances within AWS. To support customers with FIPS cryptographic requirements, the SSL-terminating load balancers in AWS GovCloud (US) are FIPS 140-2-compliant.

In addition, AWS has implemented network devices that are dedicated to managing interfacing communications with Internet service providers (ISPs). AWS employs a redundant connection to more than one communication service at each Internet-facing edge of the AWS network. These connections each have dedicated network devices.

Transmission Protection

You can connect to an AWS access point via HTTP or HTTPS using Secure Sockets Layer (SSL), a cryptographic protocol that is designed to protect against eavesdropping, tampering, and message forgery.

For customers who require additional layers of network security, AWS offers the Amazon Virtual Private Cloud (VPC), which provides a private subnet within the AWS cloud, and the ability to use an IPsec Virtual Private Network (VPN) device to provide an encrypted tunnel between the Amazon VPC and your data center. For more information about VPC configuration options, refer to the [Amazon Virtual Private Cloud \(Amazon VPC\) Security](#) section below.

Amazon Corporate Segregation

Logically, the AWS Production network is segregated from the Amazon Corporate network by means of a complex set of network security / segregation devices. AWS developers and administrators on the corporate network who need to access AWS cloud components in order to maintain them must explicitly request access through the AWS ticketing system. All requests are reviewed and approved by the applicable service owner.

Approved AWS personnel then connect to the AWS network through a bastion host that restricts access to network devices and other cloud components, logging all activity for security review. Access to bastion hosts require SSH public-key authentication for all user accounts on the host. For more information on AWS developer and administrator logical access, see *AWS Access* below.



Fault-Tolerant Design

Amazon's infrastructure has a high level of availability and provides you with the capability to deploy a resilient IT architecture. AWS has designed its systems to tolerate system or hardware failures with minimal customer impact.

Data centers are built in clusters in various global *regions*. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.

AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by region). In addition to utilizing discrete uninterruptable power supply (UPS) and onsite backup generators, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability zones are all redundantly connected to multiple tier-1 transit providers.

You should architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures. However, you should be aware of location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between regions unless proactively done so by the customer, thus allowing customers with these types of data placement and privacy requirements the ability to establish compliant environments. It should be noted that all communications between regions is across public Internet infrastructure; therefore, appropriate encryption methods should be used to protect sensitive data.

As of this writing, there are eleven regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), South America (Sao Paulo), and China (Beijing).

AWS GovCloud (US) is an isolated AWS Region designed to allow US government agencies and customers to move workloads into the cloud by helping them meet certain regulatory and compliance requirements. The AWS GovCloud (US) framework allows US government agencies and their contractors to comply with U.S. International Traffic in Arms Regulations (ITAR) regulations as well as the Federal Risk and Authorization Management Program (FedRAMP) requirements. AWS GovCloud (US) has received an Agency Authorization to Operate (ATO) from the US Department of Health and Human Services (HHS) utilizing a FedRAMP accredited Third Party Assessment Organization (3PAO) for several AWS services.

The AWS GovCloud (US) Region provides the same fault-tolerant design as other regions, with two Availability Zones. In addition, the AWS GovCloud (US) region is a mandatory AWS Virtual Private Cloud (VPC) service by default to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses. More information about GovCloud is available on the AWS website: <http://aws.amazon.com/govcloud-us/>






Figure 2: Regions and Availability Zones

Note that the number of Availability Zones may change.

Network Monitoring and Protection

AWS utilizes a wide variety of automated monitoring systems to provide a high level of service performance and availability. AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools have the ability to set custom performance metrics thresholds for unusual activity.

Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Documentation is maintained to aid and inform operations personnel in handling incidents or issues. If the resolution of an issue requires collaboration, a conferencing system is used which supports communication and logging capabilities. Trained call leaders facilitate communication and progress during the handling of operational issues that require collaboration. Post-mortems are convened after any significant operational issue, regardless of external impact, and Cause of Error (COE) documents are drafted so the root cause is captured and preventative actions are taken in the future. Implementation of the preventative measures is tracked during weekly operations meetings.

AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant telecommunication providers at each region as well as additional capacity protect against the possibility of DoS attacks.

The AWS network provides significant protection against traditional network security issues, and you can implement further protection. The following are a few examples:

- **Distributed Denial Of Service (DDoS) Attacks.** AWS API endpoints are hosted on large, Internet-scale, world-class infrastructure that benefits from the same engineering expertise that has built Amazon into the world's largest online retailer. Proprietary DDoS mitigation techniques are used. Additionally, AWS's networks are multi-homed across a number of providers to achieve Internet access diversity.
- **Man in the Middle (MITM) Attacks.** All of the AWS APIs are available via SSL-protected endpoints which provide server authentication. Amazon EC2 AMIs automatically generate new SSH host certificates on first boot and log them to the instance's console. You can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. We encourage you to use SSL for all of your interactions with AWS.
- **IP Spoofing.** Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.
- **Port Scanning.** Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on our website at: <http://aws.amazon.com/contact-us/report-abuse/>. When unauthorized port scanning is detected by AWS, it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by you. Your strict management of security groups can further mitigate the threat of port scans. If you configure the security group to allow traffic from any source to a specific port, then that specific port will be vulnerable to a port scan. In these cases, you must use appropriate security measures to protect listening services that may be essential to their application from being discovered by an unauthorized port scan. For example, a web server must clearly have port 80 (HTTP) open to the world, and the administrator of this server is responsible for the security of the HTTP server software, such as Apache. You may request permission to conduct vulnerability scans as required to meet your specific compliance requirements. These scans must be limited to your own instances and must not violate the AWS Acceptable Use Policy. Advanced approval for these types of scans can be initiated by submitting a request via the website at: <https://aws-portal.amazon.com/gp/aws/html-forms-controller/contactus/AWSecurityPenTestRequest>
- **Packet sniffing by other tenants.** It is not possible for a virtual instance running in promiscuous mode to receive or "sniff" traffic that is intended for a different virtual instance. While you can place your interfaces into promiscuous mode, the hypervisor will not deliver any traffic to them that is not addressed to them. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic. Attacks such as ARP cache poisoning do not work within Amazon EC2 and Amazon VPC. While Amazon EC2 does provide ample protection against one customer inadvertently or maliciously attempting to view another's data, as a standard practice you should encrypt sensitive traffic.

In addition to monitoring, regular vulnerability scans are performed on the host operating system, web application, and databases in the AWS environment using a variety of tools. Also, AWS Security teams subscribe to newsfeeds for applicable vendor flaws and proactively monitor vendors' websites and other relevant outlets for new patches. AWS customers also have the ability to report issues to AWS via the AWS Vulnerability Reporting website at:

<http://aws.amazon.com/security/vulnerability-reporting/>

AWS Access

The AWS Production network is segregated from the Amazon Corporate network and requires a separate set of credentials for logical access. The Amazon Corporate network relies on user IDs, passwords, and Kerberos, while the AWS Production network requires SSH public-key authentication through a bastion host.

AWS developers and administrators on the Amazon Corporate network who need to access AWS cloud components must explicitly request access through the AWS access management system. All requests are reviewed and approved by the appropriate owner or manager.

Account Review and Audit

Accounts are reviewed every 90 days; explicit re-approval is required or access to the resource is automatically revoked. Access is also automatically revoked when an employee's record is terminated in Amazon's Human Resources system. Windows and UNIX accounts are disabled and Amazon's permission management system removes the user from all systems.

Requests for changes in access are captured in the Amazon permissions management tool audit log. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked.

Background Checks

AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS conducts criminal background checks, as permitted by law, as part of pre-employment screening practices for employees and commensurate with the employee's position and level of access. The policies also identify functional responsibilities for the administration of logical access and security.

Credentials Policy

AWS Security has established a credentials policy with required configurations and expiration intervals. Passwords must be complex and are forced to be changed every 90 days.

Secure Design Principles

AWS's development process follows secure software development best practices, which include formal design reviews by the AWS Security Team, threat modeling, and completion of a risk assessment. Static code analysis tools are run as a part of the standard build process, and all deployed software undergoes recurring penetration testing performed by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through launch to ongoing operations.



Change Management

Routine, emergency, and configuration changes to existing AWS infrastructure are authorized, logged, tested, approved, and documented in accordance with industry norms for similar systems. Updates to AWS's infrastructure are done to minimize any impact on the customer and their use of the services. AWS will communicate with customers, either via email, or through the AWS Service Health Dashboard (<http://status.aws.amazon.com/>) when service use is likely to be adversely affected.

Software

AWS applies a systematic approach to managing change so that changes to customer-impacting services are thoroughly reviewed, tested, approved, and well-communicated. The AWS change management process is designed to avoid unintended service disruptions and to maintain the integrity of service to the customer. Changes deployed into production environments are:

- Reviewed: Peer reviews of the technical aspects of a change are required.
- Tested: Changes being applied are tested to help ensure they will behave as expected and not adversely impact performance.
- Approved: All changes must be authorized in order to provide appropriate oversight and understanding of business impact.

Changes are typically pushed into production in a phased deployment starting with lowest impact areas. Deployments are tested on a single system and closely monitored so impacts can be evaluated. Service owners have a number of configurable metrics that measure the health of the service's upstream dependencies. These metrics are closely monitored with thresholds and alarming in place. Rollback procedures are documented in the Change Management (CM) ticket.

When possible, changes are scheduled during regular change windows. Emergency changes to production systems that require deviations from standard change management procedures are associated with an incident and are logged and approved as appropriate.

Periodically, AWS performs self-audits of changes to key services to monitor quality, maintain high standards, and facilitate continuous improvement of the change management process. Any exceptions are analyzed to determine the root cause, and appropriate actions are taken to bring the change into compliance or roll back the change if necessary. Actions are then taken to address and remediate the process or people issue.

Infrastructure

Amazon's Corporate Applications team develops and manages software to automate IT processes for UNIX/Linux hosts in the areas of third-party software delivery, internally developed software, and configuration management. The Infrastructure team maintains and operates a UNIX/Linux configuration management framework to address hardware scalability, availability, auditing, and security management. By centrally managing hosts through the use of automated processes that manage change, Amazon is able to achieve its goals of high availability, repeatability, scalability, security, and disaster recovery. Systems and network engineers monitor the status of these automated tools on a continuous basis, reviewing reports to respond to hosts that fail to obtain or update their configuration and software.



Internally developed configuration management software is installed when new hardware is provisioned. These tools are run on all UNIX hosts to validate that they are configured and that software is installed in compliance with standards determined by the role assigned to the host. This configuration management software also helps to regularly update packages that are already installed on the host. Only approved personnel enabled through the permissions service may log in to the central configuration management servers.

AWS Account Security Features

AWS provides a variety of tools and features that you can use to keep your AWS Account and resources safe from unauthorized use. This includes credentials for access control, HTTPS endpoints for encrypted data transmission, the creation of separate IAM user accounts, user activity logging for security monitoring, and Trusted Advisor security checks. You can take advantage of all of these security tools no matter which AWS services you select.

AWS Credentials

To help ensure that only authorized users and processes access your AWS Account and resources, AWS uses several types of credentials for authentication. These include passwords, cryptographic keys, digital signatures, and certificates. We also provide the option of requiring multi-factor authentication (MFA) to log into your AWS Account or IAM user accounts. The following table highlights the various AWS credentials and their uses.

Credential Type	Use	Description
Passwords	AWS root account or IAM user account login to the AWS Management Console	A string of characters used to log into your AWS account or IAM account. AWS passwords must be a minimum of 6 characters and may be up to 128 characters.
Multi-Factor Authentication (MFA)	AWS root account or IAM user account login to the AWS Management Console	A six-digit single-use code that is required in addition to your password to log in to your AWS Account or IAM user account.
Access Keys	Digitally signed requests to AWS APIs (using the AWS SDK, CLI, or REST/Query APIs)	Includes an access key ID and a secret access key. You use access keys to digitally sign programmatic requests that you make to AWS.
Key Pairs	- SSH login to EC2 instances - CloudFront signed URLs	A key pair is required to connect to an EC2 instance launched from a public AMI. The keys that Amazon EC2 uses are 1024-bit SSH-2 RSA keys. You can have a key pair generated automatically for you when you launch the instance or you can upload your own.
X.509 Certificates	- Digitally signed SOAP requests to AWS APIs - SSL server certificates for HTTPS	X.509 certificates are only used to sign SOAP-based requests (currently used only with Amazon S3). You can have AWS create an X.509 certificate and private key that you can download, or you can upload your own certificate by using the Security Credentials page.

You can download a Credential Report for your account at any time from the Security Credentials page. This report lists all of your account's users and the status of their credentials—whether they use a password, whether their password

expires and must be changed regularly, the last time they changed their password, the last time they rotated their access keys, and whether they have MFA enabled.

For security reasons, if your credentials have been lost or forgotten, you cannot recover them or re-download them. However, you can create new credentials and then disable or delete the old set of credentials.

In fact, AWS recommends that you change (rotate) your access keys and certificates on a regular basis. To help you do this without potential impact to your application's availability, AWS supports multiple concurrent access keys and certificates. With this feature, you can rotate keys and certificates into and out of operation on a regular basis without any downtime to your application. This can help to mitigate risk from lost or compromised access keys or certificates. The AWS IAM API enables you to rotate the access keys of your AWS Account as well as for IAM user accounts.

Passwords

Passwords are required to access your AWS Account, individual IAM user accounts, AWS Discussion Forums, and the AWS Support Center. You specify the password when you first create the account, and you can change it at any time by going to the Security Credentials page. AWS passwords can be up to 128 characters long and contain special characters, so we encourage you to create a strong password that cannot be easily guessed.

You can set a password policy for your IAM user accounts to ensure that strong passwords are used and that they are changed often. A password policy is a set of rules that define the type of password an IAM user can set. For more information about password policies, go to [Managing Passwords](#) in Using IAM.

AWS Multi-Factor Authentication (AWS MFA)

AWS Multi-Factor Authentication (AWS MFA) is an additional layer of security for accessing AWS services. When you enable this optional feature, you will need to provide a six-digit single-use code in addition to your standard user name and password credentials before access is granted to your AWS Account settings or AWS services and resources. You get this single-use code from an authentication device that you keep in your physical possession. This is called multi-factor authentication because more than one authentication factor is checked before access is granted: a password (something you know) and the precise code from your authentication device (something you have). You can enable MFA devices for your AWS Account as well as for the users you have created under your AWS Account with AWS IAM. In addition, you add MFA protection for access across AWS Accounts, for when you want to allow a user you've created under one AWS Account to use an IAM role to access resources under another AWS Account. You can require the user to use MFA before assuming the role as an additional layer of security.

AWS MFA supports the use of both hardware tokens and virtual MFA devices. Virtual MFA devices use the same protocols as the physical MFA devices, but can run on any mobile hardware device, including a smartphone. A virtual MFA device uses a software application that generates six-digit authentication codes that are compatible with the Time-Based One-Time Password (TOTP) standard, as described in [RFC 6238](#). Most virtual MFA applications allow you to host more than one virtual MFA device, which makes them more convenient than hardware MFA devices. However, you should be aware that because a virtual MFA might be run on a less secure device such as a smartphone, a virtual MFA might not provide the same level of security as a hardware MFA device.

You can also enforce MFA authentication for AWS service APIs in order to provide an extra layer of protection over powerful or privileged actions such as terminating Amazon EC2 instances or reading sensitive data stored in Amazon S3. You do this by adding an MFA-authentication requirement to an IAM access policy. You can attach these access policies to IAM users, IAM groups, or resources that support Access Control Lists (ACLs) like Amazon S3 buckets, SQS queues, and SNS topics.



It is easy to obtain hardware tokens from a participating third-party provider or virtual MFA applications from an AppStore and to set it up for use via the AWS website. More information about AWS MFA is available on the AWS website: <http://aws.amazon.com/mfa/>

Access Keys

AWS requires that all API requests be signed—that is, they must include a digital signature that AWS can use to verify the identity of the requestor. You calculate the digital signature using a cryptographic hash function. The input to the hash function in this case includes the text of your request and your secret access key. If you use any of the AWS SDKs to generate requests, the digital signature calculation is done for you; otherwise, you can have your application calculate it and include it in your REST or Query requests by following the directions [in our documentation](#).

Not only does the signing process help protect message integrity by preventing tampering with the request while it is in transit, it also helps protect against potential replay attacks. A request must reach AWS within 15 minutes of the time stamp in the request. Otherwise, AWS denies the request.

The most recent version of the digital signature calculation process is Signature Version 4, which calculates the signature using the HMAC-SHA256 protocol. Version 4 provides an additional measure of protection over previous versions by requiring that you sign the message using a key that is derived from your secret access key rather than using the secret access key itself. In addition, you derive the signing key based on *credential scope*, which facilitates cryptographic isolation of the signing key.

Because access keys can be misused if they fall into the wrong hands, we encourage you to save them in a safe place and not embed them in your code. For customers with large fleets of elastically scaling EC2 instances, the use of IAM roles can be a more secure and convenient way to manage the distribution of access keys. IAM roles provide temporary credentials, which not only get automatically loaded to the target instance, but are also automatically rotated multiple times a day.

Key Pairs

Amazon EC2 instances created from a public AMI use a public/private key pair rather than a password for signing in via Secure Shell (SSH). The public key is embedded in your instance, and you use the private key to sign in securely without a password. After you create your own AMIs, you can choose other mechanisms to securely log in to your new instances.

You can have a key pair generated automatically for you when you launch the instance or you can upload your own. Save the private key in a safe place on your system, and record the location where you saved it.

For Amazon CloudFront, you use key pairs to create signed URLs for private content, such as when you want to distribute restricted content that someone paid for. You create Amazon CloudFront key pairs by using the Security Credentials page. CloudFront key pairs can be created only by the root account and cannot be created by IAM users.

X.509 Certificates

X.509 certificates are used to sign SOAP-based requests. X.509 certificates contain a public key and additional metadata (like an expiration date that AWS verifies when you upload the certificate), and is associated with a private key. When you create a request, you create a digital signature with your private key and then include that signature in the request, along with your certificate. AWS verifies that you're the sender by decrypting the signature with the public key that is in your certificate. AWS also verifies that the certificate you sent matches the certificate that you uploaded to AWS.

For your AWS Account, you can have AWS create an X.509 certificate and private key that you can download, or you can upload your own certificate by using the Security Credentials page. For IAM users, you must create the X.509 certificate (signing certificate) by using third-party software. In contrast with root account credentials, AWS cannot create an X.509 certificate for IAM users. After you create the certificate, you attach it to an IAM user by using IAM.

In addition to SOAP requests, X.509 certificates are used as SSL/TLS server certificates for customers who want to use HTTPS to encrypt their transmissions. To use them for HTTPS, you can use an open-source tool like OpenSSL to create a unique private key. You'll need the private key to create the Certificate Signing Request (CSR) that you submit to a certificate authority (CA) to obtain the server certificate. You'll then use the AWS CLI to upload the certificate, private key, and certificate chain to IAM.

You'll also need an X.509 certificate to create a customized Linux AMI for EC2 instances. The certificate is only required to create an instance-backed AMI (as opposed to an EBS-backed AMI). You can have AWS create an X.509 certificate and private key that you can download, or you can upload your own certificate by using the Security Credentials page.

Individual User Accounts

AWS provides a centralized mechanism called AWS Identity and Access Management (IAM) for creating and managing individual users within your AWS Account. A user can be any individual, system, or application that interacts with AWS resources, either programmatically or through the AWS Management Console or AWS Command Line Interface (CLI). Each user has a unique name within the AWS Account, and a unique set of security credentials not shared with other users. AWS IAM eliminates the need to share passwords or keys, and enables you to minimize the use of your AWS Account credentials.

With IAM, you define policies that control which AWS services your users can access and what they can do with them. You can grant users only the minimum permissions they need to perform their jobs. See the AWS Identity and Access Management (AWS IAM) section below for more information.

Secure HTTPS Access Points

For greater communication security when accessing AWS resources, you should use HTTPS instead of HTTP for data transmissions. HTTPS uses the SSL/TLS protocol, which uses public-key cryptography to prevent eavesdropping, tampering, and forgery. All AWS services provide secure customer access points (also called API endpoints) that allow you to establish secure HTTPS communication sessions.

Several services also now offer more advanced cipher suites that use the Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) protocol. ECDHE allows SSL/TLS clients to provide Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This helps prevent the decoding of captured data by unauthorized third parties, even if the secret long-term key itself is compromised.

Security Logs

As important as credentials and encrypted endpoints are for preventing security problems, logs are just as crucial for understanding events after a problem has occurred. And to be effective as a security tool, a log must include not just a list of what happened and when, but also identify the source. To help you with your after-the-fact investigations and near-realtime intrusion detection, AWS CloudTrail provides a log of all requests for AWS resources within your account. For each event, you can see what service was accessed, what action was performed, and who made the request. CloudTrail captures information about every API call to every AWS resource you use, including sign-in events.



Once you have enabled CloudTrail, event logs are delivered every 5 minutes. You can configure CloudTrail so that it aggregates log files from multiple regions into a single Amazon S3 bucket. From there, you can then upload them to your favorite log management and analysis solutions to perform security analysis and detect user behavior patterns. By default, log files are stored securely in Amazon S3, but you can also archive them to Amazon Glacier to help meet audit and compliance requirements.

In addition to CloudTrail's user activity logs, you can use the Amazon CloudWatch Logs feature to collect and monitor system, application, and custom log files from your EC2 instances and other sources in near-real time. For example, you can monitor your web server's log files for invalid user messages to detect unauthorized login attempts to your guest OS.

AWS Trusted Advisor Security Checks

The AWS Trusted Advisor customer support service not only monitors for cloud performance and resiliency, but also cloud security. Trusted Advisor inspects your AWS environment and makes recommendations when opportunities may exist to save money, improve system performance, or close security gaps. It provides alerts on several of the most common security misconfigurations that can occur, including leaving certain ports open that make you vulnerable to hacking and unauthorized access, neglecting to create IAM accounts for your internal users, allowing public access to Amazon S3 buckets, not turning on user activity logging (AWS CloudTrail), or not using MFA on your root AWS Account. You also have the option for a Security contact at your organization to automatically receive a weekly email with an updated status of your Trusted Advisor security checks.

The AWS Trusted Advisor service provides four checks at no additional charge to all users, including three important security checks: specific ports unrestricted, IAM use, and MFA on root account. And when you sign up for Business- or Enterprise-level AWS Support, you receive full access to all Trusted Advisor checks.

AWS Service-Specific Security

Not only is security built into every layer of the AWS infrastructure, but also into each of the services available on that infrastructure. AWS services are architected to work efficiently and securely with all AWS networks and platforms. Each service provides extensive security features to enable you to protect sensitive data and applications.

Compute Services

Amazon Web Services provides a variety of cloud-based computing services that include a wide selection of compute instances that can scale up and down automatically to meet the needs of your application or enterprise.

Amazon Elastic Compute Cloud (Amazon EC2) Security

Amazon Elastic Compute Cloud (EC2) is a key component in Amazon's Infrastructure as a Service (IaaS), providing resizable computing capacity using server instances in AWS's data centers. Amazon EC2 is designed to make web-scale computing easier by enabling you to obtain and configure capacity with minimal friction. You create and launch *instances*, which are collections of platform hardware and software.

Multiple Levels of Security

Security within Amazon EC2 is provided on multiple levels: the operating system (OS) of the host platform, the virtual instance OS or guest OS, a firewall, and signed API calls. Each of these items builds on the capabilities of the others. The goal is to prevent data contained within Amazon EC2 from being intercepted by unauthorized systems or users and to



provide Amazon EC2 instances themselves that are as secure as possible without sacrificing the flexibility in configuration that customers demand.

The Hypervisor

Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor, taking advantage of paravirtualization (in the case of Linux guests). Because paravirtualized guests rely on the hypervisor to provide support for operations that normally require privileged access, the guest OS has no elevated access to the CPU. The CPU provides four separate privilege modes: 0-3, called *rings*. Ring 0 is the most privileged and 3 the least. The host OS executes in Ring 0. However, rather than executing in Ring 0 as most operating systems do, the guest OS runs in a lesser-privileged Ring 1 and applications in the least privileged Ring 3. This explicit virtualization of the physical resources leads to a clear separation between guest and hypervisor, resulting in additional security separation between the two.

Instance Isolation

Different instances running on the same physical machine are isolated from each other via the Xen hypervisor. Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.

Customer instances have no access to raw disk devices, but instead are presented with virtualized disks. The AWS proprietary disk virtualization layer automatically resets every block of storage used by the customer, so that one customer's data is never unintentionally exposed to another. In addition, memory allocated to guests is scrubbed (set to zero) by the hypervisor when it is unallocated to a guest. The memory is not returned to the pool of free memory available for new allocations until the memory scrubbing is complete.

AWS recommends customers further protect their data using appropriate means. One common solution is to run an encrypted file system on top of the virtualized disk device.




Figure 3: Amazon EC2 Multiple Layers of Security

Host Operating System: Administrators with a business need to access the management plane are required to use multi-factor authentication to gain access to purpose-built administration hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane of the cloud. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems can be revoked.

Guest Operating System: Virtual instances are completely controlled by you, the customer. You have full root access or administrative control over accounts, services, and applications. AWS does not have any access rights to your instances or the guest OS. AWS recommends a base set of security best practices to include disabling password-only access to your guests, and utilizing some form of multi-factor authentication to gain access to your instances (or at a minimum certificate-based SSH Version 2 access). Additionally, you should employ a privilege escalation mechanism with logging on a per-user basis. For example, if the guest OS is Linux, after hardening your instance you should utilize certificate-based SSHv2 to access the virtual instance, disable remote root login, use command-line logging, and use 'sudo' for privilege escalation. You should generate your own key pairs in order to guarantee that they are unique, and not shared with other customers or with AWS.

AWS also supports the use of the Secure Shell (SSH) network protocol to enable you to log in securely to your UNIX/Linux EC2 instances. Authentication for SSH used with AWS is via a public/private key pair to reduce the risk of unauthorized access to your instance. You can also connect remotely to your Windows instances using Remote Desktop Protocol (RDP) by utilizing an RDP certificate generated for your instance.

You also control the updating and patching of your guest OS, including security updates. Amazon-provided Windows and Linux-based AMIs are updated regularly with the latest patches, so if you do not need to preserve data or customizations on your running Amazon AMI instances, you can simply relaunch new instances with the latest updated AMI. In addition, updates are provided for the Amazon Linux AMI via the Amazon Linux yum repositories.

Firewall: Amazon EC2 provides a complete firewall solution; this mandatory inbound firewall is configured in a default deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic. The traffic may be restricted by protocol, by service port, as well as by source IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

The firewall can be configured in groups permitting different classes of instances to have different rules. Consider, for example, the case of a traditional three-tiered web application. The group for the web servers would have port 80 (HTTP) and/or port 443 (HTTPS) open to the Internet. The group for the application servers would have port 8000 (application specific) accessible only to the web server group. The group for the database servers would have port 3306 (MySQL) open only to the application server group. All three groups would permit administrative access on port 22 (SSH), but only from the customer's corporate network. Highly secure applications can be deployed using this expressive mechanism. See diagram below:




Figure 4: Amazon EC2 Security Group Firewall

The firewall isn't controlled through the guest OS; rather it requires your X.509 certificate and key to authorize changes, thus adding an extra layer of security. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling you to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports you open, and for what duration and purpose. The default state is to deny all incoming traffic, and you should plan carefully what you will open when building and securing your applications. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages you to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall and VPNs. This can restrict both inbound and outbound traffic.

API Access: API calls to launch and terminate instances, change firewall parameters, and perform other functions are all signed by your Amazon Secret Access Key, which could be either the AWS Accounts Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to your Secret Access Key, Amazon EC2 API calls cannot be

made on your behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints.

Permissions: AWS IAM also enables you to further control what APIs a user has permissions to call.

Elastic Block Storage (Amazon EBS) Security

Amazon Elastic Block Storage (EBS) allows you to create storage volumes from 1 GB to 16 TB that can be mounted as devices by Amazon EC2 instances. Storage volumes behave like raw, unformatted block devices, with user supplied device names and a block device interface. You can create a file system on top of Amazon EBS volumes, or use them in any other way you would use a block device (like a hard drive). Amazon EBS volume access is restricted to the AWS Account that created the volume, and to the users under the AWS Account created with AWS IAM if the user has been granted access to the EBS operations, thus denying all other AWS Accounts and users the permission to view or access the volume.

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability. For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed. AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

You can make Amazon EBS volume snapshots publicly available to other AWS Accounts to use as the basis for creating your own volumes. Sharing Amazon EBS volume snapshots does not provide other AWS Accounts with the permission to alter or delete the original snapshot, as that right is explicitly reserved for the AWS Account that created the volume. An EBS snapshot is a block-level view of an entire EBS volume. Note that data that is not visible through the file system on the volume, such as files that have been deleted, may be present in the EBS snapshot. If you want to create shared snapshots, you should do so carefully. If a volume has held sensitive data or has had files deleted from it, a new EBS volume should be created. The data to be contained in the shared snapshot should be copied to the new volume, and the snapshot created from the new volume.

Amazon EBS volumes are presented to you as raw unformatted block devices that have been wiped prior to being made available for use. Wiping occurs immediately before reuse so that you can be assured that the wipe process completed. If you have procedures requiring that all data be wiped via a specific method, such as those detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization"), you have the ability to do so on Amazon EBS. You should conduct a specialized wipe procedure prior to deleting the volume for compliance with your established requirements.

Encryption of sensitive data is generally a good security practice, and AWS provides the ability to encrypt EBS volumes and their snapshots with AES-256. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. In order to be able to do this efficiently and with low latency, the EBS encryption feature is only available on EC2's more powerful instance types (e.g., M3, C3, R3, G2).

Auto Scaling Security

Auto Scaling allows you to automatically scale your Amazon EC2 capacity up or down according to conditions you define, so that the number of Amazon EC2 instances you are using scales up seamlessly during demand spikes to maintain performance, and scales down automatically during demand lulls to minimize costs.



Like all AWS services, Auto Scaling requires that every request made to its control API be authenticated so only authenticated users can access and manage Auto Scaling. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. However, getting credentials out to new EC2 instances launched with Auto Scaling can be challenging for large or elastically scaling fleets. To simplify this process, you can use *roles* within IAM, so that any new instances launched with a role will be given credentials automatically. When you launch an EC2 instance with an IAM role, temporary AWS security credentials with permissions specified by the role will be securely provisioned to the instance and will be made available to your application via the Amazon EC2 Instance Metadata Service. The Metadata Service will make new temporary security credentials available prior to the expiration of the current active credentials, so that valid credentials are always available on the instance. In addition, the temporary security credentials are automatically rotated multiple times per day, providing enhanced security. You can further control access to Auto Scaling by creating users under your AWS Account using AWS IAM, and controlling what Auto Scaling APIs these users have permission to call. More information about using roles when launching instances is available in the Amazon EC2 User Guide on the AWS website: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM>

Networking Services

Amazon Web Services provides a range of networking services that enable you to create a logically isolated network that you define, establish a private network connection to the AWS cloud, use a highly available and scalable DNS service and deliver content to your end users with low latency at high data transfer speeds with a content delivery web service.

Amazon Elastic Load Balancing Security

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

- Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer
- Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network
- When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options
- Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance.

HTTPS/TLS uses a long-term secret key to generate a short-term session key to be used between the server and the browser to create the ciphered (encrypted) message. Amazon Elastic Load Balancing configures your load balancer with a pre-defined cipher set that is used for TLS negotiation when a connection is established between a client and your load balancer. The pre-defined cipher set provides compatibility with a broad range of clients and uses strong cryptographic algorithms. However, some customers may have requirements for allowing only specific ciphers and protocols (such as PCI, SOX, etc.) from clients to ensure that standards are met. In these cases, Amazon Elastic Load Balancing provides options for selecting different configurations for TLS protocols and ciphers. You can choose to enable or disable the ciphers depending on your specific requirements.

To help ensure the use of newer and stronger cipher suites when establishing a secure connection, you can configure the load balancer to have the final say in the cipher suite selection during the client-server negotiation. When the Server Order Preference option is selected, the load balancer will select a cipher suite based on the server's prioritization of cipher suites rather than the client's. This gives you more control over the level of security that clients use to connect to your load balancer.

For even greater communication privacy, Amazon Elastic Load Balancer allows the use of Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This prevents the decoding of captured data, even if the secret long-term key itself is compromised.

Amazon Elastic Load Balancing allows you to identify the originating IP address of a client connecting to your servers, whether you're using HTTPS or TCP load balancing. Typically, client connection information, such as IP address and port, is lost when requests are proxied through a load balancer. This is because the load balancer sends requests to the server on behalf of the client, making your load balancer appear as though it is the requesting client. Having the originating client IP address is useful if you need more information about visitors to your applications in order to gather connection statistics, analyze traffic logs, or manage whitelists of IP addresses.

Amazon Elastic Load Balancing access logs contain information about each HTTP and TCP request processed by your load balancer. This includes the IP address and port of the requesting client, the backend IP address of the instance that processed the request, the size of the request and response, and the actual request line from the client (for example, GET http://www.example.com: 80/HTTP/1.1). All requests sent to the load balancer are logged, including requests that never made it to back-end instances.

Amazon Virtual Private Cloud (Amazon VPC) Security

Normally, each Amazon EC2 instance you launch is randomly assigned a public IP address in the Amazon EC2 address space. Amazon VPC enables you to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses in the range of your choice (e.g., 10.0.0.0/16). You can define subnets within your VPC, grouping similar kinds of instances based on IP address range, and then set up routing and security to control the flow of traffic in and out of the instances and subnets.

AWS offers a variety of VPC architecture templates with configurations that provide varying levels of public access:

- **VPC with a single public subnet only.** Your instances run in a private, isolated section of the AWS cloud with direct access to the Internet. Network ACLs and security groups can be used to provide strict control over inbound and outbound network traffic to your instances.
- **VPC with public and private subnets.** In addition to containing a public subnet, this configuration adds a private subnet whose instances are not addressable from the Internet. Instances in the private subnet can establish outbound connections to the Internet via the public subnet using Network Address Translation (NAT).
- **VPC with public and private subnets and hardware VPN access.** This configuration adds an IPsec VPN connection between your Amazon VPC and your data center, effectively extending your data center to the cloud while also providing direct access to the Internet for public subnet instances in your Amazon VPC. In this configuration, customers add a VPN appliance on their corporate data center side.
- **VPC with private subnet only and hardware VPN access.** Your instances run in a private, isolated section of the AWS cloud with a private subnet whose instances are not addressable from the Internet. You can connect this

private subnet to your corporate data center via an IPsec VPN tunnel.

You can also connect two VPCs using a private IP address, which allows instances in the two VPCs to communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region.

Security features within Amazon VPC include security groups, network ACLs, routing tables, and external gateways. Each of these items is complementary to providing a secure, isolated network that can be extended through selective enabling of direct Internet access or private connectivity to another network. Amazon EC2 instances running within an Amazon VPC inherit all of the benefits described below related to the guest OS and protection against packet sniffing. Note, however, that you must create VPC security groups specifically for your Amazon VPC; any Amazon EC2 security groups you have created will not work inside your Amazon VPC. Also, Amazon VPC security groups have additional capabilities that Amazon EC2 security groups do not have, such as being able to change the security group after the instance is launched and being able to specify any protocol with a standard protocol number (as opposed to just TCP, UDP, or ICMP).

Each Amazon VPC is a distinct, isolated network within the cloud; network traffic within each Amazon VPC is isolated from all other Amazon VPCs. At creation time, you select an IP address range for each Amazon VPC. You may create and attach an Internet gateway, virtual private gateway, or both to establish external connectivity, subject to the controls below.

API Access: Calls to create and delete Amazon VPCs, change routing, security group, and network ACL parameters, and perform other functions are all signed by your Amazon Secret Access Key, which could be either the AWS Account's Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to your Secret Access Key, Amazon VPC API calls cannot be made on your behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints. AWS IAM also enables a customer to further control what APIs a newly created user has permissions to call.

Subnets and Route Tables: You create one or more subnets within each Amazon VPC; each instance launched in the Amazon VPC is connected to one subnet. Traditional Layer 2 security attacks, including MAC spoofing and ARP spoofing, are blocked.

Each subnet in an Amazon VPC is associated with a routing table, and all network traffic leaving the subnet is processed by the routing table to determine the destination.

Firewall (Security Groups): Like Amazon EC2, Amazon VPC supports a complete firewall solution enabling filtering on both ingress and egress traffic from an instance. The default group enables inbound communication from other members of the same group and outbound communication to any destination. Traffic can be restricted by any IP protocol, by service port, as well as source/destination IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

The firewall isn't controlled through the guest OS; rather, it can be modified only through the invocation of Amazon VPC APIs. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling you to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports you open, and for what duration and purpose. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages you to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall.






Figure 5: Amazon VPC Network Architecture

Network Access Control Lists: To add a further layer of security within Amazon VPC, you can configure network ACLs. These are stateless traffic filters that apply to all traffic inbound or outbound from a subnet within Amazon VPC. These ACLs can contain ordered rules to allow or deny traffic based upon IP protocol, by service port, as well as source/destination IP address.

Like security groups, network ACLs are managed through Amazon VPC APIs, adding an additional layer of protection and enabling additional security through separation of duties. The diagram below depicts how the security controls above inter-relate to enable flexible network topologies while providing complete control over network traffic flows.




Figure 6: Flexible Network Topologies

Virtual Private Gateway: A virtual private gateway enables private connectivity between the Amazon VPC and another network. Network traffic within each virtual private gateway is isolated from network traffic within all other virtual private gateways. You can establish VPN connections to the virtual private gateway from gateway devices at your premises. Each connection is secured by a pre-shared key in conjunction with the IP address of the customer gateway device.

Internet Gateway: An Internet gateway may be attached to an Amazon VPC to enable direct connectivity to Amazon S3, other AWS services, and the Internet. Each instance desiring this access must either have an Elastic IP associated with it or route traffic through a NAT instance. Additionally, network routes are configured (see above) to direct traffic to the Internet gateway. AWS provides reference NAT AMIs that you can extend to perform network logging, deep packet inspection, application-layer filtering, or other security controls.

This access can only be modified through the invocation of Amazon VPC APIs. AWS supports the ability to grant granular access to different administrative functions on the instances and the Internet gateway, therefore enabling you to implement additional security through separation of duties.

Dedicated Instances: Within a VPC, you can launch Amazon EC2 instances that are physically isolated at the host hardware level (i.e., they will run on single-tenant hardware). An Amazon VPC can be created with ‘dedicated’ tenancy, so that all instances launched into the Amazon VPC will utilize this feature. Alternatively, an Amazon VPC may be created with ‘default’ tenancy, but you can specify dedicated tenancy for particular instances launched into it.

Elastic Network Interfaces: Each Amazon EC2 instance has a default network interface that is assigned a private IP address on your Amazon VPC network. You can create and attach an additional network interface, known as an elastic network interface (ENI), to any Amazon EC2 instance in your Amazon VPC for a total of two network interfaces per instance. Attaching more than one network interface to an instance is useful when you want to create a management network, use network and security appliances in your Amazon VPC, or create dual-homed instances with workloads/roles on distinct subnets. An ENI's attributes, including the private IP address, elastic IP addresses, and MAC address, will follow the ENI as it is attached or detached from an instance and reattached to another instance. More information about Amazon VPC is available on the AWS website: <http://aws.amazon.com/vpc/>

Additional Network Access Control with EC2-VPC

If you launch instances in a region where you did not have instances before AWS launched the new EC2-VPC feature (also called Default VPC), all instances are automatically provisioned in a ready-to-use default VPC. You can choose to create additional VPCs, or you can create VPCs for instances in regions where you already had instances before we launched EC2-VPC.

If you create a VPC later, using regular VPC, you specify a CIDR block, create subnets, enter the routing and security for those subnets, and provision an Internet gateway or NAT instance if you want one of your subnets to be able to reach the Internet. When you launch EC2 instances into an EC2-VPC, most of this work is automatically performed for you. When you launch an instance into a default VPC using EC2-VPC, we do the following to set it up for you:

- Create a default subnet in each Availability Zone
- Create an Internet gateway and connect it to your default VPC
- Create a main route table for your default VPC with a rule that sends all traffic destined for the Internet to the Internet gateway
- Create a default security group and associate it with your default VPC
- Create a default network access control list (ACL) and associate it with your default VPC
- Associate the default DHCP options set for your AWS account with your default VPC

In addition to the default VPC having its own private IP range, EC2 instances launched in a default VPC can also receive a public IP.

The following table summarizes the differences between instances launched into EC2-Classic, instances launched into a default VPC, and instances launched into a non-default VPC.

Characteristic	EC2-Classic	EC2-VPC (Default VPC)	Regular VPC
Public IP address	Your instance receives a public IP address.	Your instance launched in a default subnet receives a public	Your instance doesn't receive a public IP address by default,

Characteristic	EC2-Classic	EC2-VPC (Default VPC)	Regular VPC
		IP address by default, unless you specify otherwise during launch.	unless you specify otherwise during launch.
Private IP address	Your instance receives a private IP address from the EC2-Classic range each time it's started.	Your instance receives a static private IP address from the address range of your default VPC.	Your instance receives a static private IP address from the address range of your VPC.
Multiple private IP addresses	We select a single IP address for your instance. Multiple IP addresses are not supported.	You can assign multiple private IP addresses to your instance.	You can assign multiple private IP addresses to your instance.
Elastic IP address	An EIP is disassociated from your instance when you stop it.	An EIP remains associated with your instance when you stop it.	An EIP remains associated with your instance when you stop it.
DNS hostnames	DNS hostnames are enabled by default.	DNS hostnames are enabled by default.	DNS hostnames are disabled by default.
Security group	A security group can reference security groups that belong to other AWS accounts.	A security group can reference security groups for your VPC only.	A security group can reference security groups for your VPC only.
Security group association	You must terminate your instance to change its security group.	You can change the security group of your running instance.	You can change the security group of your running instance.
Security group rules	You can add rules for inbound traffic only.	You can add rules for inbound and outbound traffic.	You can add rules for inbound and outbound traffic.
Tenancy	Your instance runs on shared hardware; you cannot run an instance on single-tenant hardware.	You can run your instance on shared hardware or single-tenant hardware.	You can run your instance on shared hardware or single-tenant hardware.

Note that security groups for instances in EC2-Classic are slightly different than security groups for instances in EC2-VPC. For example, you can add rules for inbound traffic for EC2-Classic, but you can add rules for both inbound and outbound traffic to EC2-VPC. In EC2-Classic, you can't change the security groups assigned to an instance after it's launched, but in EC2-VPC, you can change security groups assigned to an instance after it's launched. In addition, you can't use the security groups that you've created for use with EC2-Classic with instances in your VPC. You must create security groups specifically for use with instances in your VPC. The rules you create for use with a security group for a VPC can't reference a security group for EC2-Classic, and vice versa.

Amazon Route 53 Security

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) service that answers DNS queries, translating domain names into IP addresses so computers can communicate with each other. Route 53 can be used to connect user requests to infrastructure running in AWS – such as an Amazon EC2 instance or an Amazon S3 bucket – or to infrastructure outside of AWS.

Amazon Route 53 lets you manage the IP addresses (records) listed for your domain names and it answers requests (queries) to translate specific domain names into their corresponding IP addresses. Queries for your domain are automatically routed to a nearby DNS server using anycast in order to provide the lowest latency possible. Route 53



makes it possible for you to manage traffic globally through a variety of routing types, including Latency Based Routing (LBR), Geo DNS, and Weighted Round-Robin (WRR) —all of which can be combined with DNS Failover in order to help create a variety of low-latency, fault-tolerant architectures. The failover algorithms implemented by Amazon Route 53 are designed not only to route traffic to endpoints that are healthy, but also to help avoid making disaster scenarios worse due to misconfigured health checks and applications, endpoint overloads, and partition failures.

Route 53 also offers Domain Name Registration – you can purchase and manage domain names such as example.com and Route 53 will automatically configure default DNS settings for your domains. You can buy, manage, and transfer (both in and out) domains from a wide selection of generic and country-specific top-level domains (TLDs). During the registration process, you have the option to enable privacy protection for your domain. This option will hide most of your personal information from the public Whois database in order to help thwart scraping and spamming.

Amazon Route 53 is built using AWS's highly available and reliable infrastructure. The distributed nature of the AWS DNS servers helps ensure a consistent ability to route your end users to your application. Route 53 also helps ensure the availability of your website by providing health checks and DNS failover capabilities. You can easily configure Route 53 to check the health of your website on a regular basis (even secure web sites that are available only over SSL), and to switch to a backup site if the primary one is unresponsive.

Like all AWS Services, Amazon Route 53 requires that every request made to its control API be authenticated so only authenticated users can access and manage Route 53. API requests are signed with an HMAC-SHA1 or HMAC-SHA256 signature calculated from the request and the user's AWS Secret Access key. Additionally, the Amazon Route 53 control API is only accessible via SSL-encrypted endpoints. It supports both IPv4 and IPv6 routing.

You can control access to Amazon Route 53 DNS management functions by creating users under your AWS Account using AWS IAM, and controlling which Route 53 operations these users have permission to perform.

Amazon CloudFront Security

Amazon CloudFront gives customers an easy way to distribute content to end users with low latency and high data transfer speeds. It delivers dynamic, static, and streaming content using a global network of edge locations. Requests for customers' objects are automatically routed to the nearest edge location, so content is delivered with the best possible performance. Amazon CloudFront is optimized to work with other AWS services, like Amazon S3, Amazon EC2, Amazon Elastic Load Balancing, and Amazon Route 53. It also works seamlessly with any non-AWS origin server that stores the original, definitive versions of your files.

Amazon CloudFront requires every request made to its control API be authenticated so only authorized users can create, modify, or delete their own Amazon CloudFront distributions. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Additionally, the Amazon CloudFront control API is only accessible via SSL-enabled endpoints.

There is no guarantee of durability of data held in Amazon CloudFront edge locations. The service may from time to time remove objects from edge locations if those objects are not requested frequently. Durability is provided by Amazon S3, which works as the origin server for Amazon CloudFront holding the original, definitive copies of objects delivered by Amazon CloudFront.

If you want control over who is able to download content from Amazon CloudFront, you can enable the service's private content feature. This feature has two components: the first controls how content is delivered from the Amazon CloudFront edge location to viewers on the Internet. The second controls how the Amazon CloudFront edge locations



access objects in Amazon S3. CloudFront also supports Geo Restriction, which restricts access to your content based on the geographic location of your viewers.

To control access to the original copies of your objects in Amazon S3, Amazon CloudFront allows you to create one or more “Origin Access Identities” and associate these with your distributions. When an Origin Access Identity is associated with an Amazon CloudFront distribution, the distribution will use that identity to retrieve objects from Amazon S3. You can then use Amazon S3’s ACL feature, which limits access to that Origin Access Identity so the original copy of the object is not publicly readable.

To control who is able to download objects from Amazon CloudFront edge locations, the service uses a signed-URL verification system. To use this system, you first create a public-private key pair, and upload the public key to your account via the AWS Management Console. Second, you configure your Amazon CloudFront distribution to indicate which accounts you would authorize to sign requests – you can indicate up to five AWS Accounts you trust to sign requests. Third, as you receive requests you will create policy documents indicating the conditions under which you want Amazon CloudFront to serve your content. These policy documents can specify the name of the object that is requested, the date and time of the request, and the source IP (or CIDR range) of the client making the request. You then calculate the SHA1 hash of your policy document and sign this using your private key. Finally, you include both the encoded policy document and the signature as query string parameters when you reference your objects. When Amazon CloudFront receives a request, it will decode the signature using your public key. Amazon CloudFront will only serve requests that have a valid policy document and matching signature.

Note that private content is an optional feature that must be enabled when you set up your CloudFront distribution. Content delivered without this feature enabled will be publicly readable.

Amazon CloudFront provides the option to transfer content over an encrypted connection (HTTPS). By default, CloudFront will accept requests over both HTTP and HTTPS protocols. However, you can also configure CloudFront to require HTTPS for all requests or have CloudFront redirect HTTP requests to HTTPS. You can even configure CloudFront distributions to allow HTTP for some objects but require HTTPS for other objects.




Figure 7: Amazon CloudFront Encrypted Transmission

You can configure one or more CloudFront origins to require CloudFront fetch objects from your origin using the protocol that the viewer used to request the objects. For example, when you use this CloudFront setting and the viewer uses HTTPS to request an object from CloudFront, CloudFront also uses HTTPS to forward the request to your origin.

Amazon CloudFront uses the SSLv3 or TLSv1 protocols and a selection of cipher suites that includes the Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) protocol on connections to both viewers and the origin. ECDHE allows SSL/TLS clients

to provide Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This helps prevent the decoding of captured data by unauthorized third parties, even if the secret long-term key itself is compromised.

Note that if you're using your own server as your origin, and you want to use HTTPS both between viewers and CloudFront and between CloudFront and your origin, you must install a valid SSL certificate on the HTTP server that is signed by a third-party certificate authority, for example, VeriSign or DigiCert.

By default, you can deliver content to viewers over HTTPS by using your CloudFront distribution domain name in your URLs; for example, <https://dxxxxx.cloudfront.net/image.jpg>. If you want to deliver your content over HTTPS using your own domain name and your own SSL certificate, you can use SNI Custom SSL or Dedicated IP Custom SSL. With Server Name Identification (SNI) Custom SSL, CloudFront relies on the SNI extension of the TLS protocol, which is supported by most modern web browsers. However, some users may not be able to access your content because some older browsers do not support SNI. (For a list of supported browsers, visit <http://aws.amazon.com/cloudfront/faqs/>.) With Dedicated IP Custom SSL, CloudFront dedicates IP addresses to your SSL certificate at each CloudFront edge location so that CloudFront can associate the incoming requests with the proper SSL certificate.

Amazon CloudFront access logs contain a comprehensive set of information about requests for content, including the object requested, the date and time of the request, the edge location serving the request, the client IP address, the referrer, and the user agent. To enable access logs, just specify the name of the Amazon S3 bucket to store the logs in when you configure your Amazon CloudFront distribution.

AWS Direct Connect Security

With AWS Direct Connect, you can provision a direct link between your internal network and an AWS region using a high-throughput, dedicated connection. Doing this may help reduce your network costs, improve throughput, or provide a more consistent network experience. With this dedicated connection in place, you can then create virtual interfaces directly to the AWS cloud (for example, to Amazon EC2 and Amazon S3) and Amazon VPC.

With Direct Connect, you bypass Internet service providers in your network path. You can procure rack space within the facility housing the AWS Direct Connect location and deploy your equipment nearby. Once deployed, you can connect this equipment to AWS Direct Connect using a cross-connect. Each AWS Direct Connect location enables connectivity to the geographically nearest AWS region as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

Using industry standard 802.1q VLANs, the dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon VPC using private IP space, while maintaining network separation between the public and private environments.

Amazon Direct Connect requires the use of the Border Gateway Protocol (BGP) with an Autonomous System Number (ASN). To create a virtual interface, you use an MD5 cryptographic key for message authorization. MD5 creates a keyed hash using your secret key. You can have AWS automatically generate a BGP MD5 key or you can provide your own.



Storage Services

Amazon Web Services provides low-cost data storage with high durability and availability. AWS offers storage choices for backup, archiving, and disaster recovery, as well as block and object storage.

Amazon Simple Storage Service (Amazon S3) Security

Amazon Simple Storage Service (S3) allows you to upload and retrieve data at any time, from anywhere on the web. Amazon S3 stores data as *objects* within *buckets*. An object can be any kind of file: a text file, a photo, a video, etc. When you add a file to Amazon S3, you have the option of including metadata with the file and setting permissions to control access to the file. For each bucket, you can control access to the bucket (who can create, delete, and list objects in the bucket), view access logs for the bucket and its objects, and choose the geographical region where Amazon S3 will store the bucket and its contents.

Data Access

Access to data stored in Amazon S3 is restricted by default; only bucket and object owners have access to the Amazon S3 resources they create (note that a bucket/object owner is the AWS Account owner, not the user who created the bucket/object). There are multiple ways to control access to buckets and objects:

- **Identity and Access Management (IAM) Policies.** AWS IAM enables organizations with many employees to create and manage multiple users under a single AWS Account. IAM policies are attached to the users, enabling centralized control of permissions for users under your AWS Account to access buckets or objects. With IAM policies, you can only grant *users within your own AWS account* permission to access your Amazon S3 resources.
- **Access Control Lists (ACLs).** Within Amazon S3, you can use ACLs to give read or write access on buckets or objects to groups of users. With ACLs, you can only grant *other AWS accounts* (not specific users) access to your Amazon S3 resources.
- **Bucket Policies.** Bucket policies in Amazon S3 can be used to add or deny permissions across some or all of the objects within a single bucket. Policies can be attached to users, groups, or Amazon S3 buckets, enabling centralized management of permissions. With bucket policies, you can grant users within your AWS Account *or* other AWS Accounts access to your Amazon S3 resources.

Type of Access Control	AWS Account-Level Control?	User-Level Control?
IAM Policies	No	Yes
ACLs	Yes	No
Bucket Policies	Yes	Yes

You can further restrict access to specific resources based on certain conditions. For example, you can restrict access based on request time (Date Condition), whether the request was sent using SSL (Boolean Conditions), a requester's IP address (IP Address Condition), or based on the requester's client application (String Conditions). To identify these conditions, you use *policy keys*. For more information about action-specific policy keys available within Amazon S3, refer to the [Amazon Simple Storage Service Developer Guide](#).

Amazon S3 also gives developers the option to use *query string authentication*, which allows them to share Amazon S3 objects through URLs that are valid for a predefined period of time. Query string authentication is useful for giving HTTP



or browser access to resources that would normally require authentication. The signature in the query string secures the request.

Data Transfer

For maximum security, you can securely upload/download data to Amazon S3 via the SSL encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data is transferred securely both within AWS and to and from sources outside of AWS.

Data Storage

Amazon S3 provides multiple options for protecting data at rest. For customers who prefer to manage their own encryption, they can use a client encryption library like the [Amazon S3 Encryption Client](#) to encrypt data before uploading to Amazon S3. Alternatively, you can use Amazon S3 Server Side Encryption (SSE) if you prefer to have Amazon S3 manage the encryption process for you. Data is encrypted with a key generated by AWS or with a key you supply, depending on your requirements. With Amazon S3 SSE, you can encrypt data on upload simply by adding an additional request header when writing the object. Decryption happens automatically when data is retrieved.

Note that metadata, which you can include with your object, is not encrypted. Therefore, AWS recommends that customers not place sensitive information in Amazon S3 metadata.

Amazon S3 SSE uses one of the strongest block ciphers available – 256-bit Advanced Encryption Standard (AES-256). With Amazon S3 SSE, every protected object is encrypted with a unique encryption key. This object key itself is then encrypted with a regularly rotated master key. Amazon S3 SSE provides additional security by storing the encrypted data and encryption keys in different hosts. Amazon S3 SSE also makes it possible for you to enforce encryption requirements. For example, you can create and apply bucket policies that require that only encrypted data can be uploaded to your buckets.

For long-term storage, you can automatically archive the contents of your Amazon S3 buckets to AWS's archival service called Glacier. You can have data transferred at specific intervals to Glacier by creating lifecycle rules in Amazon S3 that describe which objects you want to be archived to Glacier and when. As part of your data management strategy, you can also specify how long Amazon S3 should wait after the objects are put into Amazon S3 to delete them.

When an object is deleted from Amazon S3, removal of the mapping from the public name to the object starts immediately, and is generally processed across the distributed system within several seconds. Once the mapping is removed, there is no remote access to the deleted object. The underlying storage area is then reclaimed for use by the system.

Data Durability and Reliability

Amazon S3 is designed to provide 99.999999999% durability and 99.99% availability of objects over a given year. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 region. To help provide durability, Amazon S3 PUT and COPY operations synchronously store customer data across multiple facilities before returning SUCCESS. Once stored, Amazon S3 helps maintain the durability of the objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. In addition, Amazon S3 calculates checksums on all network traffic to detect corruption of data packets when storing or retrieving data.

Amazon S3 provides further protection via Versioning. You can use Versioning to preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket. With Versioning, you can easily recover from both unintended



user actions and application failures. By default, requests will retrieve the most recently written version. Older versions of an object can be retrieved by specifying a version in the request. You can further protect versions using Amazon S3 Versioning's MFA Delete feature. Once enabled for an Amazon S3 bucket, each version deletion request must include the six-digit code and serial number from your multi-factor authentication device.

Access Logs

An Amazon S3 bucket can be configured to log access to the bucket and objects within it. The access log contains details about each access request including request type, the requested resource, the requestor's IP, and the time and date of the request. When logging is enabled for a bucket, log records are periodically aggregated into log files and delivered to the specified Amazon S3 bucket.

Cross-Origin Resource Sharing (CORS)

AWS customers who use Amazon S3 to host static web pages or store objects used by other web pages can load content securely by configuring an Amazon S3 bucket to explicitly enable cross-origin requests. Modern browsers use the Same Origin policy to block JavaScript or HTML5 from allowing requests to load content from another site or domain as a way to help ensure that malicious content is not loaded from a less reputable source (such as during cross-site scripting attacks). With the Cross-Origin Resource Sharing (CORS) policy enabled, assets such as web fonts and images stored in an Amazon S3 bucket can be safely referenced by external web pages, style sheets, and HTML5 applications.

AWS Glacier Security

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where Amazon S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often and for which retrieval times of several hours are suitable.

Amazon Glacier stores files as *archives* within *vaults*. Archives can be any data such as a photo, video, or document, and can contain one or several files. You can store an unlimited number of archives in a single vault and can create up to 1,000 vaults per region. Each archive can contain up to 40 TB of data.

Data Upload

To transfer data into Amazon Glacier vaults, you can upload an archive in a single upload operation or a multipart operation. In a single upload operation, you can upload archives up to 4 GB in size. However, customers can achieve better results using the Multipart Upload API to upload archives greater than 100 MB. Using the Multipart Upload API allows you to upload large archives, up to about 40,000 GB. The Multipart Upload API call is designed to improve the upload experience for larger archives; it enables the parts to be uploaded independently, in any order, and in parallel. If a multipart upload fails, you only need to upload the failed part again and not the entire archive.

When you upload data to Glacier, you must compute and supply a tree hash. Glacier checks the hash against the data to help ensure that it has not been altered en route. A tree hash is generated by computing a hash for each megabyte-sized segment of the data, and then combining the hashes in tree fashion to represent ever-growing adjacent segments of the data.

As an alternate to using the Multipart Upload feature, customers with very large uploads to Amazon Glacier may consider using the AWS Import/Export service instead to transfer the data. AWS Import/Export facilitates moving large amounts of data into AWS using portable storage devices for transport. AWS transfers your data directly off of storage devices using Amazon's high-speed internal network, bypassing the Internet.



You can also set up Amazon S3 to transfer data at specific intervals to Glacier. You can create lifecycle rules in Amazon S3 that describe which objects you want to be archived to Glacier and when. You can also specify how long Amazon S3 should wait after the objects are put into Amazon S3 to delete them.

To achieve even greater security, you can securely upload/download data to Amazon Glacier via the SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data is transferred securely both within AWS and to and from sources outside of AWS.

Data Retrieval

Retrieving archives from Amazon Glacier requires the initiation of a retrieval job, which is generally completed in 3 to 5 hours. You can then access the data via HTTP GET requests. The data will remain available to you for 24 hours.

You can retrieve an entire archive or several files from an archive. If you want to retrieve only a subset of an archive, you can use one retrieval request to specify the range of the archive that contains the files you are interested or you can initiate multiple retrieval requests, each with a range for one or more files. You can also limit the number of vault inventory items retrieved by filtering on an archive creation date range or by setting a maximum items limit. Whichever method you choose, when you retrieve portions of your archive, you can use the supplied checksum to help ensure the integrity of the files provided that the range that is retrieved is aligned with the tree hash of the overall archive.

Data Storage

Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks and is built to be automatically self-healing.

Data Access

Only your account can access your data in Amazon Glacier. To control access to your data in Amazon Glacier, you can use AWS IAM to specify which users within your account have rights to operations on a given vault.

AWS Storage Gateway Security

The AWS Storage Gateway service connects your on-premises software appliance with cloud-based storage to provide seamless and secure integration between your IT environment and AWS's storage infrastructure. The service enables you to securely upload data to AWS' scalable, reliable, and secure Amazon S3 storage service for cost-effective backup and rapid disaster recovery.

AWS Storage Gateway transparently backs up data off-site to Amazon S3 in the form of Amazon EBS snapshots. Amazon S3 redundantly stores these snapshots on multiple devices across multiple facilities, detecting and repairing any lost redundancy. The Amazon EBS snapshot provides a point-in-time backup that can be restored on-premises or used to instantiate new Amazon EBS volumes. Data is stored within a single region that you specify.

AWS Storage Gateway offers three options:

- **Gateway-Stored Volumes (where the cloud is backup).** In this option, your volume data is stored locally and then pushed to Amazon S3, where it is stored in redundant, encrypted form, and made available in the form of Elastic Block Storage (EBS) snapshots. When you use this model, the on-premises storage is primary, delivering low-latency access to your entire dataset, and the cloud storage is the backup.



- **Gateway-Cached Volumes (where the cloud is primary).** In this option, your volume data is stored encrypted in Amazon S3, visible within your enterprise's network via an iSCSI interface. Recently accessed data is cached on-premises for low-latency local access. When you use this model, the cloud storage is primary, but you get low-latency access to your active working set in the cached volumes on premises.
- **Gateway-Virtual Tape Library (VTL).** In this option, you can configure a Gateway-VTL with up to 10 virtual tape drives per gateway, 1 media changer and up to 1500 virtual tape cartridges. Each virtual tape drive responds to the SCSI command set, so your existing on-premises backup applications (either disk-to-tape or disk-to-disk-to-tape) will work without modification.

No matter which option you choose, data is asynchronously transferred from your on-premises storage hardware to AWS over SSL. The data is stored encrypted in Amazon S3 using Advanced Encryption Standard (AES) 256, a symmetric-key encryption standard using 256-bit encryption keys. The AWS Storage Gateway only uploads data that has changed, minimizing the amount of data sent over the Internet.

The AWS Storage Gateway runs as a virtual machine (VM) that you deploy on a host in your data center running VMware ESXi Hypervisor v 4.1 or v 5 or Microsoft Hyper-V (you download the VMware software during the setup process). You can also run within EC2 using a gateway AMI. During the installation and configuration process, you can create up to 12 stored volumes, 20 Cached volumes, or 1500 virtual tape cartridges per gateway. Once installed, each gateway will automatically download, install, and deploy updates and patches. This activity takes place during a maintenance window that you can set on a per-gateway basis.

The iSCSI protocol supports authentication between targets and initiators via CHAP (Challenge-Handshake Authentication Protocol). CHAP provides protection against man-in-the-middle and playback attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a storage volume target. To set up CHAP, you must configure it in both the AWS Storage Gateway console and in the iSCSI initiator software you use to connect to the target.

After you deploy the AWS Storage Gateway VM, you must activate the gateway using the AWS Storage Gateway console. The activation process associates your gateway with your AWS Account. Once you establish this connection, you can manage almost all aspects of your gateway from the console. In the activation process, you specify the IP address of your gateway, name your gateway, identify the AWS region in which you want your snapshot backups stored, and specify the gateway time zone.

AWS Import/Export Security

AWS Import/Export is a simple, secure method for physically transferring large amounts of data to Amazon S3, EBS, or Glacier storage. This service is typically used by customers who have over 100 GB of data and/or slow connection speeds that would result in very slow transfer rates over the Internet. With AWS Import/Export, you prepare a portable storage device that you ship to a secure AWS facility. AWS transfers the data directly off of the storage device using Amazon's high-speed internal network, thus bypassing the Internet. Conversely, data can also be exported from AWS to a portable storage device.

Like all other AWS services, the AWS Import/Export service requires that you securely identify and authenticate your storage device. In this case, you will submit a job request to AWS that includes your Amazon S3 bucket, Amazon EBS region, AWS Access Key ID, and return shipping address. You then receive a unique identifier for the job, a digital signature for authenticating your device, and an AWS address to ship the storage device to. For Amazon S3, you place



the signature file on the root directory of your device. For Amazon EBS, you tape the signature barcode to the exterior of the device. The signature file is used only for authentication and is not uploaded to Amazon S3 or EBS.

For transfers to Amazon S3, you specify the specific buckets to which the data should be loaded and ensure that the account doing the loading has write permission for the buckets. You should also specify the access control list to be applied to each object loaded to Amazon S3.

For transfers to EBS, you specify the target region for the EBS import operation. If the storage device is less than or equal to the maximum volume size of 1 TB, its contents are loaded directly into an Amazon EBS snapshot. If the storage device's capacity exceeds 1 TB, a device image is stored within the specified S3 log bucket. You can then create a RAID of Amazon EBS volumes using software such as Logical Volume Manager, and copy the image from S3 to this new volume.

For added protection, you can encrypt the data on your device before you ship it to AWS. For Amazon S3 data, you can use a PIN-code device with hardware encryption or TrueCrypt software to encrypt your data before sending it to AWS. For EBS and Glacier data, you can use any encryption method you choose, including a PIN-code device. AWS will decrypt your Amazon S3 data before importing using the PIN code and/or TrueCrypt password you supply in your import manifest. AWS uses your PIN to access a PIN-code device, but does not decrypt software-encrypted data for import to Amazon EBS or Amazon Glacier. The following table summarizes your encryption options for each type of import/export job.

Import to Amazon S3		
Source	Target	Result
- Files on a device file system - Encrypt data using PIN-code device and/or TrueCrypt before shipping device	- Objects in an existing Amazon S3 bucket - AWS decrypts the data before performing the import	- One object for each file. - AWS erases your device after every import job prior to shipping
Export from Amazon S3		
Source	Target	Result
- Objects in one or more Amazon S3 buckets - Provide a PIN code and/or password that AWS will use to encrypt your data	- Files on your storage device - AWS formats your device - AWS copies your data to an encrypted file container on your device	- One file for each object - AWS encrypts your data prior to shipping - Use PIN-code device and/or TrueCrypt to decrypt the files
Import to Amazon Glacier		
Source	Target	Result
- Entire device - Encrypt the data using the encryption method of your choice before shipping	- One archive in an existing Amazon Glacier vault - AWS does not decrypt your device	- Device image stored as a single archive - AWS erases your device after every import job prior to shipping
Import to Amazon EBS (Device Capacity < 1 TB)		
Source	Target	Result

- Entire device - Encrypt the data using the encryption method of your choice before shipping	- One Amazon EBS snapshot - AWS does not decrypt your device	- Device image is stored as a single snapshot - If the device was encrypted, the image is encrypted - AWS erases your device after every import job prior to shipping
Import to Amazon EBS (Device Capacity > 1 TB)		
Source	Target	Result
- Entire device - Encrypt the data using the encryption method of your choice before shipping	- Multiple objects in an existing Amazon S3 bucket - AWS does not decrypt your device	- Device image chunked into series of 1 TB snapshots stored as objects in Amazon S3 bucket specified in manifest file - If the device was encrypted, the image is encrypted - AWS erases your device after every import job prior to shipping

After the import is complete, AWS Import/Export will erase the contents of your storage device to safeguard the data during return shipment. AWS overwrites all writable blocks on the storage device with zeroes. You will need to repartition and format the device after the wipe. If AWS is unable to erase the data on the device, it will be scheduled for destruction and our support team will contact you using the email address specified in the manifest file you ship with the device.

When shipping a device internationally, the customs option and certain required subfields are required in the manifest file sent to AWS. AWS Import/Export uses these values to validate the inbound shipment and prepare the outbound customs paperwork. Two of these options are whether the data on the device is encrypted or not and the encryption software's classification. When shipping encrypted data to or from the United States, the encryption software must be classified as 5D992 under the United States Export Administration Regulations.

Database Services

Amazon Web Services provides a number of database solutions for developers and businesses—from managed relational and NoSQL database services, to in-memory caching as a service and petabyte-scale data-warehouse service.

Amazon DynamoDB Security


Amazon DynamoDB is a managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables you to offload the administrative burdens of operating and scaling distributed databases to AWS, so you don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling.

You can create a database table that can store and retrieve any amount of data, and serve any level of request traffic. DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity you specified and the amount of data stored, while maintaining consistent, fast performance. All data items are stored on Solid State Drives (SSDs) and are automatically replicated across multiple availability zones in a region to provide built-in high availability and data durability.



You can set up automatic backups using a special template in AWS Data Pipeline that was created just for copying DynamoDB tables. You can choose full or incremental backups to a table in the same region or a different region. You can use the copy for disaster recovery (DR) in the event that an error in your code damages the original table, or to federate DynamoDB data across regions to support a multi-region application.

To control who can use the DynamoDB resources and API, you set up permissions in AWS IAM. In addition to controlling access at the resource-level with IAM, you can also control access at the database level—you can create database-level permissions that allow or deny access to items (rows) and attributes (columns) based on the needs of your application. These database-level permissions are called *fine-grained access controls*, and you create them using an IAM policy that specifies under what circumstances a user or application can access a DynamoDB table. The IAM policy can restrict access to individual items in a table, access to the attributes in those items, or both at the same time.



You can optionally use web identity federation to control access by application users who are authenticated by Login with Amazon, Facebook, or Google. Web identity federation removes the need for creating individual IAM users; instead, users can sign in to an identity provider and then obtain temporary security credentials from AWS Security Token Service (AWS STS). AWS STS returns temporary AWS credentials to the application and allows it to access the specific DynamoDB table.

In addition to requiring database and user permissions, each request to the DynamoDB service must contain a valid HMAC-SHA256 signature, or the request is rejected. The AWS SDKs automatically sign your requests; however, if you want to write your own HTTP POST requests, you must provide the signature in the header of your request to Amazon DynamoDB. To calculate the signature, you must request temporary security credentials from the AWS Security Token Service. Use the temporary security credentials to sign your requests to Amazon DynamoDB.

Amazon DynamoDB is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2.

Amazon Relational Database Service (Amazon RDS) Security

Amazon RDS allows you to quickly create a relational database (DB) instance and flexibly scale the associated compute resources and storage capacity to meet application demand. Amazon RDS manages the database instance on your behalf by performing backups, handling failover, and maintaining the database software. Currently, Amazon RDS is available for MySQL, Oracle, Microsoft SQL Server, and PostgreSQL database engines.

Amazon RDS has multiple features that enhance reliability for critical production databases, including DB security groups, permissions, SSL connections, automated backups, DB snapshots, and multi-AZ deployments. DB instances can also be deployed in an Amazon VPC for additional network isolation.

Access Control

When you first create a DB Instance within Amazon RDS, you will create a master user account, which is used only within the context of Amazon RDS to control access to your DB Instance(s). The master user account is a native database user account that allows you to log on to your DB Instance with all database privileges. You can specify the master user name and password you want associated with each DB Instance when you create the DB Instance. Once you have created your DB Instance, you can connect to the database using the master user credentials. Subsequently, you can create additional user accounts so that you can restrict who can access your DB Instance.

You can control Amazon RDS DB Instance access via DB Security Groups, which are similar to Amazon EC2 Security Groups but not interchangeable. DB Security Groups act like a firewall controlling network access to your DB Instance. Database Security Groups default to a “deny all” access mode and customers must specifically authorize network ingress. There are two ways of doing this: authorizing a network IP range or authorizing an existing Amazon EC2 Security Group. DB Security Groups only allow access to the database server port (all others are blocked) and can be updated without restarting the Amazon RDS DB Instance, which allows a customer seamless control of their database access. Using AWS IAM, you can further control access to your RDS DB instances. AWS IAM enables you to control what RDS operations each individual AWS IAM user has permission to call.

Network Isolation

For additional network access control, you can run your DB Instances in an Amazon VPC. Amazon VPC enables you to isolate your DB Instances by specifying the IP range you wish to use, and connect to your existing IT infrastructure through industry-standard encrypted IPsec VPN. Running Amazon RDS in a VPC enables you to have a DB instance within a private subnet. You can also set up a virtual private gateway that extends your corporate network into your VPC, and allows access to the RDS DB instance in that VPC. Refer to the [Amazon VPC User Guide](#) for more details.

For Multi-AZ deployments, defining a subnet for all availability zones in a region will allow Amazon RDS to create a new standby in another availability zone should the need arise. You can create DB Subnet Groups, which are collections of subnets that you may want to designate for your RDS DB Instances in a VPC. Each DB Subnet Group should have at least one subnet for every availability zone in a given region. In this case, when you create a DB Instance in a VPC, you select a DB Subnet Group; Amazon RDS then uses that DB Subnet Group and your preferred availability zone to select a subnet and an IP address within that subnet. Amazon RDS creates and associates an Elastic Network Interface to your DB Instance with that IP address.

DB Instances deployed within an Amazon VPC can be accessed from the Internet or from Amazon EC2 Instances outside the VPC via VPN or bastion hosts that you can launch in your public subnet. To use a bastion host, you will need to set up a public subnet with an EC2 instance that acts as a SSH Bastion. This public subnet must have an Internet gateway and routing rules that allow traffic to be directed via the SSH host, which must then forward requests to the private IP address of your Amazon RDS DB instance.

DB Security Groups can be used to help secure DB Instances within an Amazon VPC. In addition, network traffic entering and exiting each subnet can be allowed or denied via network ACLs. All network traffic entering or exiting your Amazon VPC via your IPsec VPN connection can be inspected by your on-premises security infrastructure, including network firewalls and intrusion detection systems.



Encryption

You can encrypt connections between your application and your DB Instance using SSL. For MySQL and SQL Server, RDS creates an SSL certificate and installs the certificate on the DB instance when the instance is provisioned. For MySQL, you launch the mysql client using the `--ssl_ca` parameter to reference the public key in order to encrypt connections. For SQL Server, download the public key and import the certificate into your Windows operating system. Oracle RDS uses Oracle native network encryption with a DB instance. You simply add the native network encryption option to an option group and associate that option group with the DB instance. Once an encrypted connection is established, data transferred between the DB Instance and your application will be encrypted during transfer. You can also require your DB instance to only accept encrypted connections.

Amazon RDS supports Transparent Data Encryption (TDE) for SQL Server (SQL Server Enterprise Edition) and Oracle (part of the Oracle Advanced Security option available in Oracle Enterprise Edition). The TDE feature automatically encrypts data before it is written to storage and automatically decrypts data when it is read from storage. If you require your MySQL data to be encrypted while “at rest” in the database, your application must manage the encryption and decryption of data.

Note that SSL support within Amazon RDS is for encrypting the connection between your application and your DB Instance; it should not be relied on for authenticating the DB Instance itself.

While SSL offers security benefits, be aware that SSL encryption is a compute intensive operation and will increase the latency of your database connection. To learn more about how SSL works with MySQL, you can refer directly to the MySQL documentation found [here](#). To learn how SSL works with SQL Server, you can read more in the [RDS User Guide](#).

Automated Backups and DB Snapshots

Amazon RDS provides two different methods for backing up and restoring your DB Instance(s): automated backups and database snapshots (DB Snapshots).

Turned on by default, the automated backup feature of Amazon RDS enables point-in-time recovery for your DB Instance. Amazon RDS will back up your database and transaction logs and store both for a user-specified retention period. This allows you to restore your DB Instance to any second during your retention period, up to the last 5 minutes. Your automatic backup retention period can be configured to up to 35 days.

During the backup window, storage I/O may be suspended while your data is being backed up. This I/O suspension typically lasts a few minutes. This I/O suspension is avoided with Multi-AZ DB deployments, since the backup is taken from the standby.

DB Snapshots are user-initiated backups of your DB Instance. These full database backups are stored by Amazon RDS until you explicitly delete them. You can copy DB snapshots of any size and move them between any of AWS’s public regions, or copy the same snapshot to multiple regions simultaneously. You can then create a new DB Instance from a DB Snapshot whenever you desire.

DB Instance Replication

Amazon cloud computing resources are housed in highly available data center facilities in different regions of the world, and each region contains multiple distinct locations called Availability Zones. Each Availability Zone is engineered to be isolated from failures in other Availability Zones, and to provide inexpensive, low-latency network connectivity to other Availability Zones in the same region.



To architect for high availability of your Oracle, PostgreSQL, or MySQL databases, you can run your RDS DB instance in several Availability Zones, an option called a Multi-AZ deployment. When you select this option, Amazon automatically provisions and maintains a synchronous standby replica of your DB instance in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to the standby replica. In the event of DB instance or Availability Zone failure, Amazon RDS will automatically failover to the standby so that database operations can resume quickly without administrative intervention.

For customers who use MySQL and need to scale beyond the capacity constraints of a single DB Instance for read-heavy database workloads, Amazon RDS provides a Read Replica option. Once you create a read replica, database updates on the source DB instance are replicated to the read replica using MySQL's native, asynchronous replication. You can create multiple read replicas for a given source DB instance and distribute your application's read traffic among them. Read replicas can be created with Multi-AZ deployments to gain read scaling benefits in addition to the enhanced database write availability and data durability provided by Multi-AZ deployments.

Automatic Software Patching

Amazon RDS will make sure that the relational database software powering your deployment stays up-to-date with the latest patches. When necessary, patches are applied during a maintenance window that you can control. You can think of the Amazon RDS maintenance window as an opportunity to control when DB Instance modifications (such as scaling DB Instance class) and software patching occur, in the event either are requested or required. If a "maintenance" event is scheduled for a given week, it will be initiated and completed at some point during the 30-minute maintenance window you identify.

The only maintenance events that require Amazon RDS to take your DB Instance offline are scale compute operations (which generally take only a few minutes from start-to-finish) or required software patching. Required patching is automatically scheduled only for patches that are security and durability related. Such patching occurs infrequently (typically once every few months) and should seldom require more than a fraction of your maintenance window. If you do not specify a preferred weekly maintenance window when creating your DB Instance, a 30-minute default value is assigned. If you wish to modify when maintenance is performed on your behalf, you can do so by modifying your DB Instance in the [AWS Management Console](#) or by using the `ModifyDBInstance` API. Each of your DB Instances can have different preferred maintenance windows, if you so choose.

Running your DB Instance as a Multi-AZ deployment can further reduce the impact of a maintenance event, as Amazon RDS will conduct maintenance via the following steps: 1) Perform maintenance on standby, 2) Promote standby to primary, and 3) Perform maintenance on old primary, which becomes the new standby.

When an Amazon RDS DB Instance deletion API (`DeleteDBInstance`) is run, the DB Instance is marked for deletion. Once the instance no longer indicates 'deleting' status, it has been removed. At this point the instance is no longer accessible and unless a final snapshot copy was asked for, it cannot be restored and will not be listed by any of the tools or APIs.

Event Notification

You can receive notifications of a variety of important events that can occur on your RDS instance, such as whether the instance was shut down, a backup was started, a failover occurred, the security group was changed, or your storage space is low. The Amazon RDS service groups events into categories that you can subscribe to so that you can be notified when an event in that category occurs. You can subscribe to an event category for a DB instance, DB snapshot, DB security group, or for a DB parameter group. RDS events are published via AWS SNS and sent to you as an email or text message. For more information about RDS notification event categories, refer to the [RDS User Guide](#).



Amazon Redshift Security

Amazon Redshift is a petabyte-scale SQL data warehouse service that runs on highly optimized and managed AWS compute and storage resources. The service has been architected to not only scale up or down rapidly, but to significantly improve query speeds—even on extremely large datasets. To increase performance, Redshift uses techniques such as columnar storage, data compression, and zone maps to reduce the amount of IO needed to perform queries. It also has a massively parallel processing (MPP) architecture, parallelizing and distributing SQL operations to take advantage of all available resources.

When you create a Redshift data warehouse, you provision a single-node or multi-node cluster, specifying the type and number of nodes that will make up the cluster. The node type determines the storage size, memory, and CPU of each node. Each multi-node cluster includes a leader node and two or more compute nodes. A leader node manages connections, parses queries, builds execution plans, and manages query execution in the compute nodes. The compute nodes store data, perform computations, and run queries as directed by the leader node. The leader node of each cluster is accessible through ODBC and JDBC endpoints, using standard PostgreSQL drivers. The compute nodes run on a separate, isolated network and are never accessed directly.

After you provision a cluster, you can upload your dataset and perform data analysis queries by using common SQL-based tools and business intelligence applications.

Cluster Access

By default, clusters that you create are closed to everyone. Amazon Redshift enables you to configure firewall rules (security groups) to control network access to your data warehouse cluster. You can also run Redshift inside an Amazon VPC to isolate your data warehouse cluster in your own virtual network and connect it to your existing IT infrastructure using industry-standard encrypted IPsec VPN.

The AWS account that creates the cluster has full access to the cluster. Within your AWS account, you can use AWS IAM to create user accounts and manage permissions for those accounts. By using IAM, you can grant different users permission to perform only the cluster operations that are necessary for their work.

Like all databases, you must grant permission in Redshift at the database level in addition to granting access at the resource level. Database users are named user accounts that can connect to a database and are authenticated when they login to Amazon Redshift. In Redshift, you grant database user permissions on a per-cluster basis instead of on a per-table basis. However, a user can see data only in the table rows that were generated by his own activities; rows generated by other users are not visible to him.

The user who creates a database object is its owner. By default, only a superuser or the owner of an object can query, modify, or grant permissions on the object. For users to use an object, you must grant the necessary permissions to the user or the group that contains the user. And only the owner of an object can modify or delete it.

Data Backups

Amazon Redshift distributes your data across all compute nodes in a cluster. When you run a cluster with at least two compute nodes, data on each node will always be mirrored on disks on another node, reducing the risk of data loss. In addition, all data written to a node in your cluster is continuously backed up to Amazon S3 using snapshots. Redshift stores your snapshots for a user-defined period, which can be from one to thirty-five days. You can also take your own snapshots at any time; these snapshots leverage all existing system snapshots and are retained until you explicitly delete them.



Amazon Redshift continuously monitors the health of the cluster and automatically re-replicates data from failed drives and replaces nodes as necessary. All of this happens without any effort on your part, although you may see a slight performance degradation during the re-replication process.

You can use any system or user snapshot to restore your cluster using the AWS Management Console or the Amazon Redshift APIs. Your cluster is available as soon as the system metadata has been restored and you can start running queries while user data is spooled down in the background.

Data Encryption

When creating a cluster, you can choose to encrypt it in order to provide additional protection for your data at rest. When you enable encryption in your cluster, Amazon Redshift stores all data in user-created tables in an encrypted format using hardware-accelerated AES-256 block encryption keys. This includes all data written to disk as well as any backups.

Amazon Redshift uses a four-tier, key-based architecture for encryption. These keys consist of data encryption keys, a database key, a cluster key, and a master key:

- *Data encryption keys* encrypt data blocks in the cluster. Each data block is assigned a randomly-generated AES-256 key. These keys are encrypted by using the database key for the cluster.
- The *database key* encrypts data encryption keys in the cluster. The database key is a randomly-generated AES-256 key. It is stored on disk in a separate network from the Amazon Redshift cluster and encrypted by a master key. Amazon Redshift passes the database key across a secure channel and keeps it in memory in the cluster.
- The *cluster key* encrypts the database key for the Amazon Redshift cluster. You can use either AWS or a hardware security module (HSM) to store the cluster key. HSMs provide direct control of key generation and management, and make key management separate and distinct from the application and the database.
- The *master key* encrypts the cluster key if it is stored in AWS. The master key encrypts the cluster-key-encrypted database key if the cluster key is stored in an HSM.

You can have Redshift rotate the encryption keys for your encrypted clusters at any time. As part of the rotation process, keys are also updated for all of the cluster's automatic and manual snapshots.

Note that enabling encryption in your cluster will impact performance, even though it is hardware accelerated. Encryption also applies to backups. When restoring from an encrypted snapshot, the new cluster will be encrypted as well.

To encrypt your table load data files when you upload them to Amazon S3, you can use Amazon S3 server-side encryption. When you load the data from Amazon S3, the COPY command will decrypt the data as it loads the table.

Database Audit Logging

Amazon Redshift logs all SQL operations, including connection attempts, queries, and changes to your database. You can access these logs using SQL queries against system tables or choose to have them downloaded to a secure Amazon S3 bucket. You can then use these audit logs to monitor your cluster for security and troubleshooting purposes.



Automatic Software Patching

Amazon Redshift manages all the work of setting up, operating, and scaling your data warehouse, including provisioning capacity, monitoring the cluster, and applying patches and upgrades to the Amazon Redshift engine. Patches are applied only during specified maintenance windows.

SSL Connections

To protect your data in transit within the AWS cloud, Amazon Redshift uses hardware-accelerated SSL to communicate with Amazon S3 or Amazon DynamoDB for COPY, UNLOAD, backup, and restore operations. You can encrypt the connection between your client and the cluster by specifying SSL in the parameter group associated with the cluster. To have your clients also authenticate the Redshift server, you can install the public key (.pem file) for the SSL certificate on your client and use the key to connect to your clusters.

Amazon Redshift offers the newer, stronger cipher suites that use the Elliptic Curve Diffie-Hellman Ephemeral protocol. ECDHE allows SSL clients to provide Perfect Forward Secrecy between the client and the Redshift cluster. Perfect Forward Secrecy uses session keys that are ephemeral and not stored anywhere, which prevents the decoding of captured data by unauthorized third parties, even if the secret long-term key itself is compromised. You do not need to configure anything in Amazon Redshift to enable ECDHE; if you connect from a SQL client tool that uses ECDHE to encrypt communication between the client and server, Amazon Redshift will use the provided cipher list to make the appropriate connection.

Amazon ElastiCache Security

Amazon ElastiCache is a web service that makes it easy to set up, manage, and scale distributed in-memory cache environments in the cloud. The service improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. It can be used to significantly improve latency and throughput for many read-heavy application workloads (such as social networking, gaming, media sharing, and Q&A portals) or compute-intensive workloads (such as a recommendation engine). Caching improves application performance by storing critical pieces of data in memory for low-latency access. Cached information may include the results of I/O-intensive database queries or the results of computationally-intensive calculations.

The Amazon ElastiCache service automates time-consuming management tasks for in-memory cache environments, such as patch management, failure detection, and recovery. It works in conjunction with other Amazon Web Services (such as Amazon EC2, Amazon CloudWatch, and Amazon SNS) to provide a secure, high-performance, and managed in-memory cache. For example, an application running in Amazon EC2 can securely access an Amazon ElastiCache Cluster in the same region with very low latency.

Using the Amazon ElastiCache service, you create a Cache Cluster, which is a collection of one or more Cache Nodes, each running an instance of the Memcached service. A Cache Node is a fixed-size chunk of secure, network-attached RAM. Each Cache Node runs an instance of the Memcached service, and has its own DNS name and port. Multiple types of Cache Nodes are supported, each with varying amounts of associated memory. A Cache Cluster can be set up with a specific number of Cache Nodes and a Cache Parameter Group that controls the properties for each Cache Node. All Cache Nodes within a Cache Cluster are designed to be of the same Node Type and have the same parameter and security group settings.

Amazon ElastiCache allows you to control access to your Cache Clusters using Cache Security Groups. A Cache Security Group acts like a firewall, controlling network access to your Cache Cluster. By default, network access is turned off to



your Cache Clusters. If you want your applications to access your Cache Cluster, you must explicitly enable access from hosts in specific EC2 security groups. Once ingress rules are configured, the same rules apply to all Cache Clusters associated with that Cache Security Group.

To allow network access to your Cache Cluster, create a Cache Security Group and use the Authorize Cache Security Group Ingress API or CLI command to authorize the desired EC2 security group (which in turn specifies the EC2 instances allowed). IP-range based access control is currently not enabled for Cache Clusters. All clients to a Cache Cluster must be within the EC2 network, and authorized via Cache Security Groups.

ElastiCache for Redis provides backup and restore functionality, where you can create a snapshot of your entire Redis cluster as it exists at a specific point in time. You can schedule automatic, recurring daily snapshots or you can create a manual snapshot at any time. For automatic snapshots, you specify a retention period; manual snapshots are retained until you delete them. The snapshots are stored in Amazon S3 with high durability, and can be used for warm starts, backups, and archiving.

Application Services

Amazon Web Services offers a variety of managed services to use with your applications, including services that provide application streaming, queueing, push notification, email delivery, search, and transcoding.

Amazon CloudSearch Security

Amazon CloudSearch is a managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. It enables you to quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.

An Amazon CloudSearch domain encapsulates a collection of data you want to search, the search instances that process your search requests, and a configuration that controls how your data is indexed and searched. You create a separate search domain for each collection of data you want to make searchable. For each domain, you configure indexing options that describe the fields you want to include in your index and how you want to use them, text options that define domain-specific stopwords, stems, and synonyms, rank expressions that you can use to customize how search results are ranked, and access policies that control access to the domain's document and search endpoints.

Access to your search domain's endpoints is restricted by IP address so that only authorized hosts can submit documents and send search requests. IP address authorization is used only to control access to the document and search endpoints. All Amazon CloudSearch configuration requests must be authenticated using standard AWS authentication.

Amazon CloudSearch provides separate endpoints for accessing the configuration, search, and document services:

- The configuration service is accessed through a general endpoint: `cloudsearch.us-east-1.amazonaws.com`
- The document service endpoint is used to submit documents to the domain for indexing and is accessed through a domain-specific endpoint: <http://doc-domainname-domainid.us-east-1.cloudsearch.amazonaws.com>
- The search endpoint is used to submit search requests to the domain and is accessed through a domain-specific endpoint: <http://search-domainname-domainid.us-east-1.cloudsearch.amazonaws.com>

Note that if you do not have a static IP address, you must re-authorize your computer whenever your IP address changes. If your IP address is assigned dynamically, it is also likely that you're sharing that address with other computers on your network. This means that when you authorize the IP address, all computers that share it will be able to access your search domain's document service endpoint.

Like all AWS Services, Amazon CloudSearch requires that every request made to its control API be authenticated so only authenticated users can access and manage your CloudSearch domain. API requests are signed with an HMAC-SHA1 or HMAC-SHA256 signature calculated from the request and the user's AWS Secret Access key. Additionally, the Amazon CloudSearch control API is accessible via SSL-encrypted endpoints. You can control access to Amazon CloudSearch management functions by creating users under your AWS Account using AWS IAM, and controlling which CloudSearch operations these users have permission to perform.



Amazon Simple Queue Service (Amazon SQS) Security

Amazon SQS is a highly reliable, scalable message queuing service that enables asynchronous message-based communication between distributed components of an application. The components can be computers or Amazon EC2 instances or a combination of both. With Amazon SQS, you can send any number of messages to an Amazon SQS queue at any time from any component. The messages can be retrieved from the same component or a different one right away or at a later time (within 4 days). Messages are highly durable; each message is persistently stored in highly available, highly reliable queues. Multiple processes can read/write from/to an Amazon SQS queue at the same time without interfering with each other.

Amazon SQS access is granted based on an AWS Account or a user created with AWS IAM. Once authenticated, the AWS Account has full access to all user operations. An AWS IAM user, however, only has access to the operations and queues for which they have been granted access via policy. By default, access to each individual queue is restricted to the AWS Account that created it. However, you can allow other access to a queue, using either an SQS-generated policy or a policy you write.

Amazon SQS is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2. Data stored within Amazon SQS is not encrypted by AWS; however, the user can encrypt data before it is uploaded to Amazon SQS, provided that the application utilizing the queue has a means to decrypt the message when retrieved. Encrypting messages before sending them to Amazon SQS helps protect against access to sensitive customer data by unauthorized persons, including AWS.

Amazon Simple Notification Service (Amazon SNS) Security

Amazon Simple Notification Service (Amazon SNS) is a web service that makes it easy to set up, operate, and send notifications from the cloud. It provides developers with a highly scalable, flexible, and cost-effective capability to publish messages from an application and immediately deliver them to subscribers or other applications.

Amazon SNS provides a simple web services interface that can be used to create topics that customers want to notify applications (or people) about, subscribe clients to these topics, publish messages, and have these messages delivered over clients' protocol of choice (i.e., HTTP/HTTPS, email, etc.). Amazon SNS delivers notifications to clients using a "push" mechanism that eliminates the need to periodically check or "poll" for new information and updates. Amazon SNS can be leveraged to build highly reliable, event-driven workflows and messaging applications without the need for complex middleware and application management. The potential uses for Amazon SNS include monitoring applications, workflow systems, time-sensitive information updates, mobile applications, and many others. Amazon SNS provides access control mechanisms so that topics and messages are secured against unauthorized access. Topic owners can set policies for a topic that restrict who can publish or subscribe to a topic. Additionally, topic owners can encrypt transmission by specifying that the delivery mechanism must be HTTPS.

Amazon SNS access is granted based on an AWS Account or a user created with AWS IAM. Once authenticated, the AWS Account has full access to all user operations. An AWS IAM user, however, only has access to the operations and topics for which they have been granted access via policy. By default, access to each individual topic is restricted to the AWS Account that created it. However, you can allow other access to SNS, using either an SNS-generated policy or a policy you write.

Amazon Simple Workflow Service (Amazon SWF) Security

The Amazon Simple Workflow Service (SWF) makes it easy to build applications that coordinate work across distributed components. Using Amazon SWF, you can structure the various processing steps in an application as “tasks” that drive work in distributed applications, and Amazon SWF coordinates these tasks in a reliable and scalable manner. Amazon SWF manages task execution dependencies, scheduling, and concurrency based on a developer’s application logic. The service stores tasks, dispatches them to application components, tracks their progress, and keeps their latest state.

Amazon SWF provides simple API calls that can be executed from code written in any language and run on your EC2 instances, or any of your machines located anywhere in the world that can access the Internet. Amazon SWF acts as a coordination hub with which your application hosts interact. You create desired workflows with their associated tasks and any conditional logic you wish to apply and store them with Amazon SWF.

Amazon SWF access is granted based on an AWS Account or a user created with AWS IAM. All actors that participate in the execution of a workflow—deciders, activity workers, workflow administrators—must be IAM users under the AWS Account that owns the Amazon SWF resources. You cannot grant users associated with other AWS Accounts access to your Amazon SWF workflows. An AWS IAM user, however, only has access to the workflows and resources for which they have been granted access via policy.

Amazon Simple Email Service (Amazon SES) Security

Amazon Simple Email Service (SES) is an outbound-only email-sending service built on Amazon’s reliable and scalable infrastructure. Amazon SES helps you maximize email deliverability and stay informed of the delivery status of your emails. Amazon SES integrates with other AWS services, making it easy to send emails from applications being hosted on services such as Amazon EC2.

Unfortunately, with other email systems, it's possible for a spammer to falsify an email header and spoof the originating email address so that it appears as though the email originated from a different source. To mitigate these problems, Amazon SES requires users to verify their email address or domain in order to confirm that they own it and to prevent others from using it. To verify a domain, Amazon SES requires the sender to publish a DNS record that Amazon SES supplies as proof of control over the domain. Amazon SES periodically reviews domain verification status, and revokes verification in cases where it is no longer valid.

Amazon SES takes proactive steps to prevent questionable content from being sent, so that ISPs receive consistently high-quality email from our domains and therefore view Amazon SES as a trusted email origin. Below are some of the features that maximize deliverability and dependability for all of our senders:

- Amazon SES uses content-filtering technologies to help detect and block messages containing viruses or malware before they can be sent.
- Amazon SES maintains complaint feedback loops with major ISPs. Complaint feedback loops indicate which emails a recipient marked as spam. Amazon SES provides you access to these delivery metrics to help guide your sending strategy.
- Amazon SES uses a variety of techniques to measure the quality of each user’s sending. These mechanisms help identify and disable attempts to use Amazon SES for unsolicited mail, and detect other sending patterns that would harm Amazon SES’s reputation with ISPs, mailbox providers, and anti-spam services.

- Amazon SES supports authentication mechanisms such as Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). When you authenticate an email, you provide evidence to ISPs that you own the domain. Amazon SES makes it easy for you to authenticate your emails. If you configure your account to use Easy DKIM, Amazon SES will DKIM-sign your emails on your behalf, so you can focus on other aspects of your email-sending strategy. To ensure optimal deliverability, we recommend that you authenticate your emails.

As with other AWS services, you use security credentials to verify who you are and whether you have permission to interact with Amazon SES. For information about which credentials to use, see [Using Credentials with Amazon SES](#). Amazon SES also integrates with AWS IAM so that you can specify which Amazon SES API actions a user can perform.

If you choose to communicate with Amazon SES through its SMTP interface, you are required to encrypt your connection using TLS. Amazon SES supports two mechanisms for establishing a TLS-encrypted connection: STARTTLS and TLS Wrapper. If you choose to communicate with Amazon SES over HTTP, then all communication will be protected by TLS through Amazon SES's HTTPS endpoint. When delivering email to its final destination, Amazon SES encrypts the email content with opportunistic TLS, if supported by the receiver.

Amazon Elastic Transcoder Service Security

The Amazon Elastic Transcoder service simplifies and automates what is usually a complex process of converting media files from one format, size, or quality to another. The Elastic Transcoder service converts standard-definition (SD) or high-definition (HD) video files as well as audio files. It reads input from an Amazon S3 bucket, transcodes it, and writes the resulting file to another Amazon S3 bucket. You can use the same bucket for input and output, and the buckets can be in any AWS region. The Elastic Transcoder accepts input files in a wide variety of web, consumer, and professional formats. Output file types include the MP3, MP4, OGG, TS, WebM, HLS using MPEG-2 TS, and Smooth Streaming using fmp4 container types, storing H.264 or VP8 video and AAC, MP3, or Vorbis audio.

You'll start with one or more input files, and create transcoding jobs in a type of workflow called a transcoding pipeline for each file. When you create the pipeline you'll specify input and output buckets as well as an IAM role. Each job must reference a media conversion template called a transcoding preset, and will result in the generation of one or more output files. A preset tells the Elastic Transcoder what settings to use when processing a particular input file. You can specify many settings when you create a preset, including the sample rate, bit rate, resolution (output height and width), the number of reference and keyframes, a video bit rate, some thumbnail creation options, etc.

A best effort is made to start jobs in the order in which they're submitted, but this is not a hard guarantee and jobs typically finish out of order since they are worked on in parallel and vary in complexity. You can pause and resume any of your pipelines if necessary.

Elastic Transcoder supports the use of SNS notifications when it starts and finishes each job, and when it needs to tell you that it has detected an error or warning condition. The SNS notification parameters are associated with each pipeline. It can also use the List Jobs By Status function to find all of the jobs with a given status (e.g., "Completed") or the Read Job function to retrieve detailed information about a particular job.

Like all other AWS services, Elastic Transcoder integrates with AWS Identity and Access Management (IAM), which allows you to control access to the service and to other AWS resources that Elastic Transcoder requires, including Amazon S3 buckets and Amazon SNS topics. By default, IAM users have no access to Elastic Transcoder or to the resources that it uses. If you want IAM users to be able to work with Elastic Transcoder, you must explicitly grant them permissions.



Amazon Elastic Transcoder requires every request made to its control API be authenticated so only authenticated processes or users can create, modify, or delete their own Amazon Transcoder pipelines and presets. Requests are signed with an HMAC-SHA256 signature calculated from the request and a key derived from the user's secret key. Additionally, the Amazon Elastic Transcoder API is only accessible via SSL-encrypted endpoints.

Durability is provided by Amazon S3, where media files are redundantly stored on multiple devices across multiple facilities in an Amazon S3 region. For added protection against users accidentally deleting media files, you can use the Versioning feature in Amazon S3 to preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket. You can further protect versions using Amazon S3 Versioning's MFA Delete feature. Once enabled for an Amazon S3 bucket, each version deletion request must include the six-digit code and serial number from your multi-factor authentication device.

Amazon AppStream Security

The Amazon AppStream service provides a framework for running streaming applications, particularly applications that require lightweight clients running on mobile devices. It enables you to store and run your application on powerful, parallel-processing GPUs in the cloud and then stream input and output to any client device. This can be a pre-existing application that you modify to work with Amazon AppStream or a new application that you design specifically to work with the service.

The Amazon AppStream SDK simplifies the development of interactive streaming applications and client applications. The SDK provides APIs that connect your customers' devices directly to your application, capture and encode audio and video, stream content across the Internet in near real-time, decode content on client devices, and return user input to the application. Because your application's processing occurs in the cloud, it can scale to handle extremely large computational loads.

Amazon AppStream deploys streaming applications on Amazon EC2. When you add a streaming application through the AWS Management Console, the service creates the AMI required to host your application and makes your application available to streaming clients. The service scales your application as needed within the capacity limits you have set to meet demand. Clients using the Amazon AppStream SDK automatically connect to your streamed application.

In most cases, you'll want to ensure that the user running the client is authorized to use your application before letting him obtain a session ID. We recommend that you use some sort of entitlement service, which is a service that authenticates clients and authorizes their connection to your application. In this case, the entitlement service will also call into the Amazon AppStream REST API to create a new streaming session for the client. After the entitlement service creates a new session, it returns the session identifier to the authorized client as a single-use entitlement URL. The client then uses the entitlement URL to connect to the application. Your entitlement service can be hosted on an Amazon EC2 instance or on [AWS Elastic Beanstalk](#).

Amazon AppStream utilizes an AWS CloudFormation template that automates the process of deploying a GPU EC2 instance that has the AppStream Windows Application and Windows Client SDK libraries installed; is configured for SSH, RDC, or VPN access; and has an elastic IP address assigned to it. By using this template to deploy your standalone streaming server, all you need to do is upload your application to the server and run the command to launch it. You can then use the Amazon AppStream Service Simulator tool to test your application in standalone mode before deploying it into production.

Amazon AppStream also utilizes the STX Protocol to manage the streaming of your application from AWS to local devices. The Amazon AppStream STX Protocol is a proprietary protocol used to stream high-quality application video



over varying network conditions; it monitors network conditions and automatically adapts the video stream to provide a low-latency and high-resolution experience to your customers. It minimizes latency while syncing audio and video as well as capturing input from your customers to be sent back to the application running in AWS.

Analytics Services

Amazon Web Services provides cloud-based analytics services to help you process and analyze any volume of data, whether your need is for managed Hadoop clusters, real-time streaming data, petabyte scale data warehousing, or orchestration.

Amazon Elastic MapReduce (Amazon EMR) Security

Amazon Elastic MapReduce (Amazon EMR) is a managed web service you can use to run Hadoop clusters that process vast amounts of data by distributing the work and data among several servers. It utilizes an enhanced version of the Apache Hadoop framework running on the web-scale infrastructure of Amazon EC2 and Amazon S3. You simply upload your input data and a data processing application into Amazon S3. Amazon EMR then launches the number of Amazon EC2 instances you specify. The service begins the job flow execution while pulling the input data from Amazon S3 into the launched Amazon EC2 instances. Once the job flow is finished, Amazon EMR transfers the output data to Amazon S3, where you can then retrieve it or use it as input in another job flow.

When launching job flows on your behalf, Amazon EMR sets up two Amazon EC2 security groups: one for the master nodes and another for the slaves. The master security group has a port open for communication with the service. It also has the SSH port open to allow you to SSH into the instances, using the key specified at startup. The slaves start in a separate security group, which only allows interaction with the master instance. By default both security groups are set up to not allow access from external sources, including Amazon EC2 instances belonging to other customers. Since these are security groups within your account, you can reconfigure them using the standard EC2 tools or dashboard. To protect customer input and output datasets, Amazon EMR transfers data to and from Amazon S3 using SSL.

Amazon EMR provides several ways to control access to the resources of your cluster. You can use AWS IAM to create user accounts and roles and configure permissions that control which AWS features those users and roles can access. When you launch a cluster, you can associate an Amazon EC2 key pair with the cluster, which you can then use when you connect to the cluster using SSH. You can also set permissions that allow users other than the default Hadoop user to submit jobs to your cluster.

By default, if an IAM user launches a cluster, that cluster is hidden from other IAM users on the AWS account. This filtering occurs on all Amazon EMR interfaces—the console, CLI, API, and SDKs—and helps prevent IAM users from accessing and inadvertently changing clusters created by other IAM users. It is useful for clusters that are intended to be viewed by only a single IAM user and the main AWS account. You also have the option to make a cluster visible and accessible to all IAM users under a single AWS account.

For an additional layer of protection, you can launch the EC2 instances of your EMR cluster into an Amazon VPC, which is like launching it into a private subnet. This allows you to control access to the entire subnetwork. You can also launch the cluster into a VPC and enable the cluster to access resources on your internal network using a VPN connection. You can encrypt the input data before you upload it to Amazon S3 using any common data encryption tool. If you do encrypt the data before it's uploaded, you then need to add a decryption step to the beginning of your job flow when Amazon Elastic MapReduce fetches the data from Amazon S3.



Amazon Kinesis Security

Amazon Kinesis is a managed service designed to handle real-time streaming of big data. It can accept any amount of data, from any number of sources, scaling up and down as needed. You can use Kinesis in situations that call for large-scale, real-time data ingestion and processing, such as server logs, social media or market data feeds, and web clickstream data.

Applications read and write data records to Amazon Kinesis in *streams*. You can create any number of Kinesis streams to capture, store, and transport data. Amazon Kinesis automatically manages the infrastructure, storage, networking, and configuration needed to collect and process your data at the level of throughput your streaming applications need. You don't have to worry about provisioning, deployment, or ongoing-maintenance of hardware, software, or other services to enable real-time capture and storage of large-scale data. Amazon Kinesis also synchronously replicates data across three facilities in an AWS Region, providing high availability and data durability.

In Amazon Kinesis, data records contain a sequence number, a partition key, and a data blob, which is an un-interpreted, immutable sequence of bytes. The Amazon Kinesis service does not inspect, interpret, or change the data in the blob in any way. Data records are accessible for only 24 hours from the time they are added to an Amazon Kinesis stream, and then they are automatically discarded.

Your application is a consumer of an Amazon Kinesis stream, which typically runs on a fleet of Amazon EC2 instances. A Kinesis application uses the Amazon Kinesis Client Library to read from the Amazon Kinesis stream. The Kinesis Client Library takes care of a variety of details for you including failover, recovery, and load balancing, allowing your application to focus on processing the data as it becomes available. After processing the record, your consumer code can pass it along to another Kinesis stream; write it to an [Amazon S3](#) bucket, a [Redshift](#) data warehouse, or a [DynamoDB](#) table; or simply discard it. A connector library is available to help you integrate Kinesis with other AWS services (such as DynamoDB, Redshift, and Amazon S3) as well as third-party products like Apache Storm.

You can control logical access to Kinesis resources and management functions by creating users under your AWS Account using AWS IAM, and controlling which Kinesis operations these users have permission to perform. To facilitate running your producer or consumer applications on an Amazon EC2 instance, you can configure that instance with an IAM role. That way, AWS credentials that reflect the permissions associated with the IAM role are made available to applications on the instance, which means you don't have to use your long-term AWS security credentials. Roles have the added benefit of providing temporary credentials that expire within a short timeframe, which adds an additional measure of protection. See the [Using IAM](#) guide for more information about IAM roles.

The Amazon Kinesis API is only accessible via an SSL-encrypted endpoint (kinesis.us-east-1.amazonaws.com) to help ensure secure transmission of your data to AWS. You must connect to that endpoint to access Kinesis, but you can then use the API to direct AWS Kinesis to create a stream in any AWS Region

AWS Data Pipeline Security

The AWS Data Pipeline service helps you process and move data between different data sources at specified intervals using data-driven workflows and built-in dependency checking. When you create a pipeline, you define data sources, preconditions, destinations, processing steps, and an operational schedule. Once you define and activate a pipeline, it will run automatically according to the schedule you specified.

With AWS Data Pipeline, you don't have to worry about checking resource availability, managing inter-task dependencies, retrying transient failures/timeouts in individual tasks, or creating a failure notification system. AWS Data



Pipeline takes care of launching the AWS services and resources your pipeline needs to process your data (e.g., Amazon EC2 or EMR) and transferring the results to storage (e.g., Amazon S3, RDS, DynamoDB, or EMR).

When you use the console, AWS Data Pipeline creates the necessary IAM roles and policies, including a trusted entities list for you. IAM roles determine what your pipeline can access and the actions it can perform. Additionally, when your pipeline creates a resource, such as an EC2 instance, IAM roles determine the EC2 instance's permitted resources and actions. When you create a pipeline, you specify one IAM role that governs your pipeline and another IAM role to govern your pipeline's resources (referred to as a "resource role"), which can be the same role for both. As part of the security best practice of least privilege, we recommend that you consider the minimum permissions necessary for your pipeline to perform work and define the IAM roles accordingly.

Like most AWS services, AWS Data Pipeline also provides the option of secure (HTTPS) endpoints for access via SSL.

Deployment and Management Services

Amazon Web Services provides a variety of tools to help with the deployment and management of your applications. This includes services that allow you to create individual user accounts with credentials for access to AWS services. It also includes services for creating and updating stacks of AWS resources, deploying applications on those resources, and monitoring the health of those AWS resources. Other tools help you manage cryptographic keys using hardware security modules (HSMs) and log AWS API activity for security and compliance purposes.

AWS Identity and Access Management (AWS IAM)

AWS IAM allows you to create multiple users and manage the permissions for each of these users within your AWS Account. A user is an identity (within an AWS Account) with unique security credentials that can be used to access AWS Services. AWS IAM eliminates the need to share passwords or keys, and makes it easy to enable or disable a user's access as appropriate.

AWS IAM enables you to implement security best practices, such as least privilege, by granting unique credentials to every user within your AWS Account and only granting permission to access the AWS services and resources required for the users to perform their jobs. AWS IAM is secure by default; new users have no access to AWS until permissions are explicitly granted.

AWS IAM is also integrated with the AWS Marketplace, so that you can control who in your organization can subscribe to the software and services offered in the Marketplace. Since subscribing to certain software in the Marketplace launches an EC2 instance to run the software, this is an important access control feature. Using AWS IAM to control access to the AWS Marketplace also enables AWS Account owners to have fine-grained control over usage and software costs.

AWS IAM enables you to minimize the use of your AWS Account credentials. Once you create AWS IAM user accounts, all interactions with AWS Services and resources should occur with AWS IAM user security credentials. More information about AWS IAM is available on the AWS website: <http://aws.amazon.com/iam/>

Roles

An IAM *role* uses temporary security credentials to allow you to delegate access to users or services that normally don't have access to your AWS resources. A role is a set of permissions to access specific AWS resources, but these permissions are not tied to a specific IAM user or group. An authorized entity (e.g., mobile user, EC2 instance) assumes a role and receives temporary security credentials for authenticating to the resources defined in the role. Temporary



security credentials provide enhanced security due to their short life-span (the default expiration is 12 hours) and the fact that they cannot be reused after they expire. This can be particularly useful in providing limited, controlled access in certain situations:

- **Federated (non-AWS) User Access.** Federated users are users (or applications) who do not have AWS Accounts. With roles, you can give them access to your AWS resources for a limited amount of time. This is useful if you have non-AWS users that you can authenticate with an external service, such as Microsoft Active Directory, LDAP, or Kerberos. The temporary AWS credentials used with the roles provide identity federation between AWS and your non-AWS users in your corporate identity and authorization system.

If your organization supports SAML 2.0 (Security Assertion Markup Language 2.0), you can create trust between your organization as an identity provider (IdP) and other organizations as service providers. In AWS, you can configure AWS as the service provider and use SAML to provide your users with federated single-sign on (SSO) to the AWS Management Console or to get federated access to call AWS APIs.

Roles are also useful if you create a mobile or web-based application that accesses AWS resources. AWS resources require security credentials for programmatic requests; however, you shouldn't embed long-term security credentials in your application because they are accessible to the application's users and can be difficult to rotate. Instead, you can let users sign in to your application using Login with Amazon, Facebook, or Google, and then use their authentication information to assume a role and get temporary security credentials.

- **Cross-Account Access.** For organizations who use multiple AWS Accounts to manage their resources, you can set up roles to provide users who have permissions in one account to access resources under another account. For organizations who have personnel who only rarely need access to resources under another account, using roles helps ensure that credentials are provided temporarily, only as needed.
- **Applications Running on EC2 Instances that Need to Access AWS Resources.** If an application runs on an Amazon EC2 instance and needs to make requests for AWS resources such as Amazon S3 buckets or an DynamoDB table, it must have security credentials. Using roles instead of creating individual IAM accounts for each application on each instance can save significant time for customers who manage a large number of instances or an elastically scaling fleet using AWS Auto Scaling.

The temporary credentials include a security token, an Access Key ID, and a Secret Access Key. To give a user access to certain resources, you distribute the temporary security credentials to the user you are granting temporary access to. When the user makes calls to your resources, the user passes in the token and Access Key ID, and signs the request with the Secret Access Key. The token will not work with different access keys. How the user passes in the token depends on the API and version of the AWS product the user is making calls to. More information about temporary security credentials is available on the AWS website: <http://docs.amazonwebservices.com/STS>

The use of temporary credentials means additional protection for you because you don't have to manage or distribute long-term credentials to temporary users. In addition, the temporary credentials get automatically loaded to the target instance so you don't have to embed them somewhere unsafe like your code. Temporary credentials are automatically rotated or changed multiple times a day without any action on your part, and are stored securely by default.

More information about using IAM roles to auto-provision keys on EC2 instances is available in the *Using IAM* guide on the AWS website: <http://docs.amazonwebservices.com/IAM>

Amazon CloudWatch Security

Amazon CloudWatch is a web service that provides monitoring for AWS cloud resources, starting with Amazon EC2. It provides customers with visibility into resource utilization, operational performance, and overall demand patterns—



including metrics such as CPU utilization, disk reads and writes, and network traffic. You can set up CloudWatch alarms to notify you if certain thresholds are crossed, or to take other automated actions such as adding or removing EC2 instances if Auto-Scaling is enabled.

CloudWatch captures and summarizes utilization metrics natively for AWS resources, but you can also have other logs sent to CloudWatch to monitor. You can route your guest OS, application, and custom log files for the software installed on your EC2 instances to CloudWatch, where they will be stored in durable fashion for as long as you'd like. You can configure CloudWatch to monitor the incoming log entries for any desired symbols or messages and to surface the results as CloudWatch metrics. You could, for example, monitor your web server's log files for 404 errors to detect bad inbound links or invalid user messages to detect unauthorized login attempts to your guest OS.

Like all AWS Services, Amazon CloudWatch requires that every request made to its control API be authenticated so only authenticated users can access and manage CloudWatch. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Additionally, the Amazon CloudWatch control API is only accessible via SSL-encrypted endpoints.

You can further control access to Amazon CloudWatch by creating users under your AWS Account using AWS IAM, and controlling what CloudWatch operations these users have permission to call.

AWS CloudHSM Security

The AWS CloudHSM service provides customers with dedicated access to a hardware security module (HSM) appliance designed to provide secure cryptographic key storage and operations within an intrusion-resistant, tamper-evident device. You can generate, store, and manage the cryptographic keys used for data encryption so that they are accessible only by you. AWS CloudHSM appliances are designed to securely store and process cryptographic key material for a wide variety of uses such as database encryption, Digital Rights Management (DRM), Public Key Infrastructure (PKI), authentication and authorization, document signing, and transaction processing. They support some of the strongest cryptographic algorithms available, including AES, RSA, and ECC, and many others.

The AWS CloudHSM service is designed to be used with Amazon EC2 and VPC, providing the appliance with its own private IP within a private subnet. You can connect to CloudHSM appliances from your EC2 servers through SSL/TLS, which uses two-way digital certificate authentication and 256-bit SSL encryption to provide a secure communication channel.

Selecting CloudHSM service in the same region as your EC2 instance decreases network latency, which can improve your application performance. You can configure a client on your EC2 instance that allows your applications to use the APIs provided by the HSM, including PKCS#11, MS CAPI and Java JCA/JCE (Java Cryptography Architecture/Java Cryptography Extensions).

Before you begin using an HSM, you must set up at least one partition on the appliance. A cryptographic partition is a logical and physical security boundary that restricts access to your keys, so only you control your keys and the operations performed by the HSM. AWS has administrative credentials to the appliance, but these credentials can only be used to manage the appliance, not the HSM partitions on the appliance. AWS uses these credentials to monitor and maintain the health and availability of the appliance. AWS cannot extract your keys nor can AWS cause the appliance to perform any cryptographic operation using your keys.

The HSM appliance has both physical and logical tamper detection and response mechanisms that erase the cryptographic key material and generate event logs if tampering is detected. The HSM is designed to detect tampering if



the physical barrier of the HSM appliance is breached. In addition, after three unsuccessful attempts to access an HSM partition with HSM Admin credentials, the HSM appliance erases its HSM partitions.

When your CloudHSM subscription ends and you have confirmed that the contents of the HSM are no longer needed, you must delete each partition and its contents as well as any logs. As part of the decommissioning process, AWS zeroizes the appliance, permanently erasing all key material.

AWS CloudTrail Security

AWS CloudTrail provides a log of all requests for AWS resources within your account. For each event recorded, you can see what service was accessed, what action was performed, any parameters for the action, and who made the request. Not only can you see which one of your users or services performed an action on an AWS service, but you can see whether it was as the AWS root account user or an IAM user, or whether it was with temporary security credentials for a role or federated user.

CloudTrail basically captures information about every API call to an AWS resource, whether that call was made from the AWS Management Console, CLI, or an SDK. If the API request returned an error, CloudTrail provides the description of the error, including messages for authorization failures. It even captures AWS Management Console sign-in events, creating a log record every time an AWS account owner, a federated user, or an IAM user simply signs into the console.

Once you have enabled CloudTrail, event logs are delivered every 5 minutes to the Amazon S3 bucket of your choice. The log files are organized by AWS Account ID, region, service name, date, and time. You can configure CloudTrail so that it aggregates log files from multiple regions into a single Amazon S3 bucket. From there, you can then upload them to your favorite log management and analysis solutions to perform security analysis and detect user behavior patterns.

By default, log files are stored indefinitely. The log files are automatically encrypted using Amazon [S3's Server Side Encryption](#) and will remain in the bucket until you choose to delete or archive them. You can use Amazon S3 lifecycle configuration rules to automatically delete old log files or archive them to Amazon Glacier for additional longevity at significant savings.

Like every other AWS service, you can limit access to CloudTrail to only certain users. You can use IAM to control which AWS users can create, configure, or delete AWS CloudTrail trails as well as which users can start and stop logging. You can control access to the log files by applying IAM or Amazon S3 bucket policies. You can also add an additional layer of security by enabling [MFA Delete](#) on your Amazon S3 bucket.

Mobile Services

AWS mobile services make it easier for you to build, ship, run, monitor, optimize, and scale cloud-powered applications for mobile devices. These services also help you authenticate users to your mobile application, synchronize data, and collect and analyze application usage.

Amazon Cognito

Amazon Cognito provides identity and sync services for mobile and web-based applications. It simplifies the task of authenticating users and storing, managing, and syncing their data across multiple devices, platforms, and applications. It provides temporary, limited-privilege credentials for both authenticated and unauthenticated users without having to manage any backend infrastructure.



Cognito works with well-known identity providers like Google, Facebook, and Amazon to authenticate end users of your mobile and web applications. You can take advantage of the identification and authorization features provided by these services instead of having to build and maintain your own. Your application authenticates with one of these identity providers using the provider's SDK. Once the end user is authenticated with the provider, an OAuth or OpenID Connect token returned from the provider is passed by your application to Cognito, which returns a new Cognito ID for the user and a set of temporary, limited-privilege AWS credentials.

To begin using Amazon Cognito, you create an *identity pool* through the Amazon Cognito console. The identity pool is a store of user identity information that is specific to your AWS account. During the creation of the identity pool, you will be asked to create a new [IAM role](#) or pick an existing one for your end users. An IAM role is a set of permissions to access specific AWS resources, but these permissions are not tied to a specific IAM user or group. An authorized entity (e.g., mobile user, EC2 instance) assumes a role and receives temporary security credentials for authenticating to the AWS resources defined in the role. Temporary security credentials provide enhanced security due to their short life-span (the default expiration is 12 hours) and the fact that they cannot be reused after they expire. The role you select has an impact on which AWS services your end users will be able to access with the temporary credentials. By default, Amazon Cognito creates a new role with limited permissions – end users only have access to the Cognito Sync service and Amazon Mobile Analytics. If your application needs access to other AWS resources such as Amazon S3 or DynamoDB, you can modify your roles directly from the IAM management console.

With Amazon Cognito, there's no need to create individual AWS accounts or even IAM accounts for every one of your web/mobile app's end users who will need to access your AWS resources. In conjunction with IAM roles, mobile users can securely access AWS resources and application features, and even save data to the AWS cloud without having to create an account or log in. However, if they choose to do this later, Cognito will merge data and identification information.

Because Amazon Cognito stores data locally as well as in the service, your end users can continue to interact with their data even when they are offline. Their offline data may be stale, but anything they put into the dataset, they can immediately retrieve whether they are online or not. The client SDK manages a local SQLite store so that the application can work even when it is not connected. The SQLite store functions as a cache and is the target of all read and write operations. Cognito's sync facility compares the local version of the data to the cloud version, and pushes up or pulls down deltas as needed. Note that in order to sync data across devices, your identity pool must support authenticated identities. Unauthenticated identities are tied to the device, so unless an end user authenticates, no data can be synced across multiple devices.

With Cognito, your application communicates directly with a supported public identity provider (Amazon, Facebook, or Google) to authenticate users. Amazon Cognito does not receive or store user credentials—only the OAuth or OpenID Connect token received from the identity provider. Once Cognito receives the token, it returns a new Cognito ID for the user and a set of temporary, limited-privilege AWS credentials.

Each Cognito identity has access only to its own data in the sync store, and this data is encrypted when stored. In addition, all identity data is transmitted over HTTPS. The unique Amazon Cognito identifier on the device is stored in the appropriate secure location—on iOS for example, the Cognito identifier is stored in the iOS keychain. User data is cached in a local SQLite database within the application's sandbox; if you require additional security, you can encrypt this identity data in the local cache by implementing encryption in your application.



Amazon Mobile Analytics

Amazon Mobile Analytics is a service for collecting, visualizing, and understanding mobile application usage data. It enables you to track customer behaviors, aggregate metrics, and identify meaningful patterns in your mobile applications. Amazon Mobile Analytics automatically calculates and updates usage metrics as the data is received from client devices running your app and displays the data in the console.

You can integrate Amazon Mobile Analytics with your application without requiring users of your app to be authenticated with an identity provider (like Google, Facebook, or Amazon). For these unauthenticated users, Mobile Analytics works with Amazon Cognito to provide temporary, limited-privilege credentials. To do this, you first create an identity pool in Cognito. The identity pool will use IAM *roles*, which is a set of permissions not tied to a specific IAM user or group but which allows an entity to access specific AWS resources. The entity assumes a role and receives temporary security credentials for authenticating to the AWS resources defined in the role. By default, Amazon Cognito creates a new role with limited permissions – end users only have access to the Cognito Sync service and Amazon Mobile Analytics. If your application needs access to other AWS resources such as Amazon S3 or DynamoDB, you can modify your roles directly from the IAM management console.

You can integrate the AWS Mobile SDK for Android or iOS into your application or use the Amazon Mobile Analytics REST API to send events from any connected device or service and visualize data in the reports. The Amazon Mobile Analytics API is only accessible via an SSL-encrypted endpoint (<https://mobileanalytics.us-east-1.amazonaws.com>).

Applications

AWS applications are managed services that enable you to provide your users with secure, centralized storage and work areas in the cloud.

Amazon WorkSpaces

Amazon WorkSpaces is a managed desktop service that allows you to quickly provision cloud-based desktops for your users. Simply choose a Windows 7 bundle that best meets the needs of your users and the number of WorkSpaces that you would like to launch. Once the WorkSpaces are ready, users receive an email informing them where they can download the relevant client and log into their WorkSpace. They can then access their cloud-based desktops from a variety of endpoint devices, including PCs, laptops, and mobile devices. However, your organization's data is never sent to or stored on the end-user device because Amazon WorkSpaces uses PC-over-IP ([PCoIP](#)), which provides an interactive video stream without transmitting actual data. The PCoIP protocol compresses, encrypts, and encodes the users' desktop computing experience and transmits 'pixels only' across any standard IP network to end-user devices.

In order to access their WorkSpace, users must sign in using a set of unique credentials or their regular Active Directory credentials. When you integrate Amazon WorkSpaces with your corporate Active Directory, each WorkSpace joins your Active Directory domain and can be managed just like any other desktop in your organization. This means that you can use Active Directory Group Policies to manage your users' WorkSpaces to specify configuration options that control the desktop. If you choose not to use Active Directory or other type of on-premises directory to manage your user WorkSpaces, you can create a private cloud directory within Amazon WorkSpaces that you can use for administration.

To provide an additional layer of security, you can also require the use of multi-factor authentication upon signin in the form of a hardware or software token. Amazon WorkSpaces supports MFA using an on-premise Remote Authentication Dial In User Service (RADIUS) server or any security provider that supports RADIUS authentication. It currently supports the PAP, CHAP, MS-CHAP1, and MS-CHAP2 protocols, along with RADIUS proxies.



Each Workspace resides on its own EC2 instance within a VPC. You can create WorkSpaces in a VPC you already own or have the WorkSpaces service create one for you automatically using the WorkSpaces Quick Start option. When you use the Quick Start option, WorkSpaces not only creates the VPC, but it performs several other provisioning and configuration tasks for you, such as creating an Internet Gateway for the VPC, setting up a directory within the VPC that is used to store user and Workspace information, creating a directory administrator account, creating the specified user accounts and adding them to the directory, and creating the Workspace instances. Or the VPC can be connected to an on-premises network using a secure VPN connection to allow access to an existing on-premises Active Directory and other intranet resources. You can add a security group that you create in your Amazon VPC to all the WorkSpaces that belong to your Directory. This allows you to control network access from Amazon WorkSpaces in your VPC to other resources in your Amazon VPC and on-premises network.

Persistent storage for WorkSpaces is provided by Amazon EBS and is automatically backed up twice a day to Amazon S3. If WorkSpaces Sync is enabled on a Workspace, the folder a user chooses to sync will be continuously backed up and stored in Amazon S3. You can also use WorkSpaces Sync on a Mac or PC to sync documents to or from your Workspace so that you can always have access to your data regardless of the desktop computer you are using.

Because it's a managed service, AWS takes care of several security and maintenance tasks like daily backups and patching. Updates are delivered automatically to your WorkSpaces during a weekly maintenance window. You can control how patching is configured for a user's Workspace. By default, Windows Update is turned on, but you have the ability to customize these settings, or use an alternative patch management approach if you desire. For the underlying OS, Windows Update is enabled by default on WorkSpaces, and configured to install updates on a weekly basis. You can use an alternative patching approach or to configure Windows Update to perform updates at a time of your choosing.

You can use IAM to control who on your team can perform administrative functions like creating or deleting WorkSpaces or setting up user directories. You can also set up a Workspace for directory administration, install your favorite Active Directory administration tools, and create organizational units and Group Policies in order to more easily apply Active Directory changes for all your WorkSpaces users.

Amazon WorkDocs

Amazon WorkDocs is a managed enterprise storage and sharing service with feedback capabilities for user collaboration. Users can store any type of file in a WorkDocs folder and allow others to view and download them. Commenting and annotation capabilities work on certain file types such as MS Word, and without requiring the application that was used to originally create the file. WorkDocs notifies contributors about review activities and deadlines via email and performs versioning of files that you have synced using the WorkDocs Sync application.

User information is stored in an Active Directory-compatible network directory. You can either create a new directory in the cloud, or connect Amazon WorkDocs to your on-premises directory. When you create a cloud directory using WorkDocs' quick start setup, it also creates a directory administrator account with the administrator email as the username. An email is sent to your administrator with instructions to complete registration. The administrator then uses this account to manage your directory.

When you create a cloud directory using WorkDocs' quick start setup, it also creates and configures a VPC for use with the directory. If you need more control over the directory configuration, you can choose the standard setup, which allows you to specify your own directory domain name, as well as one of your existing VPCs to use with the directory. If you want to use one of your existing VPCs, the VPC must have an Internet gateway and at least two subnets. Each of the subnets must be in a different Availability Zone.



Using the Amazon WorkDocs Management Console, administrators can view audit logs to track file and user activity by time, IP address, and device, and choose whether to allow users to share files with others outside their organization. Users can then control who can access individual files and disable downloads of files they share.

All data in transit is encrypted using industry-standard SSL. The WorkDocs web and mobile applications and desktop sync clients transmit files directly to Amazon WorkDocs using SSL. WorkDocs users can also utilize Multi-Factor Authentication, or MFA, if their organization has deployed a Radius server. MFA uses the following factors: username, password, and methods supported by the Radius server. The protocols supported are PAP, CHAP, MS-CHAPv1, and MS-CHAPv2.

You choose the AWS Region where each WorkDocs site's files are stored. Amazon WorkDocs is currently available in the US-East (Virginia), US-West (Oregon), and EU (Ireland) AWS Regions. All files, comments, and annotations stored in WorkDocs are automatically encrypted with AES-256 encryption.

Appendix – Glossary of Terms

Access Key ID: A string that AWS distributes in order to uniquely identify each AWS user; it is an alphanumeric token associated with your Secret Access Key.

Access control list (ACL): A list of permissions or rules for accessing an object or network resource. In Amazon EC2, security groups act as ACLs at the instance level, controlling which users have permission to access specific instances. In Amazon S3, you can use ACLs to give read or write access on buckets or objects to groups of users. In Amazon VPC, ACLs act like network firewalls and control access at the subnet level.

AMI: An Amazon Machine Image (AMI) is an encrypted machine image stored in Amazon S3. It contains all the information necessary to boot instances of a customer's software.

API: Application Programming Interface (API) is an interface in computer science that defines the ways by which an application program may request services from libraries and/or operating systems.

Archive: An archive in Amazon Glacier is a file that you want to store and is a base unit of storage in Amazon Glacier. It can be any data such as a photo, video, or document. Each archive has a unique ID and an optional description.

Authentication: Authentication is the process of determining whether someone or something is, in fact, who or what it is declared to be. Not only do users need to be authenticated, but every program that wants to call the functionality exposed by an AWS API must be authenticated. AWS requires that you authenticate every request by digitally signing it using a cryptographic hash function.

Auto-Scaling: An AWS service that allows customers to automatically scale their Amazon EC2 capacity up or down according to conditions they define.

Availability Zone: Amazon EC2 locations are composed of regions and availability zones. Availability zones are distinct locations that are engineered to be insulated from failures in other availability zones and provide inexpensive, low latency network connectivity to other availability zones in the same region.

Bastion host: A computer specifically configured to withstand attack, usually placed on the external/public side of a demilitarized zone (DMZ) or outside the firewall. You can set up an Amazon EC2 instance as an SSH bastion by setting up a public subnet as part of an Amazon VPC.

Bucket: A container for objects stored in Amazon S3. Every object is contained within a bucket. For example, if the object named photos/puppy.jpg is stored in the johnsmith bucket, then it is addressable using the URL <http://johnsmith.s3.amazonaws.com/photos/puppy.jpg>.

Certificate: A credential that some AWS products use to authenticate AWS Accounts and users. Also known as an X.509 certificate. The certificate is paired with a private key.

CIDR Block: Classless Inter-Domain Routing Block of IP addresses.

Client-side encryption: Encrypting data on the client side before uploading it to Amazon S3.

CloudFormation: An AWS provisioning tool that lets customers record the baseline configuration of the AWS resources needed to run their applications so that they can provision and update them in an orderly and predictable fashion.



Cognito: An AWS service that simplifies the task of authenticating users and storing, managing, and syncing their data across multiple devices, platforms, and applications. It works with multiple existing identity providers and also supports unauthenticated guest users.

Credentials: Items that a user or process must have in order to confirm to AWS services during the authentication process that they are authorized to access the service. AWS credentials include passwords, secret access keys as well as X.509 certificates and multi-factor tokens.

Dedicated instance: Amazon EC2 instances that are physically isolated at the host hardware level (i.e., they will run on single-tenant hardware).

Digital signature: A digital signature is a cryptographic method for demonstrating the authenticity of a digital message or document. A valid digital signature gives a recipient reason to believe that the message was created by an authorized sender, and that it was not altered in transit. Digital signatures are used by customers for signing requests to AWS APIs as part of the authentication process.

Direct Connect Service: Amazon service that allows you to provision a direct link between your internal network and an AWS region using a high-throughput, dedicated connection. With this dedicated connection in place, you can then create logical connections directly to the AWS cloud (for example, to Amazon EC2 and Amazon S3) and Amazon VPC, bypassing Internet service providers in the network path.

DynamoDB Service: A managed NoSQL database service from AWS that provides fast and predictable performance with seamless scalability.

EBS: Amazon Elastic Block Store (EBS) provides block-level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance.

ElastiCache: An AWS web service that allows you to set up, manage, and scale distributed in-memory cache environments in the cloud. The service improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases.

Elastic Beanstalk: An AWS deployment and management tool that automates the functions of capacity provisioning, load balancing, and auto scaling for customers' applications.

Elastic IP Address: A static, public IP address that you can assign to any instance in an Amazon VPC, thereby making the instance public. Elastic IP addresses also enable you to mask instance failures by rapidly remapping your public IP addresses to any instance in the VPC.

Elastic Load Balancing: An AWS service that is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits such as taking over the encryption/decryption work from EC2 instances and managing it centrally on the load balancer.

Elastic MapReduce (EMR) Service: An AWS service that utilizes a hosted Hadoop framework running on the web-scale infrastructure of Amazon EC2 and Amazon S3. Elastic MapReduce enables customers to easily and cost-effectively process extremely large quantities of data ("big data").



Elastic Network Interface: Within an Amazon VPC, an Elastic Network Interface is an optional second network interface that you can attach to an EC2 instance. An Elastic Network Interface can be useful for creating a management network or using network or security appliances in the Amazon VPC. It can be easily detached from an instance and reattached to another instance.

Endpoint: A URL that is the entry point for an AWS service. To reduce data latency in your applications, most AWS services allow you to select a regional endpoint to make your requests. Some web services allow you to use a general endpoint that doesn't specify a region; these generic endpoints resolve to the service's us-east-1 endpoint. You can connect to an AWS endpoint via HTTP or secure HTTP (HTTPS) using SSL.

Federated users: User, systems, or applications that are not currently authorized to access your AWS services, but that you want to give temporary access to. This access is provided using the AWS Security Token Service (STS) APIs.

Firewall: A hardware or software component that controls incoming and/or outgoing network traffic according to a specific set of rules. Using firewall rules in Amazon EC2, you specify the protocols, ports, and source IP address ranges that are allowed to reach your instances. These rules specify which incoming network traffic should be delivered to your instance (e.g., accept web traffic on port 80). Amazon VPC supports a complete firewall solution enabling filtering on both ingress and egress traffic from an instance. The default group enables inbound communication from other members of the same group and outbound communication to any destination. Traffic can be restricted by any IP protocol, by service port, as well as source/destination IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

Guest OS: In a virtual machine environment, multiple operating systems can run on a single piece of hardware. Each one of these instances is considered a guest on the host hardware and utilizes its own OS.

Hash: A cryptographic hash function is used to calculate a digital signature for signing requests to AWS APIs. A cryptographic hash is a one-way function that returns a unique hash value based on the input. The input to the hash function includes the text of your request and your secret access key. The hash function returns a hash value that you include in the request as your signature.

HMAC-SHA1/HMAC-SHA256: In cryptography, a keyed-Hash Message Authentication Code (HMAC or KMAC), is a type of message authentication code (MAC) calculated using a specific algorithm involving a cryptographic hash function in combination with a secret key. As with any MAC, it may be used to simultaneously verify both the data integrity and the authenticity of a message. Any iterative cryptographic hash function, such as SHA-1 or SHA-256, may be used in the calculation of an HMAC; the resulting MAC algorithm is termed HMAC-SHA1 or HMAC-SHA256 accordingly. The cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function, on the size and quality of the key and the size of the hash output length in bits.

Hardware security module (HSM): An HSM is an appliance that provides secure cryptographic key storage and operations within a tamper-resistant hardware device. HSMs are designed to securely store cryptographic key material and use the key material without exposing it outside the cryptographic boundary of the appliance. The AWS CloudHSM service provides customers with dedicated, single-tenant access to an HSM appliance.

Hypervisor: A hypervisor, also called Virtual Machine Monitor (VMM), is computer software/hardware platform virtualization software that allows multiple operating systems to run on a host computer concurrently.



Identity and Access Management (IAM): AWS IAM enables you to create multiple users and manage the permissions for each of these users within your AWS Account.

Identity pool: A store of user identity information in Amazon Cognito that is specific to your AWS Account. Identity pools use IAM roles, which are permissions that are not tied to a specific IAM user or group and that use temporary security credentials for authenticating to the AWS resources defined in the role.

Identity Provider: An online service responsible for issuing identification information for users who would like to interact with the service or with other cooperating services. Examples of identity providers include Facebook, Google, and Amazon.

Import/Export Service: An AWS service for transferring large amounts of data to Amazon S3 or EBS storage by physically shipping a portable storage device to a secure AWS facility.

Instance: An instance is a virtualized server, also known as a virtual machine (VM), with its own hardware resources and guest OS. In EC2, an instance represents one running copy of an Amazon Machine Image (AMI).

IP address: An Internet Protocol (IP) address is a numerical label that is assigned to devices participating in a computer network utilizing the Internet Protocol for communication between its nodes.

IP spoofing: Creation of IP packets with a forged source IP address, called spoofing, with the purpose of concealing the identity of the sender or impersonating another computing system.

Key: In cryptography, a key is a parameter that determines the output of a cryptographic algorithm (called a hashing algorithm). A key pair is a set of security credentials you use to prove your identity electronically and consists of a public key and a private key.

Key rotation: The process of periodically changing the cryptographic keys used for encrypting data or digitally signing requests. Just like changing passwords, rotating keys minimizes the risk of unauthorized access if an attacker somehow obtains your key or determines the value of it. AWS supports multiple concurrent access keys and certificates, which allows customers to rotate keys and certificates into and out of operation on a regular basis without any downtime to their application.

Mobile Analytics: An AWS service for collecting, visualizing, and understanding mobile application usage data. It enables you to track customer behaviors, aggregate metrics, and identify meaningful patterns in your mobile applications.

Multi-factor authentication (MFA): The use of two or more authentication factors. Authentication factors include something you know (like a password) or something you have (like a token that generates a random number). AWS IAM allows the use of a six-digit single-use code in addition to the user name and password credentials. Customers get this single-use code from an authentication device that they keep in their physical possession (either a physical token device or a virtual token from their smart phone).

Network ACLs: Stateless traffic filters that apply to all traffic inbound or outbound from a subnet within an Amazon VPC. Network ACLs can contain ordered rules to allow or deny traffic based upon IP protocol, by service port, as well as source/destination IP address.

Object: The fundamental entities stored in Amazon S3. Objects consist of object data and metadata. The data portion is opaque to Amazon S3. The metadata is a set of name-value pairs that describe the object. These include some default metadata such as the date last modified and standard HTTP metadata such as Content-Type. The developer can also specify custom metadata at the time the Object is stored.

Paravirtualization: In computing, paravirtualization is a virtualization technique that presents a software interface to virtual machines that is similar but not identical to that of the underlying hardware.

Peering: A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network.

Port scanning: A port scan is a series of messages sent by someone attempting to break into a computer to learn which computer network services, each associated with a "well-known" port number, the computer provides.

Region: A named set of AWS resources in the same geographical area. Each region contains at least two availability zones.

Replication: The continuous copying of data from a database in order to maintain a second version of the database, usually for disaster recovery purposes. Customers can use multiple AZs for their Amazon RDS database replication needs, or use Read Replicas if using MySQL.

Relational Database Service (RDS): An AWS service that allows you to create a relational database (DB) instance and flexibly scale the associated compute resources and storage capacity to meet application demand. Amazon RDS is available for MySQL, Oracle, or Microsoft SQL Server database engines.

Role: An entity in AWS IAM that has a set of permissions that can be assumed by another entity. Use roles to enable applications running on your Amazon EC2 instances to securely access your AWS resources. You grant a specific set of permissions to a role, use the role to launch an Amazon EC2 instance, and let EC2 automatically handle AWS credential management for your applications that run on Amazon EC2.

Route 53: An authoritative DNS system that provides an update mechanism that developers can use to manage their public DNS names, answering DNS queries and translating domain names into IP address so computers can communicate with each other.

Secret Access Key: A key that AWS assigns to you when you sign up for an AWS Account. To make API calls or to work with the command line interface, each AWS user needs the Secret Access Key and Access Key ID. The user signs each request with the Secret Access Key and includes the Access Key ID in the request. To help ensure the security of your AWS Account, the Secret Access Key is accessible only during key and user creation. You must save the key (for example, in a text file that you store securely) if you want to be able to access it again.

Security group: A security group gives you control over the protocols, ports, and source IP address ranges that are allowed to reach your Amazon EC2 instances; in other words, it defines the firewall rules for your instance. These rules specify which incoming network traffic should be delivered to your instance (e.g., accept web traffic on port 80).

Security Token Service (STS): The AWS STS APIs return temporary security credentials consisting of a security token, an Access Key ID, and a Secret Access Key. You can use STS to issue security credentials to users who need temporary



access to your resources. These users can be existing IAM users, non-AWS users (federated identities), systems, or applications that need to access your AWS resources.

Server-side encryption (SSE): An option for Amazon S3 storage for automatically encrypting data at rest. With Amazon S3 SSE, customers can encrypt data on upload simply by adding an additional request header when writing the object. Decryption happens automatically when data is retrieved.

Service: Software or computing ability provided across a network (e.g., Amazon EC2, Amazon S3).

Shard: In Amazon Kinesis, a shard is a uniquely identified group of data records in an Amazon Kinesis stream. A Kinesis stream is composed of multiple shards, each of which provides a fixed unit of capacity.

Signature: Refers to a digital signature, which is a mathematical way to confirm the authenticity of a digital message. AWS uses signatures calculated with a cryptographic algorithm and your private key to authenticate the requests you send to our web services.

Simple Data Base (Simple DB): A non-relational data store that allows AWS customers to store and query data items via web services requests. Amazon SimpleDB creates and manages multiple geographically distributed replicas of the customer's data automatically to enable high availability and data durability.

Simple Email Service (SES): An AWS service that provides a scalable bulk and transactional email-sending service for businesses and developers. In order to maximize deliverability and dependability for senders, Amazon SES takes proactive steps to prevent questionable content from being sent, so that ISPs view the service as a trusted email origin.

Simple Mail Transfer Protocol (SMTP): An Internet standard for transmitting email across IP networks, SMTP is used by the Amazon Simple Email Service. Customers who used Amazon SES can use an SMTP interface to send email, but must connect to an SMTP endpoint via TLS.

Simple Notification Service (SNS): An AWS service that makes it easy to set up, operate, and send notifications from the cloud. Amazon SNS provides developers with the ability to publish messages from an application and immediately deliver them to subscribers or other applications.

Simple Queue Service (SQS): A scalable message queuing service from AWS that enables asynchronous message-based communication between distributed components of an application. The components can be computers or Amazon EC2 instances or a combination of both.

Simple Storage Service (Amazon S3): An AWS service that provides secure storage for object files. Access to objects can be controlled at the file or bucket level and can further be restricted based on other conditions such as request IP source, request time, etc. Files can also be encrypted automatically using AES-256 encryption.

Simple Workflow Service (SWF): An AWS service that allows customers to build applications that coordinate work across distributed components. Using Amazon SWF, developers can structure the various processing steps in an application as "tasks" that drive work in distributed applications. Amazon SWF coordinates these tasks, managing task execution dependencies, scheduling, and concurrency based on a developer's application logic.

Single sign-on: The capability to log in once but access multiple applications and systems. A secure single sign-on capability can be provided to your federated users (AWS and non-AWS users) by creating a URL that passes the temporary security credentials to the AWS Management Console.



Snapshot: A customer-initiated backup of an EBS volume that is stored in Amazon S3, or a customer-initiated backup of an RDS database that is stored in Amazon RDS. A snapshot can be used as the starting point for a new EBS volume or Amazon RDS database or to protect the data for long-term durability and recovery.

Secure Sockets Layer (SSL): A cryptographic protocol that provides security over the Internet at the Application Layer. Both the TLS 1.0 and SSL 3.0 protocol specifications use cryptographic mechanisms to implement the security services that establish and maintain a secure TCP/IP connection. The secure connection prevents eavesdropping, tampering, or message forgery. You can connect to an AWS endpoint via HTTP or secure HTTP (HTTPS) using SSL.

Stateful firewall: In computing, a stateful firewall (any firewall that performs stateful packet inspection (SPI) or stateful inspection) is a firewall that keeps track of the state of network connections (such as TCP streams, UDP communication) traveling across it.

Storage Gateway: An AWS service that securely connects a customer's on-premises software appliance with Amazon S3 storage by using a VM that the customer deploys on a host in their data center running VMware ESXi Hypervisor. Data is asynchronously transferred from the customer's on-premises storage hardware to AWS over SSL, and then stored encrypted in Amazon S3 using AES-256.

Temporary security credentials: AWS credentials that provide temporary access to AWS services. Temporary security credentials can be used to provide identity federation between AWS services and non-AWS users in your own identity and authorization system. Temporary security credentials consist of security token, an Access Key ID, and a Secret Access Key.

Transcoder: A system that transcodes (converts) a media file (audio or video) from one format, size, or quality to another. Amazon Elastic Transcoder makes it easy for customers to transcode video files in a scalable and cost-effective fashion.

Transport Layer Security (TLS): A cryptographic protocol that provides security over the Internet at the Application Layer. Customers who used Amazon's Simple Email Service must connect to an SMTP endpoint via TLS.

Tree hash: A tree hash is generated by computing a hash for each megabyte-sized segment of the data, and then combining the hashes in tree fashion to represent ever-growing adjacent segments of the data. Glacier checks the hash against the data to help ensure that it has not been altered en route.

Vault: In Amazon Glacier, a vault is a container for storing archives. When you create a vault, you specify a name and select an AWS region where you want to create the vault. Each vault resource has a unique address.

Versioning: Every object in Amazon S3 has a key and a version ID. Objects with the same key, but different version IDs can be stored in the same bucket. Versioning is enabled at the bucket layer using PUT Bucket versioning.

Virtual Instance: Once an AMI has been launched, the resulting running system is referred to as an instance. All instances based on the same AMI start out identical and any information on them is lost when the instances are terminated or fail.

Virtual MFA: The capability for a user to get the six-digit, single-use MFA code from their smart phone rather than from a token/fob. MFA is the use of an additional factor (the single-use code) in conjunction with a user name and password for authentication.



Virtual Private Cloud (VPC): An AWS service that enables customers to provision an isolated section of the AWS cloud, including selecting their own IP address range, defining subnets, and configuring routing tables and network gateways.

Virtual Private Network (VPN): The capability to create a private, secure network between two locations over a public network such as the Internet. AWS customers can add an IPsec VPN connection between their Amazon VPC and their data center, effectively extending their data center to the cloud while also providing direct access to the Internet for public subnet instances in their Amazon VPC. In this configuration, customers add a VPN appliance on their corporate data center side.

WorkSpaces: An AWS managed desktop service that enables you to provision cloud-based desktops for your users and allows them to sign in using a set of unique credentials or their regular Active Directory credentials.

X.509: In cryptography, X.509 is a standard for a Public Key Infrastructure (PKI) for single sign-on and Privilege Management Infrastructure (PMI). X.509 specifies standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm. Some AWS products use X.509 certificates instead of a Secret Access Key for access to certain interfaces. For example, Amazon EC2 uses a Secret Access Key for access to its Query interface, but it uses a signing certificate for access to its SOAP interface and command line tool interface.

WorkDocs: An AWS managed enterprise storage and sharing service with feedback capabilities for user collaboration.

Changes since last version (Nov 2014):

- Updated compliance programs
- Updated shared security responsibility model
- Updated AWS Account security features
- Reorganized services into categories
- Updated several services with new features: CloudWatch, CloudTrail, CloudFront, EBS, ElastiCache, Redshift, Route 53, S3, Trusted Advisor, and WorkSpaces
- Added Cognito Security
- Added Mobile Analytics Security
- Added WorkDocs Security

Changes since last version (Nov 2013):

- Updated regions
- Updated several services with new features: CloudFront, DirectConnect, DynamoDB, EBS, ELB, EMR, Glacier, IAM, OpsWorks, RDS, Redshift, Route 53, Storage Gateway, and VPC
- Added AppStream Security
- Added CloudTrail Security
- Added Kinesis Security
- Added WorkSpaces Security

Changes since last version (May/June 2013):

- Updated IAM to incorporate roles and API access
- Updated MFA for API access for customer-specified privileged actions
- Updated RDS to add event notification, multi-AZ, and SSL to SQL Server 2012
- Updated VPC to add multiple IP addresses, static routing VPN, and VPC By Default
- Updated several other services with new features: CloudFront, CloudWatch, EBS, ElastiCache, Elastic Beanstalk, Route 53, S3, Storage Gateway
- Added Glacier Security
- Added Redshift Security
- Added Data Pipeline Security
- Added Transcoder Security
- Added Trusted Advisor Security
- Added OpsWorks Security
- Added CloudHSM Security

Changes since last version (May 2011):

- Reorganization to better identify infrastructure versus service-specific security
- Changed Control Environment Summary heading to AWS Compliance Program
- Changed Information and Communication heading to Management and Communication
- Changed Employee Lifecycle heading to Logical Access
- Changed Configuration Management heading to Change Management
- Merged Environmental Safeguards section with Physical Security section
- Incorporated information in Backups section into S3, SimpleDB, and EBS sections



- Update to certifications to reflect SAS70 name change to SSAE 16 and addition of FedRAMP
- Update to Network Security section to add Secure Network Architecture and Network Monitoring and Protection
- Update to IAM to incorporate roles/key provisioning, virtual MFA, temporary security credentials, and single sign on
- Update to regions to include new regions and GovCloud description
- Updated EBS, S3, SimpleDB, RDS, and EMR to clarify service and security descriptions
- Update to VPC to add configuration options, VPN, and Elastic Network Interfaces
- Addition of Amazon Direct Connect Security section
- Addition of Amazon Elastic Load Balancing Security
- Addition of AWS Storage Gateway Security
- Addition of AWS Import/Export Security
- Addition of Auto Scaling Security
- Addition of Amazon DynamoDB Security
- Addition of Amazon ElastiCache Security
- Addition of Amazon Simple Workflow Service (Amazon SWS) Security
- Addition of Amazon Simple Email Service (Amazon SES) Security
- Addition of Amazon Route 53 Security
- Addition of Amazon CloudSearch Security
- Addition of AWS Elastic Beanstalk Security
- Addition of AWS CloudFormation Security
- Updated glossary

Changes since last version (Aug 2010):

- Addition of AWS Identity and Access Management (AWS IAM)
- Addition of Amazon Simple Notification Service (SNS) Security
- Addition of Amazon CloudWatch Security
- Addition of Auto Scaling Security
- Update to Amazon Virtual Private Cloud (Amazon VPC)
- Update to Control Environment
- Removal of Risk Management because it has been expanded in a separate whitepaper

Changes since last version (Nov 2009):

- Major revision

Changes since last version (June 2009):

- Change to Certifications and Accreditations section to reflect SAS70
- Addition of Amazon Virtual Private Cloud (Amazon VPC)
- Addition of Security Credentials section to highlight AWS Multi-Factor Authentication and Key Rotation
- Addition of Amazon Relational Database Service (Amazon RDS) Security

Changes since last version (Sep 2008):

- Addition of Security Design Principles
- Update of Physical Security information and inclusion of background checks
- Backup section updated for clarity with respect to Amazon EBS



- Update of Amazon EC2 Security section to include:
- Certificate-based SSHv2
- Multi-tier security group detail and diagram
- Hypervisor description and Instance Isolation diagram
- Fault Separation
- Addition of Configuration Management
- Amazon S3 section updated for detail and clarity
- Addition of Storage Device Decommissioning
- Addition of Amazon SQS Security
- Addition of Amazon CloudFront Security
- Addition of Amazon Elastic MapReduce Security

Notices

© 2010-2015 Amazon.com, Inc., or its affiliates. This document is provided for informational purposes only. It represents AWS's current product offerings as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.



Data Processing Amendment to Google Apps Agreement

The Customer agreeing to these terms ("**Customer**") and Google Inc., Google Ireland Limited, Google Commerce Limited or Google Asia Pacific Pte. Ltd. (as applicable, "**Google**") have entered into a Google Apps for Work Agreement, Google Apps Enterprise Agreement, Google Apps for Business Agreement, Google Apps for Work via Reseller Agreement, Google Apps Enterprise via Reseller Agreement, Google Apps for Business via Reseller Agreement, Google Apps for Education Agreement or Google Apps for Education via Reseller Agreement, as applicable (as amended to date, the "**Google Apps Agreement**"). This amendment (the "**Data Processing Amendment**") is entered into by Customer and Google as of the Amendment Effective Date and amends the Google Apps Agreement.

The "**Amendment Effective Date**" is: (a) if this Data Processing Amendment is incorporated into the Google Apps Agreement by reference, the effective date of the Google Apps Agreement, as defined in that agreement; or (b) if this Data Processing Amendment is not incorporated into the Google Apps Agreement by reference, the date Customer accepts this Data Processing Amendment by clicking to accept these terms.

If this Data Processing Amendment is not incorporated into the Google Apps Agreement by reference and you are accepting on behalf of Customer, you represent and warrant that: (i) you have full legal authority to bind your employer, or the applicable entity, to these terms; (ii) you have read and understand these terms; and (iii) you agree, on behalf of the party you represent, to this Data Processing Amendment. If you do not have the legal authority to bind Customer, please do not click the "I Accept" button.

1. **Introduction.**

This Data Processing Amendment reflects the parties' agreement with respect to terms governing the processing of Customer Data under the Google Apps Agreement.

2. **Definitions.**

2.1. Capitalized terms used but not defined in this Data Processing Amendment have the meanings given in the Google Apps Agreement. In this Data Processing Amendment, unless expressly stated otherwise:

"**Additional Products**" means products, services and applications that are not part of the Services but that may be accessible, via the Admin Console or otherwise, for use with the Services.

"**Advertising**" means online advertisements displayed by Google to End Users, excluding any advertisements Customer expressly chooses to have Google or any Google Affiliate display in connection with the Services under a separate agreement (for example, Google AdSense advertisements implemented by Customer on a website created by Customer using the "Google Sites" functionality within the Services).

"**Affiliate**" means any entity controlling, controlled by, or under common control with a party, where "control" is defined as (a) the ownership of at least fifty percent (50%) of the equity or beneficial interests of the entity; (b) the right to vote for or appoint a majority of the board of directors or other governing body of the entity; or (c) the power to exercise a controlling influence over the management or policies of the entity.

“Agreement” means the Google Apps Agreement, as amended by this Data Processing Amendment and as may be further amended from time to time in accordance with the Google Apps Agreement.

“Customer Data” means data (which may include personal data and the categories of data referred to in Appendix 1) submitted, stored, sent or received via the Services by Customer, its Affiliates or End Users.

“Data Incident” means (a) any unlawful access to Customer Data stored in the Services or systems, equipment or facilities of Google or its Sub processors, or (b) unauthorized access to such Services, systems, equipment or facilities that results in loss, disclosure or alteration of Customer Data.

“Data Privacy Officer” means Google’s Data Privacy Officer for Apps.

“Data Protection legislation” means, as applicable: (a) any national provisions adopted pursuant to the Directive that are applicable to Customer and or any Customer Affiliates as the controller(s) of the Customer Data; and or (b) the federal Data Protection Act of June 1, 2002 (Switzerland).

“Directive” means Directive 95/46/EC of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the free Movement of Such Data.

“EEA” means the European Economic Area.

“Google Group” means those Google Affiliates involved in provision of the Services to Customer.

“Instructions” means Customer’s written instructions to Google consisting of the Agreement, including instructions to Google to provide the Services and technical support for the Services as set out in the Agreement; instructions given by Customer, its Affiliates and End Users via the Admin Console and otherwise in its and their use of the Services and related technical support services; and any subsequent written instructions given by Customer to Google and acknowledged by Google.

“Model Contract Clauses” or **“MCCs”** means the standard contractual clauses (processors) for the purposes of Article 2 (2) of Directive 95/46/EC for the transfer of personal data to processors established in third countries which do not ensure an adequate level of data protection.

“Safe Harbor or Certification” means a current certification to the U.S. Department of Commerce Safe Harbor framework requirements as set out at the following URL: http://export.gov/safeharbor/eu/eg_main_01_5.asp, or any replacement framework or URL from time to time.

“Services” means, for purposes of this Data Processing Amendment, the Google Apps for Work Services which are described at www.google.com/apps/intl/en/terms/user/features.html (as such services and URL link may be updated or modified by Google from time to time in accordance with the Google Apps Agreement).

“**u processors**” means (a) all Google Group entities that have logical access to and process Customer Data (each, a “**Google Group u processor**”); and (b) all third parties (other than Google Group entities) that are engaged to provide services to Customer and that have logical access to and process Customer Data (each, a “**ird Part u processor**”).

“**erm**” means the term of the Google Apps Agreement, as defined in that agreement.

“**ird Part Auditor**” means a qualified and independent third party auditor, whose then current identity Google will disclose to Customer.

2.2. The terms “personal data”, “processing”, “data subject”, “controller” and “processor” have the meanings given to them in the Directive. The terms “data importer” and “data exporter” have the meanings given to them in the Model Contract Clauses.

erm.

This Data Processing Amendment will take effect on the Amendment Effective Date and, notwithstanding expiry or termination of the Google Apps Agreement, will remain in effect until, and automatically terminate upon, deletion by Google of all data as described in Section (Data Deletion) of this Data Processing Amendment.

Data Protection egislation.

The parties agree and acknowledge that the Data Protection Legislation may apply to the processing of Customer Data.

5. Processing of Customer Data.

5.1. **Controller and Processor.** If the Data Protection Legislation applies to the processing of Customer Data, then as between the parties, the parties acknowledge and agree that: (a) Customer is the controller of Customer Data under the Agreement; (b) Google is a processor of such data; (c) Customer will comply with its obligations as a controller under the Data Protection Legislation; and (d) Google will comply with its obligations as a processor under the Agreement. If under the Data Protection Legislation a Customer Affiliate is considered the controller (either alone or jointly with the Customer) with respect to certain Customer Data, Customer represents and warrants to Google that Customer is authorized (i) to give the Instructions to Google and otherwise act on behalf of such Customer Affiliate in relation to such Customer Data as described in this Data Processing Amendment, and (ii) to bind the Customer Affiliate to the terms of this Data Processing Amendment.

5.2. **cope of Processing.** Google will only process Customer Data in accordance with the Instructions, and will not process Customer Data for any other purpose.

5. **Processing estrictions.** otwithstanding any other term of the Agreement, Google will not process Customer Data for Advertising purposes or serve Advertising in the Services.

5. **Additional Products.** Customer acknowledges that if it installs, uses, or enables Additional Products, the Services may allow such Additional Products to access Customer Data as required for the interoperation of those Additional Products with the Services. This Data Processing Amendment does not apply to the processing of data transmitted to or from such Additional Products. Customer can enable or disable Additional Products. Customer is not required to use Additional Products in order to use the Services.

Data ecurit ecurit Compliance Audits.

.1. **Security Measures**. Google will take and implement appropriate technical and organizational measures to protect Customer Data against accidental or unlawful destruction or accidental loss or alteration or unauthorized disclosure or access or other unauthorized processing, as detailed in Appendix 2 (“**Security Measures**”). Google may update or modify the Security Measures from time to time provided that such updates and modifications do not result in the degradation of the overall security of the Services. Customer agrees that it is solely responsible for its use of the Services, including securing its account authentication credentials, and that Google has no obligation to protect Customer Data that Customer elects to store or transfer outside of Google’s and its Subprocessors’ systems (e.g., offline or on premise storage).

.2. **Security Compliance** **Google staff**. Google will take appropriate steps to ensure compliance with the Security Measures by its employees, contractors and Subprocessors to the extent applicable to their scope of performance.

. **Data Incidents**. If Google becomes aware of a Data Incident, Google will promptly notify Customer of the Data Incident, and take reasonable steps to minimize harm and secure Customer Data. Notification(s) of any Data Incident(s) will be delivered to the Notification Email Address provided by Customer in connection with the Agreement or, at Google’s discretion, by direct communication (e.g., by phone call or an in person meeting). Customer acknowledges that it is solely responsible for ensuring the contact information given for purposes of the Notification Email Address is current and valid, and for fulfilling any third party notification obligations. Customer agrees that “Data Incidents” do not include: (i) unsuccessful access attempts or similar events that do not compromise the security or privacy of Customer Data, including pings, port scans, denial of service attacks and other network attacks on firewalls or networked systems; or (ii) accidental loss or disclosure of Customer Data caused by Customer’s use of the Services or Customer’s loss of account authentication credentials. Google’s obligation to report or respond to a Data Incident under this Section will not be construed as an acknowledgement by Google of any fault or liability with respect to the Data Incident.

. **Compliance with Security and Privacy Standards** **Cloud Reports**. During the Term, Google will maintain the following:

(a) its ISO/IEC 27001:2017 Certification or a comparable certification (“**ISO Certification**”) for the Services;

(b) conformity of the Services with ISO/IEC 27017:2015 or a comparable standard (“**ISO Conformity**”), as independently verified;

(c) its confidential Service Organization Control (SOC) 2 Report (or a comparable report) on Google’s systems examining logical security controls, physical security controls, and system availability as related to the Services (the “**SOC Report**”), as produced by the Third Party Auditor and updated at least once every eighteen (18) months; and

(d) its Service Organization Control (SOC) Report (or a comparable report) as related to the Services (the “**SOC Report**”), as produced by the Third Party Auditor and updated at least once every eighteen (18) months.

.5. **Auditing Security Compliance**

.5.1. Review of Security Documentation. Google will make the following available for review by Customer:

- (a) the certificate issued in relation to Google's ISO 27001 Certification;
- (b) the then current SOC 2 Report;
- (c) a summary or redacted version of the then current confidential SOC 2 Report; and
- (d) following a request by Customer in accordance with Section 5.2 below, the then current confidential SOC 2 Report.

.5.2. Customer Audits. If Customer (or an authorized Customer Affiliate) has entered into Model Contract Clauses as described in Section 10.2 of this Data Processing Amendment, Customer or such Customer Affiliate may exercise the audit rights granted under clauses 5(f) and 12(2) of such Model Contract Clauses:

- (a) by instructing Google to execute the audit as described in Sections 5.1 and 5.1 above; and or
- (b) following a request by Customer in accordance with Section 5.2 below, by executing an audit as described in such Model Contract Clauses.

.5.3. Additional business terms for Review and Audits. Google and Customer (or an authorized Customer Affiliate if applicable) will discuss and agree in advance on:

- (a) the reasonable date(s) of and security and confidentiality controls applicable to any Customer review under Section 5.1(d); and
- (b) the identity of a suitably qualified and independent third party auditor for any audit under Section 5.2(b), and the reasonable start date, scope and duration of and security and confidentiality controls applicable to any such audit.

Google reserves the right to charge a fee (based on Google's reasonable costs) for any review under Section 5.1(d) and or audit under Section 5.2(b). For clarity, Google is not responsible for any costs incurred or fees charged by any third party auditor appointed by Customer (or an authorized Customer Affiliate) in connection with an audit under Section 5.2(b). Nothing in this Section 5.3 varies or modifies any rights or obligations of Customer (or any authorized Customer Affiliate) or Google Inc. under any Model Contract Clauses entered into as described in Section 10.2 (Transfers of Data out of the EEA) of this Data Processing Amendment.

.5.4. Requests for Review and Audits. Any requests under Section 5.1 or 5.2 must be sent to the Data Privacy Officer as described in Section 1 (Data Privacy Officer) of this Data Processing Amendment.

6. Data Deletion.

.6.1. Deletion Customer and End Users. During the Term, Google will provide Customer or End Users with the ability to delete Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. Once Customer or End User deletes Customer Data and such Customer Data cannot be recovered by the Customer or End User, such

as from the “trash” (“**Customer-Deleted Data**”), Google will delete such data from its systems as soon as soon as reasonably practicable within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

7.2. **Deletion on Standard Termination**. On expiry or termination of the Google Apps Agreement (or, if applicable, on expiry of any post-termination period during which Google may agree to continue providing the Services), Google will, subject to Section 7.3 (Deletion on Termination for Non-Payment or No Purchase) below, delete all Customer-Deleted Data from its systems as soon as reasonably practicable within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

7.3. **Deletion on Termination for Non-Payment or No Purchase**. On termination of the Google Apps Agreement due to Customer breaching its payment obligations or opting not to purchase the Services at the end of a free trial of the Services, Google will delete all Customer Data from its systems within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

8. **Access to Data.**

8.1. **Access; Export of Data**. During the Term, Google will provide Customer with access to and the ability to correct, block and export Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. To the extent Customer, in its use and administration of the Services during the Term, does not have the ability to correct or block Customer Data as required by applicable law, or to migrate Customer Data to another system or service provider, Google will comply with any reasonable requests by Customer to assist in facilitating such actions to the extent Google is legally permitted to do so and has reasonable access to the Customer Data.

8.2. **End User Requests**. During the Term, if Google receives any request from an End User for records relating to that End User’s personal data included in the Customer Data, Google will advise such End User to submit its request to Customer. Customer will be responsible for responding to any such request using the functionality of the Services.

9. **Data Privacy Officer.**

The Data Privacy Officer can be contacted by Customer Administrators at: https://support.google.com/a/contact/gfw_dpo (or via such other means as may be provided by Google). Administrators must be signed in to their Admin Account to use this address.

10. **Data Transfers.**

10.1. **Data Storage and Processing Facilities**. Google may store and process Customer Data in the United States or any other country in which Google or any of its Subprocessors maintains facilities, subject to Section 10.2 (Transfers of Data Out of the EEA) below.

10.2. **Transfers of Data Out of the EEA**. If the storage and processing of Customer Data (as set out in Section 10.1 above) involves transfers of Customer personal data out of the EEA and Data Protection Legislation applies to those transfers, Google will:

10.2.1 ensure that Google Inc. maintains its Safe Harbor Certification, and that the transfers are made in accordance with such Safe Harbor Certification; and/or

10.2.2 ensure that Google Inc. as the data importer of such Customer personal data enters into Model Contract Clauses with Customer (or an authorized Customer Affiliate) as the data exporter of such data, if Customer so requests, and that the transfers are made in accordance with any such Model Contract Clauses; and/or

10.2.3 adopt an alternative solution that achieves compliance with the terms of the Directive for transfers of personal data to a third country, and ensure that the transfers are made in accordance with any such compliance solution.

10.3. **Safe Harbor or Certification and Processing Practices.** While Google Inc. maintains its Safe Harbor Certification pursuant to Section 10.2.1, Google will ensure that: (a) the scope of such Safe Harbor Certification includes Customer Data; and (b) the Google Group's processing practices in respect of Customer Data remain consistent with those described in such Safe Harbor Certification.

10.4. **Data Center Information.** Google will make available to Customer information about the countries in which data centers used to store Customer Data are located.

11. **Subprocessors.**

11.1. **Subprocessors.** Google may engage Subprocessors to provide parts of the Services and related technical support services, subject to the restrictions in this Data Processing Amendment.

11.2. **Subprocessing Restrictions.** Google will ensure that Subprocessors only access and use Customer Data in accordance with the terms of the Agreement and that Subprocessors are bound by written agreements that require them to provide at least the level of data protection required by the following, as applicable pursuant to Section 10.2 (Transfers of Data Out of the EEA): (a) any Safe Harbor Certification maintained by Google Inc.; (b) any Model Contract Clauses entered into by Google Inc. and Customer (or an authorized Customer Affiliate); and/or (c) any alternative compliance solution adopted by Google.

11.3. **Consent to Subprocessing.** Customer consents to Google subcontracting the processing of Customer Data to Subprocessors in accordance with the Agreement. If the Model Contract Clauses have been entered into as described above, Customer (or, if applicable, an authorized Customer Affiliate) consents to Google Inc. subcontracting the processing of Customer Data in accordance with the terms of the Model Contract Clauses.

11.4. **Additional Information.** Information about Third Party Subprocessors is available at the following URL: www.google.com/intl/en/work/apps/terms/subprocessors.html, as such URL may be updated by Google from time to time. The information available at the URL is accurate at the time of publication. At the written request of the Customer, Google will provide additional information regarding Subprocessors and their locations. Any such requests must be sent to the Data Privacy Officer for Google Apps as described in Section 9 (Data Privacy Officer) of this Data Processing Amendment.

11.5. **Termination.** Google will, at least 14 days before appointing any new Third Party Subprocessor, inform Customer of the appointment (including the name and location of such subprocessor and the activities it will perform) either by sending an email to the Notification Email Address or via the Admin Console. If Customer objects to Google's use of any new Third Party Subprocessor, Customer may, as its sole and exclusive remedy, terminate the Google Apps Agreement by giving written notice to Google within 30 days of being informed by Google of the appointment of such subprocessor.

12. Liability Cap.

If Google Inc. and Customer (or an authorized Customer Affiliate) enter into Model Contract Clauses as described above, then, subject to the remaining terms of the Agreement relating to liability (including any specific exclusions from any limitation of liability), the total combined liability of Google and its Affiliates, on the one hand, and Customer and its Affiliates, on the other hand, under or in connection with the Agreement and all those MCCs combined will be limited to the maximum monetary or payment-based liability amount set out in the Agreement.

13. Third Party Beneficiary.

Notwithstanding anything to the contrary in the Agreement, where Google Inc. is not a party to the Agreement, Google Inc. will be a third party beneficiary of Section 10 (Auditing Security Compliance), Section 11.3 (Consent to Subprocessing) and Section 12 (Liability Cap) of this Data Processing Amendment.

14. Effect of Amendment.

To the extent of any conflict or inconsistency between the terms of this Data Processing Amendment and the remainder of the Agreement, the terms of this Data Processing Amendment will govern. Subject to the amendments in this Data Processing Amendment, the Agreement remains in full force and effect.

Appendix Categories of Data and Data Subjects**Categories of Data**

Personal data submitted, stored, sent or received by Customer or End Users via the Services may include the following categories of data: user IDs, email, documents, presentations, images, calendar entries, tasks and other electronic data

Data Subjects

Personal data submitted, stored, sent or received via the Services may concern the following categories of data subjects: End Users including Customer's employees and contractors; the personnel of Customer's customers, suppliers and subcontractors; and any other person who transmits data via the Services, including individuals collaborating and communicating with End Users.

Appendix Security Measures

As of the Amendment Effective Date, Google will take and implement the Security Measures set out in this Appendix to the Data Processing Amendment. Google may update or modify such Security Measures from time to time provided that such updates and modifications do not result in the degradation of the overall security of the Services.

1. Data Center Network Security.**(a) Data Centers.**

Infrastructure. Google maintains geographically distributed data centers. Google stores all production data in physically secure data centers.

Redundancy. Infrastructure systems have been designed to eliminate single points of failure and minimize the impact of anticipated environmental risks. Dual circuits, switches, networks or other necessary devices help provide this redundancy. The Services are designed to allow Google to perform certain types of preventative and corrective maintenance without interruption. All environmental equipment and facilities have documented preventative maintenance procedures that detail the process for and frequency of performance in accordance with the

manufacturer's or internal specifications. Preventative and corrective maintenance of the data center equipment is scheduled through a standard change process according to documented procedures.

Power. The data center electrical power systems are designed to be redundant and maintainable without impact to continuous operations, 24 hours a day, and 7 days a week. In most cases, a primary as well as an alternate power source, each with equal capacity, is provided for critical infrastructure components in the data center. Backup power is provided by various mechanisms such as uninterruptible power supplies (UPS) batteries, which supply consistently reliable power protection during utility brownouts, blackouts, over voltage, under voltage, and out-of-tolerance frequency conditions. If utility power is interrupted, backup power is designed to provide transitory power to the data center, at full capacity, for up to 10 minutes until the diesel generator systems take over. The diesel generators are capable of automatically starting up within seconds to provide enough emergency electrical power to run the data center at full capacity typically for a period of days.

Server Operating Systems. Google servers use a Linux based implementation customized for the application environment. Data is stored using proprietary algorithms to augment data security and redundancy. Google employs a code review process to increase the security of the code used to provide the Services and enhance the security products in production environments.

Businesses Continuity. Google replicates data over multiple systems to help to protect against accidental destruction or loss. Google has designed and regularly plans and tests its business continuity planning/disaster recovery programs.

(b) **Networks Transmission.**

Data Transmission. Data centers are typically connected via high-speed private links to provide secure and fast data transfer between data centers. This is designed to prevent data from being read, copied, altered or removed without authorization during electronic transfer or transport or while being recorded onto data storage media. Google transfers data via Internet standard protocols.

External Attack Surface. Google employs multiple layers of network devices and intrusion detection to protect its external attack surface. Google considers potential attack vectors and incorporates appropriate purpose built technologies into external facing systems.

Intrusion Detection. Intrusion detection is intended to provide insight into ongoing attack activities and provide adequate information to respond to incidents. Google's intrusion detection involves:

1. Tightly controlling the size and make-up of Google's attack surface through preventative measures;
2. Employing intelligent detection controls at data entry points; and
3. Employing technologies that automatically remedy certain dangerous situations.

Incident Response. Google monitors a variety of communication channels for security incidents, and Google's security personnel will react promptly to known incidents.

Encryption Technologies. Google makes HTTPS encryption (also referred to as SSL or TLS connection) available.

2. Access and Site Controls.

(a) Site Controls.

On-site Data Center Security Operation. Google's data centers maintain an on-site security operation responsible for all physical data center security functions 24 hours a day, 7 days a week. The on-site security operation personnel monitor Closed Circuit Television (CCTV) cameras and all alarm systems. On-site Security operation personnel perform internal and external patrols of the data center regularly.

Data Center Access Procedures. Google maintains formal access procedures for allowing physical access to the data centers. The data centers are housed in facilities that require electronic card key access, with alarms that are linked to the on-site security operation. All entrants to the data center are required to identify themselves as well as show proof of identity to on-site security operations. Only authorized employees, contractors and visitors are allowed entry to the data centers. Only authorized employees and contractors are permitted to request electronic card key access to these facilities. Data center electronic card key access requests must be made through e-mail, and require the approval of the requestor's manager and the data center director. All other entrants requiring temporary data center access must: (i) obtain approval in advance from the data center managers for the specific data center and internal areas they wish to visit; (ii) sign in at on-site security operations; and (iii) reference an approved data center access record identifying the individual as approved.

On-site Data Center Security Devices. Google's data centers employ an electronic card key and biometric access control system that is linked to a system alarm. The access control system monitors and records each individual's electronic card key and when they access perimeter doors, shipping and receiving, and other critical areas. Unauthorized activity and failed access attempts are logged by the access control system and investigated, as appropriate. Authorized access throughout the business operations and data centers is restricted based on roles and the individual's job responsibilities. The fire doors at the data centers are alarmed. CCTV cameras are in operation both inside and outside the data centers. The positioning of the cameras has been designed to cover strategic areas including, among others, the perimeter, doors to the data center building, and shipping/receiving. On-site security operations personnel manage the CCTV monitoring, recording and control equipment. Secure cables throughout the data centers connect the CCTV equipment. Cameras record on site via digital video recorders 24 hours a day, 7 days a week. The surveillance records are retained for up to 30 days based on activity.

(b) Access Control.

Infrastructure Security Personnel. Google has, and maintains, a security policy for its personnel, and requires security training as part of the training package for its personnel. Google's infrastructure security personnel are responsible for the ongoing monitoring of Google's security infrastructure, the review of the Services, and responding to security incidents.

Access Control and Privilege Management. Customer's Administrators and End Users must authenticate themselves via a central authentication system or via a single sign on system in order to use the Services. Each application checks credentials in order to allow the display of data to an authorized End User or authorized Administrator.

Internal Data Access Processes and Policies – Access Policy. Google's internal data access processes and policies are designed to prevent unauthorized persons and/or systems from gaining access to systems used to process personal data. Google aims to design its systems to: (i) only allow authorized persons to access data they are authorized to access; and (ii) ensure that personal data cannot be read, copied, altered or removed without authorization during processing, use and after recording. The systems are designed to detect any inappropriate access. Google employs a centralized access management system to control personnel access to production servers, and only provides access to a limited number of authorized personnel. LDAP, Kerberos and a proprietary system utilizing RSA keys are designed to provide Google with secure and flexible access mechanisms. These mechanisms are designed to grant only approved access rights to site hosts, logs, data and configuration information. Google requires the use of unique user IDs, strong passwords, two factor authentication and carefully monitored access lists to minimize the potential for unauthorized account use. The granting or modification of access rights is based on: the authorized personnel's job responsibilities; job duty requirements necessary to perform authorized tasks; and a need to know basis. The granting or modification of access rights must also be in accordance with Google's internal data access policies and training. Approvals are managed by workflow tools that maintain audit records of all changes. Access to systems is logged to create an audit trail for accountability. Where passwords are employed for authentication (e.g., login to workstations), password policies that follow at least industry standard practices are implemented. These standards include password expiry, restrictions on password reuse and sufficient password strength. For access to extremely sensitive information (e.g., credit card data), Google uses hardware tokens.

3. **Data.**

(a) **Data Storage, Isolation & Authentication.**

Google stores data in a multi-tenant environment on Google-owned servers. Data, the Services database and file system architecture are replicated between multiple geographically dispersed data centers. Google logically isolates data on a per End User basis at the application layer. Google logically isolates each Customer's data, and logically separates each End User's data from the data of other End Users, and data for an authenticated End User will not be displayed to another End User (unless the former End User or an Administrator allows the data to be shared). A central authentication system is used across all Services to increase uniform security of data.

The Customer will be given control over specific data sharing policies. Those policies, in accordance with the functionality of the Services, will enable Customer to determine the product sharing settings applicable to End Users for specific purposes. Customer may choose to make use of certain logging capability that Google may make available via the Services, products and APIs. Customer agrees that its use of the APIs is subject to the API Terms of Use. Google agrees that changes to the APIs will not result in the degradation of the overall security of the Services.

(b) **Decommissioned Disks and Disk Erase Policy.**

Certain disks containing data may experience performance issues, errors or hardware failure that lead them to be decommissioned ("Decommissioned Disk"). Every Decommissioned Disk is subject to a series of data destruction processes (the "Disk Erase Policy") before leaving Google's premises either for reuse or destruction. Decommissioned Disks are erased in a multi-step process and verified complete by at least two independent validators. The erase results are logged by the

Decommissioned Disk's serial number for tracking. Finally, the erased Decommissioned Disk is released to inventory for reuse and redeployment. If, due to hardware failure, the Decommissioned Disk cannot be erased, it is securely stored until it can be destroyed. Each facility is audited regularly to monitor compliance with the Disk Erase Policy

Personnel Security.

Google personnel are required to conduct themselves in a manner consistent with the company's guidelines regarding confidentiality, business ethics, appropriate usage, and professional standards. Google conducts reasonably appropriate backgrounds checks to the extent legally permissible and in accordance with applicable local labor law and statutory regulations.

Personnel are required to execute a confidentiality agreement and must acknowledge receipt of, and compliance with, Google's confidentiality and privacy policies. Personnel are provided with security training. Personnel handling Customer Data are required to complete additional requirements appropriate to their role (eg., certifications). Google's personnel will not process Customer Data without authorization.

Subprocessor Security.

Prior to onboarding Subprocessors, Google conducts an audit of the security and privacy practices of Subprocessors to ensure Subprocessors provide a level of security and privacy appropriate to their access to data and the scope of the services they are engaged to provide. Once Google has assessed the risks presented by the Subprocessor, then subject always to the requirements set out in Section 4.2 (Subprocessing Restrictions) of this Data Processing Amendment, the Subprocessor is required to enter into appropriate security, confidentiality and privacy contract terms.

Google Apps Data Processing Amendment, Version 2.0

Getting Started with Microsoft Azure Virtual Machines

Introduction

You can use a Microsoft Azure Virtual Machine when you need a scalable, cloud-based server running a Windows or Linux operating system and any application of your choosing. By taking advantage of Microsoft Azure Infrastructure as a Service (IaaS), you can run a virtual machine on an ongoing basis, or you can stop and restart it later with no loss to your data or server settings.

You can quickly provision a new virtual machine from one of the images available from Microsoft Azure. On the other hand, if you are already running applications in a VMware or Hyper-V virtualized environment, you can easily migrate your virtual machine to Microsoft Azure. Once you have created your virtual machine and added it to Microsoft Azure, you can work with it much like an on-premises server by attaching more disks for data storage or by installing and running applications on it.

Virtual machines rely on Microsoft Azure Storage for high availability. When your virtual machine is provisioned, it is replicated to three separate locations within the data center to which you assign it. You have the option to enable geo-replication to have copies of your virtual machine available in a remote data center region.

Considering Scenarios for a Virtual Machine

Microsoft Azure gives you not only the flexibility to support many application platforms, but also the flexibility to scale up and scale down to suit your requirements. Furthermore, you can quickly provision a new virtual machine in a few minutes. A Microsoft Azure virtual machine is simply a fresh machine preloaded with an operating system of your choice—you can add any needed application easily. Consider the following suggested scenarios as merely a few of the possibilities.

Existing virtual environment. You can move an existing server workload into the cloud simply by copying an existing on-premises VHD to Microsoft Azure. That way, you can leverage cloud benefits such as elasticity, redundancy, and automatic platform updates, while paying only for storage and the time the virtual machine is in use.

Cloud workstation. You can create a virtual machine for use in the cloud that contains applications you commonly use and want to access from any location. Rather than set up a desktop workstation or carry a laptop to run the applications you use frequently, you can install your applications on a virtual machine that is accessible in the cloud. With this type of environment, you can travel easily from location to location with a tablet and then connect to a virtual machine using RDP or SSH to access your routinely used applications such as Matlab, R, IPython, or SQL Server.

Linux cluster. You can create a virtual machine based on a Linux operating system and then change the configuration to run multiple virtual machines as a cluster. By using a cluster, you provide a fast, high capacity, and redundant environment for any custom Big Data and Big Compute applications you choose to install on the virtual machine.

Development or test environment. Another option is to set up virtual machines for development or test purposes that you use solely on an as-needed basis. For example, you could install Visual Studio on a

virtual machine for use during the development phase of a project without the need to allocate on-premises resources to development. Likewise, rather than incur the cost of setting up and maintaining a load test environment that you use only periodically, you can set up virtual machines to run only when you need an independent and isolated environment.

Custom machine images. Perhaps you need a team of people to work with a common set of complex tools and packages. You can take advantage of community-contributed virtual machine images available at VM Depot. Another option is to create and contribute a virtual machine that you preconfigure with selected applications or packages for an online community to reuse in the cloud. See “Using and Contributing Virtual Machines to VM Depot” at <http://research.microsoft.com/en-us/projects/azure/technical-papers.aspx> to learn more.



Creating a Virtual Machine

There are two primary ways to create a virtual machine that runs in Microsoft Azure—mounting an existing disk or using an image.



- A disk is a virtual hard disk (VHD) that already contains an operating system at minimum and optionally includes one or more applications. Any configuration changes you make are persisted on the disk and available the next time you start the VHD.



- An image is a template that contains a base operating system (Windows or Linux) and perhaps one or more applications, but no user accounts. A disk for the operating system is added to Microsoft Azure Storage and assigned to the new virtual machine during the provisioning process. You can use an image to quickly reproduce the same virtual machine multiple times.

Using the Virtual Machine Image Gallery

A simple way to get started using a virtual machine is to create virtual machine from the Image Gallery that you can access from the Microsoft Azure Management Portal. To view the available images, click the New button (which displays as a plus symbol on the bottom left side of the portal), select Virtual Machine in the Compute group, and then click From Gallery, as shown in Figure 1.




Figure 1: Creating a New Virtual Machine

The Create a Virtual Machine dialog box gives you the following three options:

- **Platform Images.** Here you can find a collection of images, as shown in Figure 2, containing only an operating system, such as Windows Server 2012 R2 or OpenLogic CentOS 6.3, or images that contain both an operating system and an application, such as SharePoint Server 2013 or Visual Studio Ultimate 2013. For a list of supported operating system roles and features as well as supported software, see <http://support.microsoft.com/kb/2721672>.
- **My Images.** This section of the Image Gallery displays a list of images that you create and store in Microsoft Azure. To create an image, you can capture an existing Microsoft Azure Virtual Machine, use SUSE Studio to build and migrate an image, or upload your own virtual machine, as explained later in this document.
- **My Disks.** You can create a virtual machine simply by selecting a disk on this list.

Note: In addition to the images available in the Virtual Machine Image Gallery, there are images created and contributed by members of the online community available at [VM Depot](http://vmdepot.net). Refer to the document "Using and Contributing Virtual Machines to VM Depot" at <http://research.microsoft.com/en-us/projects/azure/technical-papers.aspx> for more information.




Figure 2: Selecting a Virtual Machine Image from the Gallery

After you select the image you want, you assign a name to the virtual machine, specify a size, and create the administrator's login name and password for a Windows operating system or specify the authentication method for a Linux operating system. Some platform images have minimum size recommendations noted in the image description. The number of CPU cores and the amount of available memory defines the size of a virtual machine. At one end of the spectrum, a virtual machine can share a CPU core, consume only 768 MB of memory, and support one data disk. The maximum size of a data disk is 1 TB. At the other end of the spectrum, a virtual machine can be allocated eight CPU cores, up to 56 GB of memory, and up to 16 data disks at the time of this writing, but these maximum values are likely to increase as Microsoft Azure adds more powerful hardware offerings in the future. Pricing of the virtual machine depends on the size of the virtual machine you select.

Your next step is to configure the following additional settings for your virtual machine:

- **Cloud Service.** A cloud service is simply a container for your virtual machines. You can host a single virtual machine in a cloud service or assign multiple virtual machines to the same cloud service. When configuring the cloud service, you can request a new cloud service or select one that you created earlier.
- **Cloud Service DNS Name.** You must provide a unique name that you later use to connect to the virtual machine.
- **Subscription.** You must assign the virtual machine to an existing subscription.
- **Region/Affinity Group/Virtual Network.** You must select a data center region to host your virtual machine. Alternatively, if you have created an affinity group or virtual network, you can select either of those options instead.
- **Storage Account.** The disk for your virtual machine that gets created from the image you select must be assigned to a storage account. You can create a new one for your virtual machine or select a storage account that you created earlier.
- **Availability Set.** An optional step is to create an availability set or assign the virtual machine to an existing availability set to ensure high availability. That way, if your physical server is taken

down for maintenance, every service included in the availability set will be accessible through other physical hosts. When multiple virtual machines belong to the same availability set, the Microsoft Azure Fabric Controller places virtual machines across multiple racks, each of which can be considered a fault domain and separate physical points of failure, as shown in Figure 3. In addition, each virtual machine is placed in a separate upgrade domain so that virtual machines in the same availability set are not taken offline at the same time during host updates.




Figure 3: Virtual Machines Assigned to Separate Availability Sets

Last, you configure endpoints for your virtual machine. The image you select has one or more preconfigured endpoints. For example, if the image runs a Windows operating system, the Remote Desktop and PowerShell endpoints are created with public and private ports to use when you or your applications connect to the virtual machine. You can add more endpoints if necessary, such as FTP, SSH, or SMTP for example, and specify whether to assign the TCP or UDP protocol to the endpoint. Additionally, you can change the public and private port numbers from the defaults to alternate values. You can also take advantage of advanced networking features in Microsoft Azure Virtual Network to connect your virtual machine to an on-premises datacenter or a single client machine using an IPsec connection.

When you complete the configurations steps, the virtual machine is provisioned in Microsoft Azure. You can monitor this process, which typically takes only a few minutes, in the Microsoft Azure Management Portal. When provisioning is complete, the virtual machine starts. When its status changes to Running, your new virtual machine is available for you to connect to it as described in the “Using a Virtual Machine” section of this document.

If you are not ready to start working with your virtual machine, select it in the list on the Virtual Machines page in Microsoft Azure Management Portal, and then click Shut Down. That way, you can save usage costs, although there is an ongoing cost associated with the disk in your Microsoft Azure Storage account.

Note: Specific steps for creating a Windows virtual machine from the Image Gallery are available at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-windows-tutorial/>. If you

want to create a Linux virtual machine instead, refer to the instructions at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-tutorial/>.

Capturing a Microsoft Azure Virtual Machine

If you have an existing Windows-server based Microsoft Azure Virtual Machine for which you want to create an image, you can use the Sysprep utility to put the virtual machine into the necessary state. To do this, open a command prompt window, navigate to the `system32\sysprep` directory in the Windows directory of the virtual machine, and execute `sysprep.exe`. In the Sysprep tool, select the Enter System Out-of-Box Experience (OOBE) action, select the Generalize check box, and select Quit as the Shutdown Option. When the `sysprep succeeded.tag` appears in the same folder as the Sysprep tool, log off the virtual machine and then use the Shut Down button on the command bar to stop the virtual machine. When it has successfully stopped, click the Capture button to generate the image. You then name the image and confirm that you have run the Sysprep utility. When you are ready to use the image, you can find it on the My Images section of the Image Gallery described in the “Creating a Virtual Machine” section of this document.

For more detailed instructions, see <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-capture-image-windows-server/>.

Using SUSE Studio to Create a Virtual Machine Image

[SUSE Studio](#) is a web-based service that you can use to create a virtual machine (called an appliance by SUSE Studio) that you can then transfer to run in Microsoft Azure. SUSE Studio offers a wizard interface to simplify the process of creating your virtual machine. You start by selecting an open source operating system, such as SUSE Linux Enterprise 11 SP3, a 32-bit or 64-bit architecture, and a name for your virtual machine. Then you select packages to install, such as Perl or Python. Next you configure the locale, time zone, firewall, and users for the virtual machine in addition to allocating memory and other configuration tasks. You can also specify files to upload into the virtual machine.

After completing the configuration, you use a Build command to produce the virtual machine in one or more specified formats, including a Microsoft Azure-ready image. An Upload link displays for you to move the prepared image into your Microsoft Azure account after the build process completes. Instructions are available to help you download a settings file that you then upload to your SUSE account. Then you specify the Microsoft Azure storage account into which you want the image transferred, and then click the link to go to the SUSE Studio’s dashboard where you click the Upload button to send your image to Microsoft Azure. After the upload completes, you follow the instructions for creating a virtual machine from the Virtual Machine Image Gallery. Your uploaded image displays on the My Images tab of the Images Gallery wizard.

For more information about limitations of Microsoft Azure images created by using SUSE Studio, see <http://susestudio.com/help/use/windows-azure.html>.

Uploading a Virtual Machine to Microsoft Azure

Whether you want to permanently run one of your on-premises virtual machines in Microsoft Azure or simply want to test a configuration outside of your production environment, you can upload a copy of that virtual machine as a page blob to your Microsoft Azure Storage account, as shown in Figure 4. You can even move the virtual machine back on-premises later if necessary. The VHD can be a sysprepped fixed disk VHD or a customized Linux image based on supported distributions.




Figure 4: Mobility of a VHD between On-Premises and Microsoft Azure

Uploading an Existing Windows Virtual Machine

If you are already using a Hyper-V virtual machine, you can upload the associated VHD file to Microsoft Azure and begin using the virtual machine from the cloud. If you are using VMware virtualization, you can use the System Center Virtual Machine Manager to [convert a VMDK file to a VHD file](#) or use [Microsoft Virtual Machine Converter](#). Much like you do with a Microsoft Azure virtual machine, you generalize the operating system of your on-premises virtual machine by using the Sysprep utility. You launch the utility from the system32\sysprep directory in the Windows directory of the virtual machine. Then you select the Enter System Out-of-Box Experience (OOBE) action, select the Generalize check box, and select Shutdown as the Shutdown Option.

Next you must create a new storage account in the Microsoft Azure Management Portal if you do not have one already available. Click New, click Storage, and then click Quick Create. You need to supply a subdomain name as well as a location or affinity group, and choose whether to use geo-replication. You then use Microsoft Azure PowerShell cmdlets, `Get-AzurePublishSettingsFile`, `Import-AzurePublishSettingsFile`, and `Add-AzureVhd`, to upload the VHD file for your virtual machine to your storage account.

After the VHD is uploaded, you can go to the Virtual Machines page in the Microsoft Azure Management Portal, click the Images link above the list of virtual machines, and click Create to set up a new image based on your VHD. You must provide a name for the image, an optional description, the VHD location (which you can obtain by browsing your cloud storage accounts), and the operating system family (Windows, in this case).

For a complete list of steps to perform, refer to <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-create-upload-vhd-windows-server/>.

Uploading a Local Linux Virtual Machine

There are several steps required to prepare and upload a local Linux virtual machine and different steps required for different Linux distributions. At the conclusion of the preparation process, regardless of Linux version, you must install the [Microsoft Azure Linux Agent](#). In addition, you must provision a Microsoft Azure storage account. In the Microsoft Azure Management Portal, click New, click Storage, and then click Quick Create. Next, type in a subdomain name, select a location or affinity group, and specify whether to

use geo-replication. Then just as you do for a Windows VHD, you use Microsoft Azure PowerShell cmdlets, `Get-AzurePublishSettingsFile`, `Import-AzurePublishSettingsFile`, and `Add-AzureVhd`, to upload the VHD file for your virtual machine to your storage account.

You can review the details for uploading a Linux virtual machine at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-create-upload-vhd/>.

Working with a Virtual Machine

The disk for the virtual machine's operating system is limited to 127 GB. The virtual machine also has a disk allocated to temporary storage for data, but this data does not persist if the virtual machine is relocated to other hardware for any reason. Therefore, do not store data in this location if you cannot easily recreate it. These disks are assigned to the drive letters C and D, respectively. If necessary, you can add a data disk to increase your storage capacity and to hold additional applications that you want to install. You can assign any drive letter, other than C or D, to each data disk that you add.

Connecting to a Windows Virtual Machine

If the virtual machine is not currently running, you must first select it in the virtual machines list in the Microsoft Azure Management Portal and then click **Start** in the command bar at the bottom of the page. After the virtual machine is running, click the **Connect** button to generate the remote desktop connection file, click **Open** to use the file, and then click **Connect** to establish the connection to the virtual machine. Type in the administrator login name and password to authenticate, and then click **Yes** to verify the virtual machine's identity. At this point, you can interact with the virtual machine in the same way you would use a local server.

Note: If your firewall has a rule that does not allow you to connect to the RDP port of your virtual machine, you may change the public port of your virtual machine's endpoint.

Connecting to a Linux Virtual Machine

Start the virtual machine if it is not currently running. Then use an SSH client, such as [PuTTY](#) for a Windows client or [OpenSSH](#) for a Linux client, to connect to your virtual machine. To make this connection, you must know the host name and port for the virtual machine. You can get this information by clicking the virtual machine's name in the Microsoft Azure Management Portal, clicking the **Dashboard** link, and then locating the **SSH Details** section in the lower right corner of the page. After connecting to the virtual machine, you can use it just like any other local server.

Attaching a Data Disk

If you need additional storage for your virtual machine, you can attach an empty disk or a VHD that you have preloaded with data. Select your virtual machine in the list of instances, click **Attach** in the command bar, and then click **Attach Empty Disk** or **Attach Disk** as appropriate. For an existing disk, you select the disk in storage and specify a **Host Cache Preference** option. For a new empty disk, a storage location and a file name are automatically assigned, but you can override the file name if you prefer. You must specify a size for the disk in gigabytes, and then select from one of the following **Host Cache Preference** options:

- **None.** Host caching is disabled. This is the default setting for a data disk.
- **Read Only.** Host caching is enabled only for read operations as a performance optimization.
- **Read Write.** Host caching is enabled for both reads and writes. This is the default setting for an operating system disk.

When the disk is attached, you must next initialize the disk by using the Disk Management utility if you are using a Windows virtual machine, and then create a new simple volume. Once the volume is created, the disk is available for use by your applications. If you are using a Linux virtual machine, you use the `fdisk` command to create a new device and a new partition, create a file system on the new partition, add a directory, and then mount the drive.

You can review the steps to attach a data disk for either operating system at <http://azure.microsoft.com/en-us/documentation/articles/storage-windows-attach-disk/>.

Managing a Virtual Machine

You have three options for managing virtual machines—the management portal, scripting, and the Rest API.



- You can use the graphical interface of the Microsoft Azure Management Portal to perform various tasks related to managing your virtual machine. You can start and stop a virtual machine, connect to it, attach and detach disks, capture the VHD as an image, or delete the VHD. In addition, you can configure a different size for the virtual machine, assign it to an availability set, or add new endpoints.



- If you prefer, you can also download and install Microsoft Azure PowerShell cmdlets to use scripts for management tasks. Instructions are available at <http://azure.microsoft.com/en-us/documentation/articles/install-configure-powershell/>. As an example, you can use the [Add-AzureDataDisk](#) cmdlet to add a new data disk to your virtual machine.
- Another similar option is the CLI available for all platforms. You can learn more about the CLI for Windows, Linux, or Mac at <http://azure.microsoft.com/en-us/documentation/articles/xplat-cli/>.



- You can use various language-specific SDKs to manage Microsoft Azure. See the language-specific Developer Centers documentation at <http://azure.microsoft.com/en-us/documentation/>.
- Although the portal, PowerShell cmdlets, CLI, and language-specific SDKs all use common REST APIs behind the scenes, you can also access [Blob Service REST APIs](#) directly if you prefer. With this API, you can manage the containers and blobs in your Storage account, including uploading, downloading, or deleting your VHDs.

© 2013 Microsoft Corporation. All rights reserved. Except where otherwise noted, content on this site is licensed under a Creative Commons Attribution-NonCommercial 3.0 License.

AWS PUBLIC SECTOR ACCESS POLICY (STATE)

(Last Updated January 4, 2013)

This AWS Access Policy ("**Access Policy**") governs your access to and use of the Services (as defined below) of Amazon Web Services, Inc. ("**AWS**") provided to you by your systems integrator, reseller, or services provider ("**Provider**"). It sets out the additional rules, conditions and restrictions that apply to you or the entity you represent ("**you**") for use of the Services. In this Access Policy, "**we**", "**us**", or "**our**" means AWS and any of its affiliates. Please see Section 10 for definitions of capitalized terms.

1. Use of the Services.

1.1 Generally. You are provided access to the Services by your Provider. Your use of and access to the Services are governed by the agreement between you and Provider. This Access Policy supplements the terms of such agreement and may be updated by us from time to time. AWS Service Level Agreements do not apply to your use of the Services. Your continued access to and use of the Services is conditioned on your compliance with all laws, rules, regulations, policies and instructions applicable to your use of the Services, including the Policies.

1.2 Account Keys. Provider may provide you with AWS account keys which will allow you to directly access the Services via Provider's account(s). We are not responsible for any activities that occur under these account keys, regardless of whether the activities are undertaken by you, Provider or a third party (including your employees, contractors or agents) and we are also not responsible for unauthorized access to the account.

1.3 Third Party Materials. Through the use of Provider's AWS account(s), you may have access to Third Party Materials, such as software applications provided by third parties, which are made available directly to you by other companies or individuals under separate terms and conditions, including separate fees and charges. Your use of any Third Party Materials is at your sole risk.

2. Your Responsibilities.

2.1 Your Materials. You are solely responsible for the development, content, operation, maintenance, and use of Your Materials with the Services. For example, you are solely responsible for:

- (a) the technical operation of Your Materials, including ensuring that calls you make to any Service are compatible with then-current application program interfaces for that Service;
- (b) compliance of Your Materials with the Acceptable Use Policy, the other Policies, and the law;
- (c) any claims relating to Your Materials;
- (d) properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Materials violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act;
- (e) any action that you permit, assist or facilitate any person or entity to take related to this Access Policy, Your Materials or use of the Services; and
- (f) End Users' use of Your Materials and the Services and ensuring that End Users comply with your obligations under this Access Policy and that the terms of your agreement with each End User are consistent with this Access Policy.

2.2 Other Security and Backup. You or Provider are solely responsible for properly configuring and using the Services and taking steps to maintain appropriate security, protection and backup of Your Materials, including using encryption technology to protect Your Materials from unauthorized access and routinely archiving Your Materials.

2.3 End User Violations. If you become aware of any violation of your obligations under this Access Policy by an End User, you will immediately terminate such End User's access to Your Materials and the Services.

3. Service Interruption.

3.1 General. We may suspend the AWS account(s) through which you access the Services immediately if we determine your or an End User's use of the Services (i) violates the terms of this Access Policy (including the Acceptable Use Policy or Service Terms); (ii) poses a security risk to the Services or any other AWS customer, (iii) may harm our systems or the systems or Materials of any other AWS customer; or (iv) may subject us to liability as a result of any of the foregoing. We will provide notice of any suspension as soon as practicable to Provider, who is solely responsible for providing any notices to you under your agreement with them.

3.2 Scope of Interruption. To the extent practicable, we will (i) suspend your right to access or use only those instances, data, or portions of the Services that caused the suspension, and (ii) limit the suspension to those accounts that caused the suspension. If commercially feasible, access to the Services will be restored once the conditions or circumstances giving rise to the suspension have been removed or corrected. Nothing in this Section 3 will operate to limit your rights or remedies otherwise available to you against Provider under your agreement with them or applicable law.

4. Proprietary Rights.

4.1 Services. As between you and us, we or our licensors own and reserve all right, title, and interest in and to the Services. You have the right to use the Services solely as a licensee of Provider in accordance with this Access Policy and the agreement between you and Provider. We have no obligation to provide the Service to you under this Access Policy, so you must look exclusively to Provider and your agreement with Provider regarding such obligation. Except as expressly provided in this Section 4, you obtain no rights to the Services, the AWS Materials or any Third Party Materials.

4.2 Materials. As a part of the Services, you may have access to AWS Materials and Third Party Materials, which may be subject to additional terms and conditions (including the Terms of Use and Apache Software License). By using those materials, you are subject to such additional terms. You are solely responsible for securing any necessary approvals for the download and use of such materials.

4.3 Restrictions. Neither you nor any End User may use the Services in any manner or for any purpose other than as expressly permitted by this Access Policy and the agreement between you and Provider. Neither you nor any End User may, or may attempt to, (a) modify, alter, tamper with, repair, or otherwise create derivative works of any software included in the Services (except to the extent software included in the Services are provided to you under a separate license that expressly permits the creation of derivative works), (b) reverse engineer, disassemble, or decompile the software included in the Services or apply any other process or procedure to derive the source code of any software included in the Services, or (c) access or use the Services in a way intended to avoid incurring fees or exceeding usage limits or quotas. All rights and access granted to you with respect to the Services are conditioned on your continued compliance with this Access Policy, and you will immediately discontinue your use of the Services if you cannot comply with this Access Policy. You will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, vendors, business partners, or licensors, any patent infringement or other intellectual property infringement claim regarding any Services or AWS Materials that you have used.

4.4 Suggestions. If you provide any Suggestions to us when using the Services, you hereby grant to AWS and its affiliates a perpetual, irrevocable, non-exclusive, worldwide, royalty-free right and license to reproduce, distribute, make derivative works based upon, publicly display, publicly perform, make, have made, use, sell, offer for sale, and import the Suggestions, including the right to sublicense such rights through multiple tiers, alone or in combination.

4.5 Government Rights. If you are using the Services on behalf of the government and these terms fail to meet the government's needs or are inconsistent in any respect with federal or state law, you will immediately discontinue your use of the Services (including any AWS Materials).

5. Representations and Warranties. You represent and warrant that (a) you and your End Users' use of the Services (including any use by your employees and personnel) will not violate this Access Policy; (b) you or your

licensors own all right, title, and interest in and to Your Materials; (c) Your Materials (including the use, development, design, production, advertising, or marketing of your Materials) or the combination of your Materials with other applications, content or processes, do not and will not violate any applicable laws or infringe or misappropriate any third-party rights; and (d) your use of the Services will not cause harm to any End User.

6. Indemnification. Except to the extent prohibited by applicable state law, you will defend, indemnify, and hold harmless us, our licensors and each of our respective employees, officers, directors, and representatives from and against any claims, damages, losses, liabilities, costs, and expenses (including reasonable attorneys' fees) arising out of or relating to any third party claim concerning: (a) your or any End Users' use of the Services (including use by your employees and personnel); (b) Your Materials or the combination of Your Materials with other applications, content or processes, including any claim involving alleged infringement or misappropriation of third-party rights or the use, development, design, production, advertising or marketing of Your Materials; or (c) a dispute between you and any End User. If your ability to comply with the foregoing provision is limited to any extent by the absence of appropriations or government authorization, you will make good faith efforts to obtain sufficient appropriations or authorization for any liabilities arising under this Section 5.

7. Disclaimers. WE PROVIDE THE SERVICES ON AN "AS IS" BASIS TO PROVIDER. WE AND OUR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND TO YOU, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE REGARDING THE SERVICES OR ANY THIRD PARTY MATERIALS, INCLUDING ANY WARRANTY THAT THE SERVICES OR THIRD PARTY MATERIALS WILL BE UNINTERRUPTED, ERROR FREE OR FREE OF HARMFUL COMPONENTS, OR THAT ANY MATERIALS, INCLUDING YOUR MATERIALS OR THE THIRD PARTY MATERIALS, WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED. EXCEPT TO THE EXTENT PROHIBITED BY LAW, WE AND OUR LICENSORS DISCLAIM ALL WARRANTIES, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR USAGE OF TRADE.

8. Limitations of Liability. YOU MUST LOOK SOLELY TO PROVIDER AND YOUR AGREEMENT WITH THEM REGARDING ANY CLAIMS OR DAMAGES RELATED TO THE SERVICES. WE AND OUR AFFILIATES OR LICENSORS WILL NOT BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES (INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA), EVEN IF A PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. FURTHER, NEITHER WE NOR ANY OF OUR LICENSORS WILL BE RESPONSIBLE FOR ANY COMPENSATION, REIMBURSEMENT, OR DAMAGES ARISING IN CONNECTION WITH: (A) YOUR INABILITY TO USE THE SERVICES, INCLUDING AS A RESULT OF ANY (I) SUSPENSION OF YOUR USE OF OR ACCESS TO THE SERVICES, (II) OUR DISCONTINUATION OF ANY OR ALL OF THE SERVICES, OR, (III) ANY UNANTICIPATED OR UNSCHEDULED DOWNTIME OF ALL OR A PORTION OF THE SERVICES FOR ANY REASON; OR (B) ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION, DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR MATERIALS OR OTHER DATA THAT YOU OR ANY END USER SUBMITS OR USES IN CONNECTION WITH THE SERVICES (INCLUDING AS A RESULT OF YOUR OR ANY END USERS' ERRORS, ACTS OR OMISSIONS).

9. Miscellaneous.

9.1 Governing Law; Venue. Except to the extent prohibited by applicable state law, the laws of the State of Washington, without reference to conflict of law rules, govern this Access Policy and any dispute of any sort that might arise between you and us. You irrevocably consent to exclusive jurisdiction and venue of the federal courts located in King County, Washington with respect to any dispute arising in connection with the Services or this Access Policy. We may seek injunctive or other relief in any state, federal, or national court of competent jurisdiction for any actual or alleged infringement of our or any third party's intellectual property or other proprietary rights. The United Nations Convention for the International Sale of Goods does not apply.

9.2 Entire Agreement. This Access Policy is the entire agreement between you and us regarding the Services, including any materials. It supersedes all prior or contemporaneous representations, understandings, agreements, or communications between you and us, whether written or verbal, regarding the subject matter of this Access

Policy. If the terms of this document are inconsistent with the terms contained in your agreement with Provider, the terms contained in this document will control. We will not be bound by, and specifically object to, any term, condition or other provision which is different from or in addition to the provisions of this Access Policy (whether or not it would materially alter it) and which is submitted by you in any order, receipt, acceptance, confirmation, correspondence or other document.

9.3 Survival. The following provisions will survive any termination of your use of the Services: Sections 2.1, 4, 5, 6, 7, 8, 9 and 10.

10. Definitions.

“Acceptable Use Policy” means the policy currently available at <http://aws.amazon.com/aup>, as it may be updated by us from time to time.

“AWS Materials” means Materials we make available in connection with the Services or on the AWS Site to allow access to and use of the Services, including WSDLs; Documentation; sample code; software libraries; command line tools; and other related technology. AWS Materials does not include the Services.

“AWS Service Level Agreement” means all service level agreements that we offer with respect to the Services and post on the AWS Site, as they may be updated by us from time to time.

“AWS Site” means <http://aws.amazon.com> and any successor or related site designated by us.

“Documentation” means the developer guides, getting started guides, user guides, quick reference guides, and other technical and operations manuals, instructions and specifications for the Services currently located at <http://aws.amazon.com/documentation>, as such documentation may be updated by us from time to time.

“End User” means any individual or entity that directly or indirectly through another user: (a) accesses or uses Your Materials; or (b) otherwise accesses or uses the Services through you.

“Materials” means software (including machine images), data, text, audio, video, images or other content.

“Policies” means the Acceptable Use Policy, the Terms of Use, the Service Terms, all restrictions described in the AWS Materials and on the AWS Site, and any other policy or terms referenced in or incorporated into this Access Policy.

“Services” means, collectively or individually (as applicable), the web services made commercially available by us to Provider for use under this Access Policy, including (as applicable) those web services described in the Service Terms.

“Service Terms” means the rights and restrictions for particular Services located at <http://aws.amazon.com/serviceterms>, as they may be updated by us from time to time.

“Suggestions” means all suggested improvements to the Services or AWS Materials that you provide to us.

“Terms of Use” means the terms of use located at <http://aws.amazon.com/terms/>, as they may be updated by us from time to time.

“Third Party Materials” means Materials made available to you by any third party on the AWS Site or in conjunction with the Services.

“Your Materials” means Materials you or any End User (a) run on the Services, (b) cause to interface with the Services, or (c) upload to the Services or otherwise transfer, process, use or store in connection with the Services.

AWS PUBLIC SECTOR ACCESS POLICY (STATE)

(Last Updated January 4, 2013)

This AWS Access Policy ("**Access Policy**") governs your access to and use of the Services (as defined below) of Amazon Web Services, Inc. ("**AWS**") provided to you by your systems integrator, reseller, or services provider ("**Provider**"). It sets out the additional rules, conditions and restrictions that apply to you or the entity you represent ("**you**") for use of the Services. In this Access Policy, "**we**", "**us**", or "**our**" means AWS and any of its affiliates. Please see Section 10 for definitions of capitalized terms.

1. Use of the Services.

1.1 Generally. You are provided access to the Services by your Provider. Your use of and access to the Services are governed by the agreement between you and Provider. This Access Policy supplements the terms of such agreement and may be updated by us from time to time. AWS Service Level Agreements do not apply to your use of the Services. Your continued access to and use of the Services is conditioned on your compliance with all laws, rules, regulations, policies and instructions applicable to your use of the Services, including the Policies.

1.2 Account Keys. Provider may provide you with AWS account keys which will allow you to directly access the Services via Provider's account(s). We are not responsible for any activities that occur under these account keys, regardless of whether the activities are undertaken by you, Provider or a third party (including your employees, contractors or agents) and we are also not responsible for unauthorized access to the account.

1.3 Third Party Materials. Through the use of Provider's AWS account(s), you may have access to Third Party Materials, such as software applications provided by third parties, which are made available directly to you by other companies or individuals under separate terms and conditions, including separate fees and charges. Your use of any Third Party Materials is at your sole risk.

2. Your Responsibilities.

2.1 Your Materials. You are solely responsible for the development, content, operation, maintenance, and use of Your Materials with the Services. For example, you are solely responsible for:

- (a) the technical operation of Your Materials, including ensuring that calls you make to any Service are compatible with then-current application program interfaces for that Service;
- (b) compliance of Your Materials with the Acceptable Use Policy, the other Policies, and the law;
- (c) any claims relating to Your Materials;
- (d) properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Materials violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act;
- (e) any action that you permit, assist or facilitate any person or entity to take related to this Access Policy, Your Materials or use of the Services; and
- (f) End Users' use of Your Materials and the Services and ensuring that End Users comply with your obligations under this Access Policy and that the terms of your agreement with each End User are consistent with this Access Policy.

2.2 Other Security and Backup. You or Provider are solely responsible for properly configuring and using the Services and taking steps to maintain appropriate security, protection and backup of Your Materials, including using encryption technology to protect Your Materials from unauthorized access and routinely archiving Your Materials.

2.3 End User Violations. If you become aware of any violation of your obligations under this Access Policy by an End User, you will immediately terminate such End User's access to Your Materials and the Services.

3. Service Interruption.

3.1 General. We may suspend the AWS account(s) through which you access the Services immediately if we determine your or an End User's use of the Services (i) violates the terms of this Access Policy (including the Acceptable Use Policy or Service Terms); (ii) poses a security risk to the Services or any other AWS customer, (iii) may harm our systems or the systems or Materials of any other AWS customer; or (iv) may subject us to liability as a result of any of the foregoing. We will provide notice of any suspension as soon as practicable to Provider, who is solely responsible for providing any notices to you under your agreement with them.

3.2 Scope of Interruption. To the extent practicable, we will (i) suspend your right to access or use only those instances, data, or portions of the Services that caused the suspension, and (ii) limit the suspension to those accounts that caused the suspension. If commercially feasible, access to the Services will be restored once the conditions or circumstances giving rise to the suspension have been removed or corrected. Nothing in this Section 3 will operate to limit your rights or remedies otherwise available to you against Provider under your agreement with them or applicable law.

4. Proprietary Rights.

4.1 Services. As between you and us, we or our licensors own and reserve all right, title, and interest in and to the Services. You have the right to use the Services solely as a licensee of Provider in accordance with this Access Policy and the agreement between you and Provider. We have no obligation to provide the Service to you under this Access Policy, so you must look exclusively to Provider and your agreement with Provider regarding such obligation. Except as expressly provided in this Section 4, you obtain no rights to the Services, the AWS Materials or any Third Party Materials.

4.2 Materials. As a part of the Services, you may have access to AWS Materials and Third Party Materials, which may be subject to additional terms and conditions (including the Terms of Use and Apache Software License). By using those materials, you are subject to such additional terms. You are solely responsible for securing any necessary approvals for the download and use of such materials.

4.3 Restrictions. Neither you nor any End User may use the Services in any manner or for any purpose other than as expressly permitted by this Access Policy and the agreement between you and Provider. Neither you nor any End User may, or may attempt to, (a) modify, alter, tamper with, repair, or otherwise create derivative works of any software included in the Services (except to the extent software included in the Services are provided to you under a separate license that expressly permits the creation of derivative works), (b) reverse engineer, disassemble, or decompile the software included in the Services or apply any other process or procedure to derive the source code of any software included in the Services, or (c) access or use the Services in a way intended to avoid incurring fees or exceeding usage limits or quotas. All rights and access granted to you with respect to the Services are conditioned on your continued compliance with this Access Policy, and you will immediately discontinue your use of the Services if you cannot comply with this Access Policy. You will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, vendors, business partners, or licensors, any patent infringement or other intellectual property infringement claim regarding any Services or AWS Materials that you have used.

4.4 Suggestions. If you provide any Suggestions to us when using the Services, you hereby grant to AWS and its affiliates a perpetual, irrevocable, non-exclusive, worldwide, royalty-free right and license to reproduce, distribute, make derivative works based upon, publicly display, publicly perform, make, have made, use, sell, offer for sale, and import the Suggestions, including the right to sublicense such rights through multiple tiers, alone or in combination.

4.5 Government Rights. If you are using the Services on behalf of the government and these terms fail to meet the government's needs or are inconsistent in any respect with federal or state law, you will immediately discontinue your use of the Services (including any AWS Materials).

5. Representations and Warranties. You represent and warrant that (a) you and your End Users' use of the Services (including any use by your employees and personnel) will not violate this Access Policy; (b) you or your

licensors own all right, title, and interest in and to Your Materials; (c) Your Materials (including the use, development, design, production, advertising, or marketing of your Materials) or the combination of your Materials with other applications, content or processes, do not and will not violate any applicable laws or infringe or misappropriate any third-party rights; and (d) your use of the Services will not cause harm to any End User.

6. Indemnification. Except to the extent prohibited by applicable state law, you will defend, indemnify, and hold harmless us, our licensors and each of our respective employees, officers, directors, and representatives from and against any claims, damages, losses, liabilities, costs, and expenses (including reasonable attorneys' fees) arising out of or relating to any third party claim concerning: (a) your or any End Users' use of the Services (including use by your employees and personnel); (b) Your Materials or the combination of Your Materials with other applications, content or processes, including any claim involving alleged infringement or misappropriation of third-party rights or the use, development, design, production, advertising or marketing of Your Materials; or (c) a dispute between you and any End User. If your ability to comply with the foregoing provision is limited to any extent by the absence of appropriations or government authorization, you will make good faith efforts to obtain sufficient appropriations or authorization for any liabilities arising under this Section 5.

7. Disclaimers. WE PROVIDE THE SERVICES ON AN "AS IS" BASIS TO PROVIDER. WE AND OUR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND TO YOU, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE REGARDING THE SERVICES OR ANY THIRD PARTY MATERIALS, INCLUDING ANY WARRANTY THAT THE SERVICES OR THIRD PARTY MATERIALS WILL BE UNINTERRUPTED, ERROR FREE OR FREE OF HARMFUL COMPONENTS, OR THAT ANY MATERIALS, INCLUDING YOUR MATERIALS OR THE THIRD PARTY MATERIALS, WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED. EXCEPT TO THE EXTENT PROHIBITED BY LAW, WE AND OUR LICENSORS DISCLAIM ALL WARRANTIES, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR USAGE OF TRADE.

8. Limitations of Liability. YOU MUST LOOK SOLELY TO PROVIDER AND YOUR AGREEMENT WITH THEM REGARDING ANY CLAIMS OR DAMAGES RELATED TO THE SERVICES. WE AND OUR AFFILIATES OR LICENSORS WILL NOT BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES (INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA), EVEN IF A PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. FURTHER, NEITHER WE NOR ANY OF OUR LICENSORS WILL BE RESPONSIBLE FOR ANY COMPENSATION, REIMBURSEMENT, OR DAMAGES ARISING IN CONNECTION WITH: (A) YOUR INABILITY TO USE THE SERVICES, INCLUDING AS A RESULT OF ANY (I) SUSPENSION OF YOUR USE OF OR ACCESS TO THE SERVICES, (II) OUR DISCONTINUATION OF ANY OR ALL OF THE SERVICES, OR, (III) ANY UNANTICIPATED OR UNSCHEDULED DOWNTIME OF ALL OR A PORTION OF THE SERVICES FOR ANY REASON; OR (B) ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION, DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR MATERIALS OR OTHER DATA THAT YOU OR ANY END USER SUBMITS OR USES IN CONNECTION WITH THE SERVICES (INCLUDING AS A RESULT OF YOUR OR ANY END USERS' ERRORS, ACTS OR OMISSIONS).

9. Miscellaneous.

9.1 Governing Law; Venue. Except to the extent prohibited by applicable state law, the laws of the State of Washington, without reference to conflict of law rules, govern this Access Policy and any dispute of any sort that might arise between you and us. You irrevocably consent to exclusive jurisdiction and venue of the federal courts located in King County, Washington with respect to any dispute arising in connection with the Services or this Access Policy. We may seek injunctive or other relief in any state, federal, or national court of competent jurisdiction for any actual or alleged infringement of our or any third party's intellectual property or other proprietary rights. The United Nations Convention for the International Sale of Goods does not apply.

9.2 Entire Agreement. This Access Policy is the entire agreement between you and us regarding the Services, including any materials. It supersedes all prior or contemporaneous representations, understandings, agreements, or communications between you and us, whether written or verbal, regarding the subject matter of this Access

Policy. If the terms of this document are inconsistent with the terms contained in your agreement with Provider, the terms contained in this document will control. We will not be bound by, and specifically object to, any term, condition or other provision which is different from or in addition to the provisions of this Access Policy (whether or not it would materially alter it) and which is submitted by you in any order, receipt, acceptance, confirmation, correspondence or other document.

9.3 Survival. The following provisions will survive any termination of your use of the Services: Sections 2.1, 4, 5, 6, 7, 8, 9 and 10.

10. Definitions.

“Acceptable Use Policy” means the policy currently available at <http://aws.amazon.com/aup>, as it may be updated by us from time to time.

“AWS Materials” means Materials we make available in connection with the Services or on the AWS Site to allow access to and use of the Services, including WSDLs; Documentation; sample code; software libraries; command line tools; and other related technology. AWS Materials does not include the Services.

“AWS Service Level Agreement” means all service level agreements that we offer with respect to the Services and post on the AWS Site, as they may be updated by us from time to time.

“AWS Site” means <http://aws.amazon.com> and any successor or related site designated by us.

“Documentation” means the developer guides, getting started guides, user guides, quick reference guides, and other technical and operations manuals, instructions and specifications for the Services currently located at <http://aws.amazon.com/documentation>, as such documentation may be updated by us from time to time.

“End User” means any individual or entity that directly or indirectly through another user: (a) accesses or uses Your Materials; or (b) otherwise accesses or uses the Services through you.

“Materials” means software (including machine images), data, text, audio, video, images or other content.

“Policies” means the Acceptable Use Policy, the Terms of Use, the Service Terms, all restrictions described in the AWS Materials and on the AWS Site, and any other policy or terms referenced in or incorporated into this Access Policy.

“Services” means, collectively or individually (as applicable), the web services made commercially available by us to Provider for use under this Access Policy, including (as applicable) those web services described in the Service Terms.

“Service Terms” means the rights and restrictions for particular Services located at <http://aws.amazon.com/serviceterms>, as they may be updated by us from time to time.

“Suggestions” means all suggested improvements to the Services or AWS Materials that you provide to us.


“Terms of Use” means the terms of use located at <http://aws.amazon.com/terms/>, as they may be updated by us from time to time.

“Third Party Materials” means Materials made available to you by any third party on the AWS Site or in conjunction with the Services.

“Your Materials” means Materials you or any End User (a) run on the Services, (b) cause to interface with the Services, or (c) upload to the Services or otherwise transfer, process, use or store in connection with the Services.

Trusted Cloud: Microsoft Azure Security, Privacy, and Compliance


April 2015






Contents

Introduction	4
What customers want from cloud providers.....	5
Microsoft Azure: Built for trust	6
Security: Working to keep customer data safe.....	7
Security design and operations	7
Infrastructure protection.....	9
Network protection.....	10
Data protection	11
Identity and access.....	12
Privacy: Customers own and control their data	12
Customers are in control of their data.....	14
Transparency	15
Compliance: Azure conforms to global standards.....	16
Additional resources	18



Introduction

With the emergence of cloud computing, today's IT organizations are playing an increasingly important role in driving business strategy. While cost reduction is still a top priority, scalability and business agility have stepped to the forefront for IT decision makers. As a result, spending on cloud solutions is expected to grow 30 percent from 2013 to 2018, compared with 5 percent overall growth for enterprise IT. And cloud services are keeping pace¹. Analysts expect to see a ten-fold increase in the number of cloud-based solutions on the market in the next four to five years.



Sources:

- 1 Forbes, "Roundup of Cloud Computing Forecasts and Market Estimates, 2015," 1/24/2015. <http://aka.ms/forbes-cloud-2015>
- 2 Intel/McAfee, "Net Losses: Estimating the Global Cost of Cybercrime," June 2014. <http://aka.ms/mcafee-cybercrime-report>
- 3 UK Dept. for Business, Innovation and Skills, "2014 Information Security Breaches Survey," http://aka.ms/uk-gov_breach-survey
- 4 PWC, "Global State of Information Security Survey: 2015," <http://aka.ms/pwc-cybercrime>
- 5 Gemalto, 2014 Breach Level Index Report
- 6 McKinsey & Company, report for World Economic Forum, Jan. 2014

Still, many CIOs hesitate to fully embrace a cloud-first approach. Their hesitation stems in part from anxiety over a wide range of privacy and security related issues. Large-scale data breaches dominated headlines in 2014 and continue in the news today, raising a critical question for IT leaders everywhere: How can organizations build scalable cloud solutions and increase business agility while taking the necessary steps to secure our data and ensure privacy and compliance across the enterprise?

Without a clear answer, security concerns threaten to stall innovation and stifle business growth. IT and business leaders need a trusted partner to bridge the gap between innovation and security. With the right technologies and processes, even the most complex enterprise can move to the cloud with confidence.

“71% of strategic buyers cite scalability, cost and business agility as the most important drivers for using cloud services.”

Gigaom Research

What customers want from cloud providers

Every business has different needs and every business will reap distinct benefits from cloud solutions. Still, customers of all kinds have the same basic concerns about moving to the cloud. They want to retain control of their data, and they want that data to be kept secure and private, all while maintaining transparency and compliance.

Secure our data. The scale and scope of intrusions are growing. In 2014, cyber criminals compromised more than a billion data records in more than 1500 breaches.⁷ In a 2014 report for the World Economic Forum⁸, McKinsey & Company estimated the risk of cyberattacks “could materially slow the pace of technology and business innovation with as much as \$3 trillion in aggregate impact.” In any security attack, target organizations are only as safe as their weakest link. If any component is not secured, then the entire system is at risk. While acknowledging that the cloud can provide increased data security and administrative control, IT leaders are still concerned that migrating to the cloud will leave them more vulnerable to hackers than their current in-house solutions.

Keep our data private. Cloud services raise unique privacy challenges for businesses. As companies look to the cloud to save on infrastructure costs and improve their flexibility, they also worry about losing control of where their data is stored, who is accessing it, and how it gets used. Since the revelations of widespread surveillance by the US government in 2013, privacy concerns have become more accentuated, and the cloud has come under greater scrutiny as a result.

Give us control. Even as they take advantage of the cloud to deploy more innovative solutions, companies are very concerned about losing control of their data. The recent disclosures of government agencies accessing customer data, through both legal and extralegal means, make some CIOs wary of storing their data in the cloud. Many companies are therefore looking to choose where their data resides in the cloud and to control what entities have visibility into that data.

Promote transparency. While security, privacy, and control are important to business decision makers, they also want the ability to independently verify how their data is being stored, accessed, and secured. Businesses understand that they cannot control what they cannot see. To create this sort of visibility for customers, cloud providers must offer transparency of their security, privacy and compliance practices and actions to give customers the information they need to make their own decisions.

⁷ Gemalto, 2014 Breach Level Index Report

⁸ McKinsey & Company, for World Economic Forum, Jan. 2014

Maintain compliance. As companies and government agencies expand their use of cloud technologies, the complexity and scope of standards and regulations continues to evolve. Companies need to know that their compliance standards will be met, and that compliance will evolve as regulations change over time.

Microsoft Azure: Built for trust

Microsoft Azure provides cloud services for a wide range of enterprise and government customers. The core of Microsoft Azure provides four primary functions on which customers build and manage virtual environments, applications, and associated configurations.

Microsoft Azure
Unified platform for modern business



Compute



Data storage



Network services



App services

Global physical infrastructure servers/ networks/ datacenters

- Stores over 10 trillion objects
- Handles on average 127,000 requests/second
- Peak of 880,000 requests/second



Microsoft, with its unique experience and scale, delivers these services to many of the world's leading enterprises and government agencies. Today, the Microsoft cloud infrastructure supports over 1 billion customers across our enterprise and consumer services in 140 countries and supports 10 languages and 24 currencies. Drawing on this history and scale, Microsoft has implemented software development with enhanced security, operational management, and threat mitigation practices, helping it to deliver services that achieve higher levels of security, privacy, and compliance than most customers could achieve on their own.

Microsoft shares best practices with government and commercial organizations and engages in broad security efforts through the creation of centers of excellence, including the Microsoft Digital Crimes Unit, Microsoft Security Response Center, and Microsoft Malware Protection Center.

Security: Working to keep customer data safe

Azure can help reduce the cost, complexity, and risk associated with security and compliance in the cloud. A survey funded by Microsoft and performed by ComScore⁹ found that while many organizations have initial concerns about moving to the cloud, a majority of cloud adopters reported that they achieved significant security benefits. These security benefits are reported because few individual organizations can replicate the technology and operational processes that Microsoft uses to help safeguard its enterprise cloud services and comply with a wide range of international standards. When companies use Azure, they benefit from Microsoft's unmatched scale and experience running compliant online services around the globe. Microsoft's expertise becomes the customer's expertise.


Initial concern



Realized benefit



Microsoft Cloud Experience:



Security Design and Operations

Secure cloud solutions are the result of comprehensive planning, innovative design, and efficient operations. Microsoft makes security a priority at every step, from code development to incident response.

Design for security from the ground up. Azure code development adheres to Microsoft's Security Development Lifecycle (SDL). The SDL is a software development process that helps developers build more secure software and address security compliance requirements while reducing development cost. The SDL became central to Microsoft's development practices a decade ago and is shared freely with the industry and customers. It embeds security requirements into systems and software through the planning, design, development, and deployment phases.

⁹ <http://aka.ms/twc-cloud-trust-study>

“We don’t have the resources to respond to security threats all day, 365 days a year, the way that Microsoft does.”

Bo Wandschneider,
CIO and Associate Vice Principal
Queen’s University (Canada)


Enhancing operational security. Azure adheres to a rigorous set of security controls that govern operations and support. Microsoft deploys combinations of preventive, defensive, and reactive controls including the following mechanisms to help protect against unauthorized developer and/or administrative activity:

- Tight access controls on sensitive data, including a requirement for two-factor smartcard-based authentication to perform sensitive operations.
- Combinations of controls that enhance independent detection of malicious activity.
- Multiple levels of monitoring, logging, and reporting.

Additionally, Microsoft conducts background verification checks of certain operations personnel and limits access to applications, systems, and network infrastructure in proportion to the level of background verification.

Assume breach. One key operational best practice that Microsoft uses to harden its cloud services is known as the “assume breach” strategy. A dedicated “red team” of software security experts simulates real-world attacks at the network, platform, and application layers, testing Azure’s ability to detect, protect against, and recover from breaches. By constantly challenging the security capabilities of the service, Microsoft can stay ahead of emerging threats.

Incident management and response. Microsoft has a global, 24x7 incident response service that works to mitigate the effects of attacks and malicious activity. The incident response team follows established procedures for incident management, communication, and recovery, and uses discoverable and predictable interfaces with internal and external partners alike. In the event of a security incident, the security team follows these five phases:



- **Identification:** If an event indicates a security issue, the incident is assigned a severity classification and appropriately escalated within Microsoft.
- **Containment:** The immediate priority of the escalation team is to ensure the incident is contained and data is safe.
- **Eradication:** After the situation is contained, the escalation team moves toward eradicating any damage caused by the security incident and identifies the root cause of the security issue.
- **Recovery:** Software or configuration updates are applied to the system and services are returned to full working capacity.
- **Lessons Learned:** Each security incident is analyzed to ensure the appropriate mitigations are applied to protect against future recurrence.

Infrastructure Protection

Azure infrastructure includes hardware, software, networks, administrative and operations staff, and the physical data centers that house it all. Azure addresses security risks across its infrastructure.

Physical security. Azure runs in geographically distributed Microsoft facilities, sharing space and utilities with other Microsoft Online Services. Each facility is designed to run 24x7x365 and employs various measures to help protect operations from power failure, physical intrusion, and network outages. These datacenters comply with industry standards (such as ISO 27001) for physical security and availability. They are managed, monitored, and administered by Microsoft operations personnel.

Monitoring and logging. Centralized monitoring, correlation, and analysis systems manage the large amount of information generated by devices within the Azure environment, providing continuous visibility and timely alerts to the teams that manage the service. Additional monitoring, logging, and reporting capabilities provide visibility to customers.

Perimeter 	Buildings 	Computer room 
• Security staff around the clock • Facility setback requirements • Barriers • Fencing	• Alarms • Security operations center • Seismic bracing • Security cameras	• Two-factor access control: biometric and card readers • Cameras • Days of backup power

Update management. Security update management helps protect systems from known vulnerabilities. Azure uses integrated deployment systems to manage the distribution and installation of security updates for Microsoft software. Azure uses a combination of Microsoft and third-party scanning tools to run OS, web application, and database scans of the Azure environment.

Antivirus and antimalware. Azure software components must go through a virus scan prior to deployment. Code is not moved to production without a clean and successful virus scan. In addition, Microsoft provides native antimalware on all Azure VMs. Microsoft recommends that customers run some form of antimalware or antivirus on all virtual machines (VMs). Customers can install Microsoft Antimalware for Cloud Services and Virtual Machines or another antivirus solution on VMs, and VMs can be routinely reimaged to clean out intrusions that may have gone undetected.

Penetration testing. Microsoft conducts regular penetration testing to improve Azure security controls and processes. Microsoft understands that security assessment is also an important part of our customers' application development and deployment. Therefore, Microsoft has established a policy for customers to carry out authorized penetration testing on their own—and only their own—applications hosted in Azure.

DDoS Protection. Azure has a defense system against Distributed Denial-of-Service (DDoS) attacks on Azure platform services. It uses standard detection and mitigation techniques. Azure's DDoS defense system is designed to withstand attacks generated from outside and inside the platform.

Network Protection

Azure networking provides the infrastructure necessary to securely connect VMs to one another and to connect on-premises data centers with Azure VMs. Because Azure's shared infrastructure hosts hundreds of millions of active VMs, protecting the security and confidentiality of network traffic is critical.

In the traditional datacenter model, a company's IT organization controls networked systems, including physical access to networking equipment. In the cloud service model, the responsibilities for network protection and management are shared between the cloud provider and the customer. Customers do not have physical access, but they implement the logical equivalent within their cloud environment through tools such as Guest operating system (OS) firewalls, Virtual Network Gateway configuration, and Virtual Private Networks.

Network isolation. Azure is a multitenant service, meaning that multiple customers' deployments and VMs are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's data from that of others. This provides the scale and economic benefits of multitenant services while rigorously preventing customers from accessing one another's data.

Virtual networks. A customer can assign multiple deployments within a subscription to a virtual network and allow those deployments to communicate with each other through private IP addresses. Each virtual network is isolated from other virtual networks.

VPN and Express Route. Microsoft enables connections from customer sites and remote workers to Azure Virtual Networks using Site-to-Site and Point-to-Site VPNs. For even better performance, customers can use an optional ExpressRoute, a private fiber link into Azure data centers that keeps their traffic off the Internet.

Encrypting communications. Built-in cryptographic technology enables customers to encrypt communications within and between deployments, between Azure regions, and from Azure to on-premises data centers.

"If you're resisting the cloud because of security concerns, you're running out of excuses."

FORRESTER

Data Protection

Azure allows customers to encrypt data and manage keys, and safeguards customer data for applications, platform, system and storage using three specific methods: encryption, segregation, and destruction.

Data isolation. Azure is a multitenant service, meaning that multiple customers' deployments and virtual machines are stored on the same physical hardware.

Protecting data at rest. Azure offers a wide range of encryption capabilities, giving customers the flexibility to choose the solution that best meets their needs. Azure Key Vault helps customers easily and cost effectively streamline key management and maintain control of keys used by cloud applications and services to encrypt data.

Protecting data in transit. For data in transit, customers can enable encryption for traffic between their own VMs and end users. Azure protects data in transit, such as between two virtual networks. Azure uses industry standard transport protocols such as TLS between devices and Microsoft datacenters, and within datacenters themselves.

Encryption. Customers can encrypt data in storage and in transit to align with best practices for protecting confidentiality and data integrity. For data in transit, Azure uses industry-standard transport protocols between devices and Microsoft datacenters and within datacenters themselves. You can enable encryption for traffic between your own virtual machines and end users.

Data redundancy. Customers may opt for in-country storage for compliance or latency considerations or out-of-country storage for security or disaster recovery purposes. Data may be replicated within a selected geographic area for redundancy.

Data destruction. When customers delete data or leave Azure, Microsoft follows strict standards for overwriting storage resources before reuse. As part of our agreements for cloud services such as Azure Storage, Azure VMs, and Azure Active Directory, we contractually commit to specific processes for the deletion of data.

“From a security point of view, I think Azure is a demonstrably more secure environment than most banks’ datacenters.”


John Schlesinger,
Chief Enterprise Architect, Temenos (Switzerland)

Identity and Access

Microsoft has strict controls that restrict access to Azure by Microsoft employees. Azure also enables customers to control access to their environments, data and applications.

Enterprise cloud directory. Azure Active Directory is a comprehensive identity and access management solution in the cloud. It combines core directory services, advanced identity governance, security, and application access management. Azure Active Directory makes it easy for developers to build policy-based identity management into their applications. Azure Active Directory Premium includes additional features to meet the advanced identity and access needs of enterprise organizations. Azure Active Directory enables a single identity management capability across on-premises, cloud, and mobile solutions.

Active Directory



Multi-Factor Authentication. Microsoft Azure provides Multi-Factor Authentication (MFA). This helps safeguard access to data and applications and enables regulatory compliance while meeting user demand for a simple sign-in process for both on-premises and cloud applications. It delivers strong authentication via a range of easy verification options—phone call, text message, or mobile app notification—allowing users to choose the method they prefer.

Access monitoring and logging. Security reports are used to monitor access patterns and to proactively identify and mitigate potential threats. Microsoft administrative operations, including system access, are logged to provide an audit trail if unauthorized or accidental changes are made. Customers can turn on additional access monitoring functionality in Azure and use third-party monitoring tools to detect additional threats. Customers can request reports from Microsoft that provide information about user access to their environments.

Privacy: Customers own and control their data

Customers will only use cloud providers in which they have great trust. They must trust that the privacy of their information will be protected, and that their data will be used in a way that is consistent with their expectations.

We build privacy protections into Azure through Privacy by Design, a program which describes how we build and operate products and services to protect privacy. Standards and processes that support Privacy by Design principles include the Microsoft Online Services Privacy Statement (which details Microsoft's core privacy requirements and practices) and the Microsoft Secure Development Lifecycle (which includes addressing privacy requirements).

“The question is no longer: ‘How do I move to the cloud?’ Instead, it’s ‘Now that I’m in the cloud, how do I make sure I’ve optimized my investment and risk exposure?’”



We then back those protections with strong contractual commitments to safeguard customer data, including offering EU Model Clauses (which provides terms covering the processing of personal information), and complying with international standards.

Microsoft uses customer data stored in Azure only to provide the service, including purposes compatible with providing the service. Azure does not use customer data for advertising or similar commercial purposes.

Contractual commitments. Microsoft was the first major cloud service provider to make contractual privacy commitments that help assure the privacy protections built into in-scope Azure services are strong. Among the many commitments that Microsoft supports are:

- **EU Model Clauses.** EU data protection law regulates the transfer of EU customer personal data to countries outside the European Economic Area (EEA). Microsoft offers customers the EU Standard Contractual Clauses that provide specific contractual guarantees around transfers of personal data for in-scope services. Europe’s privacy regulators have determined that the contractual privacy protections Azure delivers to its enterprise cloud customers meet current EU standards for international transfers of data. Microsoft is the first cloud provider to receive this recognition.
- **US-EU Safe Harbor Framework and the US-Swiss Safe Harbor Program.** Microsoft abides by these frameworks set forth by the US Department of Commerce regarding the collection, use, and retention of data from the EEA and Switzerland.
- **ISO/IEC 27018.** Microsoft is the first major cloud provider to adopt the first international code of practice for cloud privacy. ISO/IEC 27018 was developed to establish a uniform, international approach to protecting the privacy of personal data stored in the cloud. The British Standards Institution independently verified that Microsoft Azure is aligned with the guideline’s code of practice. ISO 27018 controls include a prohibition on the use of customer data for advertising and marketing purposes without the customer’s express consent

Restricted access by Microsoft personnel. Access to customer data by Microsoft personnel is restricted. Customer data is only accessed when necessary to support the customer’s use of Azure. This may include troubleshooting aimed at preventing, detecting, or repairing problems affecting the operation of Azure and the improvement of features that involve the detection of, and protection against, emerging and evolving threats to the user (such as malware or spam). When granted, access is controlled and logged. Strong authentication, including the use of multi-factor authentication, helps limit access to authorized personnel only. Access is revoked as soon as it is no longer needed.

Notification of lawful requests for information. Microsoft believes that customers should control their data whether stored on their premises or in a cloud service. We will not disclose Azure customer data to law enforcement except as a customer directs or where required by law. When governments make a lawful demand for Azure customer data from Microsoft, we strive to be principled, limited in what we disclose, and committed to transparency.

- Microsoft does not provide any third party with direct or unfettered access to customer data. Microsoft only releases specific data mandated by the relevant legal demand.
- If a government wants customer data—including for national security purposes—it needs to follow the applicable legal process, meaning it must serve us with a warrant or court order for content or subpoena for account information. If compelled to disclose customer data, we will promptly notify the customer and provide a copy of the demand unless legally prohibited from doing so.

- Microsoft will only respond to requests for specific accounts and identifiers. There is no blanket or indiscriminate access to Microsoft's customer data. Every request is explicitly reviewed by Microsoft's legal team, who ensures that the requests are valid, rejects those that are not, and makes sure we only provide the data specified in the order.

In its commitment to transparency, Microsoft regularly publishes a Law Enforcement Requests Report that discloses the scope and number of requests we receive.

Greater transparency and simplicity of data use policies.

Microsoft keeps customers informed about the processes to protect data privacy and security, including practices and policies. Microsoft also provides the summaries of independent audits of services, which helps customers pursue their own compliance.

“Just as computer users back up their laptops in case they break or are lost, Estonia is working out how to back up the country, in case it is attacked by Russia.”

The Economist,
reporting on Estonia's Azure cloud backup

Customers are in control of their data

For many organizations, the benefits of moving to the cloud are clear. Still, fear of losing control causes their decision makers to hesitate. Where will data be stored? Who owns the organization's data? Who will be accessing the data? And what happens if the organization wants to switch providers? These are all valid questions—questions Microsoft has in mind when making a clear commitment to provide customers with control over their data. This commitment is unique among major cloud service providers.

Customers own their data. This belief is fundamental to the Microsoft approach. When a customer utilizes Azure, they retain exclusive ownership of their data. Microsoft takes steps to protect many types of data.

Microsoft defines customer data as “all data, including all text, sound, video or image files, and software that are provided to Microsoft by, or on behalf of, Customer through use of the Online Service.” For example, this includes data that you upload for storage or processing and applications that you run in Azure.

Customers can access their own customer data at any at any time and for any reason without assistance from Microsoft. Microsoft will not use customer data or derive information from it for advertising. We will use customer data only to provide the service or for purposes compatible with providing the service.

-
- **Customer data** is all data, including all text, sound, video or image files, and software that are provided to Microsoft by or on behalf of the customer through use of Azure. For example, it includes data uploaded for storage or processing and applications uploaded by the customer for hosting on Azure.
 - **Administrator data** is the information about administrators (including account contact and subscription administrators) supplied during signup, purchase, or administration of Azure, such as name, phone number, and email address.
 - **Metadata** includes configuration and technical settings and information. For example, it includes the disk configuration settings for an Azure virtual machine or the database design for an SQL Database. Metadata does not include information from which customer data could be derived.
 - **Access control data** is data that is used to manage access to other types of data or functions within Azure. It includes passwords, security certificates, and other authentication-related data.
-

“Our brand rests on the continuity of our IT systems, which are now more available running in Azure.”

Andrew Goodin,
Global Manager of Information Systems
Zespri International (New Zealand)

Control over data location. When customers entrust their data to Microsoft, they are not giving up control. For many customers, knowing and controlling the location of their data can be an important element of data privacy, compliance, and governance. Microsoft Azure offers an ever-expanding network of data centers across the globe. Most Azure services permit customers to specify the particular geography where their customer data will be stored. Data may be replicated within a selected geographic area for redundancy, but will not be replicated outside it for redundancy.

Encryption key management. To ensure control over encrypted data, customers have the option to generate and manage their own encryption keys, and determine who is authorized to use them. They also have the option to revoke Microsoft's copy of their encryption key, although this may limit Microsoft's ability to troubleshoot or repair problems and security threats.

Role based access control. Microsoft provides an approach allowing customers to restrict system access to authorized users based on role assignment, role authorization, and permission authorization. Tools in multiple Microsoft cloud services support authorization based on a user's role, simplifying access control across defined groups of users.

Control over data destruction. When customers delete data or leave a Microsoft cloud service, Microsoft follows strict standards for overwriting storage resources before reuse, as well physical destruction of decommissioned hardware, including contractual commitments to specific processes for the deletion of data and the destruction of storage hardware.

Transparency

For customers to effectively exercise their right to control their data, they must have access and visibility to that data. They must know where it is stored. They must also know, through clearly stated and readily available policies and procedures, how the cloud provider helps secure customer data, who can access it, and under what circumstances.

Where and how data is stored and used. Microsoft gives Azure customers visibility to where their customer data is stored in an ever-expanding network of datacenters around the globe. Customers can balance the need to store backups at multiple locations in case of a disaster with the need to keep their data out of certain geographies. Microsoft provides clear data maps and geographic boundary information for all datacenters.

How data is secured. Customers have access to up-to-date information regarding security policies and procedures. Microsoft promotes transparency by publishing and adhering to the Security Development Lifecycle.

Who requests access to customer data. Microsoft will never disclose Azure customer data to a government or law enforcement agency except as directed by the customer or where required by law. In response to lawful demands for Azure customer data, Microsoft strives to be principled, limited in disclosure, and committed to transparency. Microsoft regularly publishes a Law Enforcement Requests Report that discloses the scope and number of government requests received.

Breach notification. In the event that customer data is compromised, Microsoft will notify its customers. Azure has comprehensive, transparent policies that govern incident response from identification all the way through to lessons learned.

Audit standards certifications. Rigorous third-party audits, such as those conducted by the British Standards Institute, verify Azure's adherence to the strict security controls these standards mandate. As part of Microsoft's commitment to transparency, customers can verify Azure's implementation of many security controls by requesting audit results from the certifying third parties.

“By 2020 clouds will stop being referred to as ‘public’ and ‘private’. It will simply be the way business is done and IT is provisioned.”



Customer guidance. Microsoft publishes a Security Response Center Progress Report and a Security Intelligence Report to provide customers with insights into the threat landscape, and provide prescriptive guidance for managing risk to protect their assets.

Transparency Centers. Microsoft operates Transparency Centers that provide government customers with the ability to review source code, reassure themselves of its integrity, and confirm there are no back doors.

Compliance: Azure conforms to global standards

Microsoft invests heavily in the development of robust and innovative compliance processes. The Microsoft compliance framework for online services maps controls to multiple regulatory standards. This enables Microsoft to design and build services using a common set of controls, streamlining compliance across a range of regulations today and as they evolve in the future.

Microsoft compliance processes also make it easier for customers to achieve compliance across multiple services and meet their changing needs efficiently. Together, security-enhanced technology and effective compliance processes enable Microsoft to maintain and expand a rich set of third-party certifications. These help customers demonstrate compliance readiness to their customers, auditors, and regulators. As part of its commitment to transparency, Microsoft shares third-party verification results with its customers.

Certifications and attestations. Azure meets a broad set of international as well as regional and industry-specific compliance standards, such as ISO 27001, FedRAMP, SOC 1 and SOC 2. Azure's adherence to the strict security controls contained in these standards is verified by rigorous third-party audits that demonstrate Azure services work with and meet world-class industry standards, certifications, attestations, and authorizations.

Comprehensive, independently verified compliance. Azure is designed with a compliance strategy that helps customers address business objectives and industry standards and regulations. The security compliance framework includes test and audit phases, security analytics, risk management best practices, and security benchmark analysis to achieve certificates and attestations. Microsoft Azure offers the following certifications for all in-scope services.

CDSA. The Content Delivery and Security Association (CDSA) provides a Content Protection and Security (CPS) standard for compliance with anti-piracy procedures governing digital media. Azure passed the CDSA audit, enabling secure workflows for content development and distribution.

CJIS. Any US state or local agency that wants to access the FBI's Criminal Justice Information Services (CJIS) database is required to adhere to the CJIS Security Policy. Azure is the only major cloud provider that contractually commits to conformance with the CJIS Security Policy, which commits Microsoft to adhere to the same requirements that law enforcement and public safety entities must meet.

CSA CCM. The Cloud Security Alliance (CSA) is a nonprofit, member-driven organization with a mission to promote the use of best practices for providing security assurance within the cloud. The CSA Cloud Controls Matrix (CCM) provides detailed information about how Azure fulfills the security, privacy, compliance, and risk management requirements defined in the CCM version 1.2, and is published in the CSA's Security Trust and Assurance Registry (STAR).

EU Model Clauses. Microsoft offers customers EU Standard Contractual Clauses that provide contractual guarantees around transfers of personal data outside of the EU. Microsoft is the first company to receive joint approval from the EU's Article 29 Working Party that the contractual privacy protections Azure delivers to its enterprise cloud customers meet current EU standards for international transfers of data. This ensures that Azure customers can use Microsoft services to move data freely through our cloud from Europe to the rest of the world.

FDA 21 CFR Part 11. The US Food and Drug Administration (FDA) Code of Federal Regulations (CFR) Title 21 Part 11 lists requirements for the security of electronic records of companies that sell food and drugs manufactured or consumed in the United States. The compliance reports produced by Azure's independent third party SSAE and ISO auditors identify the procedural and technical controls established at Microsoft and can be used to satisfy the requirements of CFR Title 21 Part 11. Microsoft is able to show how relevant controls within these reports have an impact on compliance with the FDA 21 CFR 11 regulations.

FedRAMP. Azure has been granted a Provisional Authority to Operate (P-ATO) from the Federal Risk and Authorization Management Program (FedRAMP) Joint Authorization Board (JAB) at a Moderate impact level based upon the FIPS 199 classification. FedRAMP is a US government program that provides a standard approach to security assessment, authorization, and monitoring for cloud services used by federal agencies and thereby saves the taxpayer and individual organizations the time and cost of conducting their own independent reviews.

FERPA. The Family Educational Rights and Privacy Act (FERPA) is a US federal law that protects the privacy of student educational records. Microsoft agrees to use and disclosure restrictions imposed by FERPA.

FIPS 140-2. Azure complies with the Federal Information Processing Standard (FIPS) Publication 140-2, a US government standard that defines a minimum set of security requirements for products and systems that implement cryptography.

HIPAA. The Health Insurance Portability and Accountability Act (HIPAA) is a US federal law that regulates patient Protected Health Information (PHI). Azure offers customers a HIPAA Business Associate Agreement (BAA), stipulating adherence to certain security and privacy provisions in HIPAA and the HITECH Act. To assist customers in their individual compliance efforts, Microsoft offers a BAA to Azure customers as a contract addendum.

IRAP. Azure has been assessed against the Australian Government Information Security Registered Assessors Program (IRAP), which provides assurance for public sector customers that Microsoft has appropriate and effective security controls.

ISO/IEC 27018. Microsoft is the first cloud provider to have adopted the ISO/IEC 27018 code of practice, covering the processing of personal information by cloud service providers.

ISO/IEC 27001/27002:2013. Azure complies with this standard, which defines the security controls required of an information security management system.

MLPS. Multi-Level Protection Scheme (MLPS) is based on the Chinese state standard issued by the Ministry of Public Security. Azure operated by 21Vianet adheres to this standard, which provides assurance for both the management and technical security of cloud systems.

MTCS. Azure has achieved Level-1 certification with the Multi-Tier Cloud Security Standard for Singapore (MTCS SS), a cloud security standard covering areas such as data security, confidentiality, business impact, and operational transparency, developed under the Singapore Information Technology Standards Committee.

PCI DSS. Azure is Level 1 compliant with Payment Card Industry (PCI) Data Security Standards (DSS) version 3.0, the global certification standard for organizations that accept most payments cards, as well store, process, or transmit cardholder data.

SOC 1 and SOC 2. Azure has been audited against the Service Organization Control (SOC) reporting framework for both SOC 1 Type 2 and SOC 2 Type 2. Both reports are available to customers to meet a wide range of US and international auditing requirements.

The SOC 1 Type 2 audit report attests to the design and operating effectiveness of Azure controls. The SOC 2 Type 2 audit included a further examination of Azure controls related to security, availability, and confidentiality. Azure is audited annually to ensure that security controls are maintained.

TCS CCCPPF. Azure operated by 21Vianet is among the first cloud providers in China to pass the Trusted Cloud Service certification developed by the China Cloud Computing Promotion and Policy Forum (CCCPF).

UK G-Cloud. The UK Government G-Cloud is a cloud computing certification for services used by government entities in the United Kingdom. Azure has received OFFICIAL accreditation from the UK Government Pan Government Accreditor.

Additional resources

Azure Trust Center

<http://azure.microsoft.com/trustcenter>

Cloud Security Alliance Cloud Controls Matrix

<https://cloudsecurityalliance.org/research/ccm/>

Microsoft Cloud Security Readiness Tool


<http://www.microsoft.com/trustedcloud>

Microsoft Online Services Privacy Statement

<http://aka.ms/oneservices-privacy>

Microsoft Privacy Practices


<http://aka.ms/privacy-practices>





NOTE: Certain recommendations contained herein may result in increased data, network, or compute resource usage, and increase your license or subscription costs.

© 2015 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

The background is a solid blue color. Overlaid on this is a grid of thin, light green lines. At the intersections of these lines, there are small circles. Some circles are solid green, while others are hollow with a yellow-green outline. The circles are scattered across the page, with a higher concentration in the upper half. The text is centered in the middle of the page.

Google for Work Security and Compliance Whitepaper

How Google protects your data.

Google™ for Work

This whitepaper applies to the
following Google Apps products

*Google Apps for Work, Education, Government,
Nonprofit, Drive for Work, and Google Apps Unlimited*

Table of Contents



Introduction 1

Google Has a Strong Security Culture 2

- Employee background checks
- Security training for all employees
- Internal security and privacy events
- Our dedicated security team
- Our dedicated privacy team
- Internal audit and compliance specialists
- Collaboration with the security research community

Operational Security 4

- Vulnerability management
- Malware prevention
- Monitoring
- Incident management

Technology with Security at Its Core 6

- State-of-the-art data centers
 - Powering our data centers
 - Environmental impact
- Custom server hardware and software
- Hardware tracking and disposal
- A global network with unique security benefits
- Securing data in transit
- Low latency and highly available solution
- Service availability

Independent Third-Party Certifications 10

- ISO 27001
- SOC 2/3
- FISMA

Data Usage 11

- Our philosophy
- No advertising in Google Apps

Data Access and Restrictions 12

- Administrative access
- For customer administrators
- Law enforcement data requests
- Third-party suppliers

Regulatory compliance 14

- Data processing amendment
- EU Data Protection Directive
 - U.S.-EU and U.S.-Swiss Safe Harbor Frameworks
 - EU model contract clauses
- U.S. Health Insurance Portability and Accountability Act (HIPAA)
- U.S. Family Educational Rights and Privacy Act (FERPA)
- Children's Online Privacy Protection Act of 1998 (COPPA)

Empowering Users and Administrators to Improve Security and Compliance 16

User authentication/authorization features

- 2-step verification
- Security Key
- Single sign-on (SAML 2.0)
- OAuth 2.0 and OpenID Connect

Data management features

- Information Rights Management (IRM)
- Drive audit log
- Drive content compliance / alerting
- Trusted domains for Drive sharing

Email security features

- Secure transport (TLS) enforcement
- Phishing prevention
- Email content compliance
- Objectionable content
- Restricted email delivery

eDiscovery features

- Email retention policy
- Legal holds
- Search/discovery
- Evidence export
- Support for third-party email platforms

Securing endpoints

- Mobile device management (MDM)
- Policy-based Chrome browser security
- Chrome device management

Data recovery

- Restore a recently deleted user
- Restore a user's Drive or Gmail data

Security reports

Conclusion 23



Introduction

Cloud computing offers many advantages and conveniences for today's organizations. Employees can work together in documents in real time from their phone or tablet from any location, and communicate instantly with teammates via video, voice, instant message, or email. No longer tied to a single machine, they have the freedom to work together from anywhere, using any device they choose. Meanwhile, their employers don't shoulder the cost or burden of maintaining servers and constantly updating software. It's no surprise, then, that so many organizations around the world are storing their information and getting work done in the cloud.



The growth of the cloud has thrust the issue of security and trust into the spotlight. That's because cloud services operate very differently from traditional on-premises technology. Rather than residing on local servers, content is now managed on Google servers that are part of our global data center network. In the past, organizations felt that they had complete control over how infrastructure was run and who operated it. Organizations moving to the cloud will rely on cloud suppliers to manage the infrastructure, operations, and delivery of services. In this new world, companies will still control company data, but via cloud-based tools and dashboards. Rather than only using desktop computers, users can now access work files on their personal mobile devices. Customers must assess whether the security controls and compliance of any cloud solution meet their individual requirements. Customers must therefore understand how these solutions protect and process their data. The goal of this whitepaper is to provide an introduction to Google's technology in the context of security and compliance.

As a cloud pioneer, Google fully understands the security implications of the cloud model. Our cloud services are designed to deliver better security than many traditional on-premises solutions. We make security a priority to protect our own operations, but because Google runs on the same infrastructure that we make available to our customers, your organization can directly benefit from these protections. That's why we focus on security, and protection of data is among our primary design criteria. Security drives our organizational structure, training priorities and hiring processes. It shapes our data centers and the technology they house. It's central to our everyday operations and disaster planning, including how we address threats. It's prioritized in the way we handle customer data. And it's the cornerstone of our account controls, our compliance audits and the certifications we offer our customers.

This paper outlines Google's approach to security and compliance for Google Apps, our cloud-based productivity suite. Used by more than five million organizations worldwide, from large banks and retailers with hundreds of thousands of people to fast-growing startups, Google Apps for Work and Education includes Gmail, Calendar, Groups, Drive, Docs, Sheets, Slides, Hangouts, Sites, Talk, Contacts and Vault. Google Apps is designed to help teams work together in new, more efficient ways, no matter where members are located or what device they happen to be using.

This whitepaper will be divided into two main sections: security and compliance. The security section will include details on organizational and technical controls regarding how Google protects your data. The second section on compliance will cover how your data is processed and details on how organizations can meet regulatory requirements.

Google Has a Strong Security Culture

Google has created a vibrant and inclusive security culture for all employees. The influence of this culture is apparent during the hiring process, employee onboarding, as part of ongoing training and in company-wide events to raise awareness.

Employee background checks

Before they join our staff, Google will verify an individual's education and previous employment, and perform internal and external reference checks. Where local labor law or statutory regulations permit, Google may also conduct criminal, credit, immigration, and security checks. The extent of these background checks is dependent on the desired position.

Security training for all employees

All Google employees undergo security training as part of the orientation process and receive ongoing security training throughout their Google careers. During orientation, new employees agree to our [Code of Conduct](#), which highlights our commitment to keep customer information safe and secure. Depending on their job role, additional training on specific aspects of security may be required. For instance, the information security team instructs new engineers on topics like secure coding practices, product design and automated vulnerability testing tools. Engineers also attend technical presentations on security-related topics and receive a security newsletter that covers new threats, attack patterns, mitigation techniques and more.

Internal security and privacy events

Google hosts regular internal conferences to raise awareness and drive innovation in security and data privacy, which are open to all employees. Security and privacy is an ever-evolving area, and Google recognizes that dedicated employee engagement is a key means of raising awareness. One example is "Privacy Week," during which Google hosts events across global offices to raise awareness of privacy in all facets, from software development, data handling and policy enforcement to living our [privacy principles](#). Google also hosts regular "Tech Talks" focusing on subjects that often include security and privacy.



Google employs more than 500 full-time security and privacy professionals, who are part of our software engineering and operations division. Our team includes some of the world's foremost experts in information, application and network security.

Our dedicated security team

Google employs more than 500 full-time security and privacy professionals, who are part of our software engineering and operations division. Our team includes some of the world's foremost experts in information, application and network security. This team is tasked with maintaining the company's defense systems, developing security review processes, building security infrastructure and implementing Google's security policies. Google's dedicated security team actively scans for security threats using commercial and custom tools, penetration tests, quality assurance (QA) measures and software security reviews.

Within Google, members of the information security team review security plans for all networks, systems and services. They provide project-specific consulting services to Google's product and engineering teams. They monitor for suspicious activity on Google's networks, address information security threats, perform routine security evaluations and audits, and engage outside experts to conduct regular security assessments. We specifically built a full-time team, known as [Project Zero](#), that aims to prevent targeted attacks by reporting bugs to software vendors and filing them in an external database.

The security team also takes part in research and outreach activities to protect the wider community of Internet users, beyond just those who choose Google solutions. Some examples of this research would be the discovery of the [POODLE SSL 3.0 exploit](#) and [cipher suite weaknesses](#). The security team also publishes security research papers, [available to the public](#). The security team also organizes and participates in [open-source projects](#) and academic conferences.

Our dedicated privacy team

The Google Privacy team operates independently from product development and security organizations, but participates in every Google product launch. The team reviews design documentation and code audits to ensure that privacy requirements are followed. The Privacy team has built a set of automated monitoring tools to help ensure that products with Customer Data operate as designed and in accordance with our privacy policy. They help release products that reflect strong privacy standards: transparent collection of user data and providing users and administrators with meaningful privacy configuration options, while continuing to be good stewards of any information stored on our platform. After products launch, the privacy team oversees automated processes that audit data traffic to verify appropriate data usage. In addition, the privacy team conducts research providing thought leadership on privacy best practices for our emerging technologies.

Internal audit and compliance specialists

Google has a dedicated internal audit team that reviews compliance with security laws and regulations around the world. As new auditing standards are created, the internal audit team determines what controls, processes, and systems are needed to meet them. This team facilitates and supports independent audits and assessments by third parties.

Collaboration with the security research community

Google has long enjoyed a close relationship with the security research community, and we greatly value their help identifying vulnerabilities in Google Apps and other Google products. Our [Vulnerability Reward Program](#) encourages researchers to report design and implementation issues that may put customer data at risk, offering rewards in the tens of thousands of dollars. In Chrome, for instance, we warn users against malware and phishing, and offer rewards for finding security bugs. Due to our collaboration with the research community, we've squashed more than 700 Chrome security bugs and have rewarded more than \$1.25 million — more than \$2 million has been awarded across Google's various vulnerability rewards programs. We publicly [thank these individuals](#) and list them as contributors to our products and services.

Operational Security

Far from being an afterthought or the focus of occasional initiatives, security is an integral part of our operations.

Vulnerability management

Google administrates a vulnerability management process that actively scans for security threats using a combination of commercially available and purpose-built in-house tools, intensive automated and manual penetration efforts, quality assurance processes, software security reviews and external audits. The vulnerability management team is responsible for tracking and following up on vulnerabilities. Once a vulnerability requiring remediation has been identified, it is logged, prioritized according to severity, and assigned an owner. The vulnerability management team tracks such issues and follows up frequently until they can verify that the issues have been remediated.

Google also maintains relationships and interfaces with members of the security research community to track reported issues in Google services and open-source tools. More information about reporting security issues can be found at [Google Application Security](#).



Malware prevention

An effective malware attack can lead to account compromise, data theft, and possibly additional access to a network. Google takes these threats to its networks and its customers very seriously and uses a variety of methods to prevent, detect and eradicate malware. Google helps tens of millions of people every day to protect themselves from harm by showing warnings to users of Google Chrome, Mozilla Firefox and Apple Safari when they attempt to navigate to websites that would steal their personal information or install software designed to take over their computers. Malware sites or email attachments install malicious software on users' machines to steal private information, perform identity theft, or attack other computers. When people visit these sites, software that takes over their computer is downloaded without their knowledge. Google's malware strategy begins with infection prevention by using manual and automated scanners to scour Google's search index for websites that may be vehicles for malware or phishing. Approximately one billion people use [Google's Safe Browsing](#) on a regular basis. Google's Safe Browsing technology examines billions of URLs per day looking for unsafe websites. Every day, we discover thousands of new unsafe sites, many of which are legitimate websites that have been compromised. When we detect unsafe sites, we show warnings on Google Search and in web browsers. In addition to our Safe Browsing solution, Google operates [VirusTotal](#), a free online service that analyzes files and URLs enabling the identification of viruses, worms, trojans and other kinds of malicious content detected by antivirus engines and website scanners. VirusTotal's mission is to help in improving the antivirus and security industry and make the Internet a safer place through the development of free tools and services.

Google makes use of multiple antivirus engines in Gmail, Drive, servers and workstations to help identify malware that may be missed by antivirus signatures.

Monitoring

Google's security monitoring program is focused on information gathered from internal network traffic, employee actions on systems and outside knowledge of vulnerabilities. At many points across our global network, internal traffic is inspected for suspicious behavior, such as the presence of traffic that might indicate botnet connections. This analysis is performed using a combination of open-source and commercial tools for traffic capture and parsing. A proprietary correlation system built on top of Google technology also supports this analysis. Network analysis is supplemented by examining system logs to identify unusual behavior, such as attempted access of customer data. Google security engineers place standing search alerts on public data repositories to look for security incidents that might affect the company's infrastructure. They actively review inbound security reports and monitor public mailing lists, blog posts, and wikis. Automated network analysis helps determine when an unknown threat may exist and escalates to Google security staff, and network analysis is supplemented by automated analysis of system logs.

Google helps tens of millions of people every day to protect themselves from harm by showing warnings to users of Google Chrome, Mozilla Firefox and Apple Safari when they attempt to navigate to websites that would steal their personal information or install software designed to take over their computers.



Incident management

We have a rigorous incident management process for security events that may affect the confidentiality, integrity, or availability of systems or data. If an incident occurs, the security team logs and prioritizes it according to its severity. Events that directly impact customers are assigned the highest priority. This process specifies courses of action, procedures for notification, escalation, mitigation, and documentation. Google's security incident management program is structured around the NIST guidance on handling incidents (NIST SP 800–61). Key staff are trained in forensics and handling evidence in preparation for an event, including the use of third-party and proprietary tools. Testing of incident response plans is performed for key areas, such as systems that store sensitive customer information. These tests take into consideration a variety of scenarios, including insider threats and software vulnerabilities. To help ensure the swift resolution of security incidents, the Google security team is available 24/7 to all employees. If an incident involves customer data, Google or its partners will inform the customer and support investigative efforts via our support team.

Technology with Security at Its Core

Google Apps runs on a technology platform that is conceived, designed and built to operate securely. Google is an innovator in hardware, software, network and system management technologies. We custom-designed our servers, proprietary operating system, and geographically distributed data centers. Using the principles of “defense in depth,” we’ve created an IT infrastructure that is more secure and easier to manage than more traditional technologies.

State-of-the-art data centers

Google's focus on security and protection of data is among [our primary design criteria](#). Google data center physical security features a layered security model, including safeguards like custom-designed electronic access cards, alarms, vehicle access barriers, perimeter fencing, metal detectors, and biometrics, and the data center floor features laser beam intrusion detection. Our data centers are monitored 24/7 by high-resolution interior and exterior cameras that can detect and track intruders. Access logs,

activity records, and camera footage are available in case an incident occurs. Data centers are also routinely patrolled by experienced security guards who have undergone rigorous background checks and training. As you get closer to the data center floor, security measures also increase. Access to the data center floor is only possible via a security corridor which implements multifactor access control using security badges and biometrics. Only approved employees with specific roles may enter. Less than one percent of Googlers will ever step foot in one of our data centers.

Powering our data centers

To keep things running 24/7 and ensure uninterrupted services, Google's data centers feature redundant power systems and environmental controls. Every critical component has a primary and alternate power source, each with equal power. Diesel engine backup generators can provide enough emergency electrical power to run each data center at full capacity. Cooling systems maintain a constant operating temperature for servers and other hardware, reducing the risk of service outages. Fire detection and suppression equipment helps prevent damage to hardware. Heat, fire, and smoke detectors trigger audible and visible alarms in the affected zone, at security operations consoles, and at remote monitoring desks.

Environmental impact

Google reduces environmental impact of running our data centers by designing and building our own facilities. We install smart temperature controls, use “free-cooling” techniques like using outside air or reused water for cooling, and redesign how power is distributed to reduce unnecessary energy loss. To gauge improvements, we calculate the performance of each facility using comprehensive efficiency measurements. We're the first major Internet services company to gain external certification of our high environmental, workplace safety and energy management standards throughout our data centers. Specifically, we received voluntary [ISO 14001](#), [OHSAS 18001](#) and [ISO 50001](#) certifications. In a nutshell, these standards are built around a very simple concept: Say what you're going to do, then do what you say—and then keep improving.

Custom server hardware and software

Google's data centers house energy-efficient custom, purpose-built servers and network equipment that we design and manufacture ourselves. Unlike much commercially available hardware, Google servers don't include unnecessary components such as video cards, chipsets, or peripheral connectors, which can introduce vulnerabilities. Our production servers run a custom-designed operating system (OS) based on a stripped-down and hardened version of Linux. Google's servers and their OS are designed for the sole purpose of providing Google services. Server resources are dynamically allocated, allowing for flexibility in growth and the ability to adapt quickly and efficiently, adding or reallocating resources based on customer demand. This homogeneous environment is maintained by proprietary software that continually monitors systems for binary modifications. If a modification is found that differs from the standard Google image, the system is automatically returned to its official state. These automated, self-healing mechanisms are designed to enable Google to monitor and remediate destabilizing events, receive notifications about incidents, and slow down potential compromise on the network.

Hardware tracking and disposal

Google meticulously tracks the location and status of all equipment within our data centers from acquisition to installation to retirement to destruction, via bar codes and asset tags. Metal detectors and video surveillance are implemented to help make sure no equipment leaves the data center floor without authorization. If a component fails to pass a performance test at any point during its lifecycle, it is removed from inventory and retired. Google hard drives leverage technologies like FDE (full disk encryption) and drive locking, to protect data at rest. When a hard drive is retired, authorized individuals verify that the disk is erased by writing zeros to the drive and performing a multiple-step verification process to ensure the drive contains no data. If the drive cannot be erased for any reason, it is stored securely until it can be physically destroyed. Physical destruction of disks is a multistage process beginning with a crusher that deforms the drive, followed by a shredder that breaks the drive into small pieces, which are then recycled at a secure facility. Each data center adheres to a strict disposal policy and any variances are immediately addressed.

A global network with unique security benefits

Google's IP data network consists of our own fiber, public fiber, and undersea cables. This allows us to deliver highly available and low latency services across the globe.

In other cloud services and on-premises solutions, customer data must make several journeys between devices, known as "hops," across the public Internet. The number of hops depends on the distance between the customer's ISP and the solution's data center. Each additional hop introduces a new opportunity for data to be attacked or intercepted. Because it's linked to most ISPs in the world, Google's global network improves the security of data in transit by limiting hops across the public Internet.

Defense in depth describes the multiple layers of defense that protect Google's network from external attacks. Only authorized services and protocols that meet our security requirements are allowed to traverse it; anything else is automatically dropped. Industry-standard firewalls and access control lists (ACLs) are used to enforce network segregation. All traffic is routed through custom GFE (Google Front End) servers to detect and stop malicious requests and Distributed Denial of Service (DDoS) attacks. Additionally, GFE servers are only allowed to communicate with a controlled list of servers internally; this "default deny" configuration prevents GFE servers from accessing unintended resources. Logs are routinely examined to reveal any exploitation of programming errors. Access to networked devices is restricted to authorized personnel.

Google's IP data network consists of our own fiber, public fiber, and undersea cables. This allows us to deliver highly available and low latency services across the globe.

Securing data in transit

Data is most vulnerable to unauthorized access as it travels across the Internet or within networks. For this reason, securing data in transit is a high priority for Google. Data traveling between a customer's device and Google is encrypted using HTTPS/TLS (Transport Layer Security). In fact, Google was the first major cloud provider to enable HTTPS/TLS by default. When sending to or receiving email from a non-Google user, all links of the chain (device, browser, provider of the email service) have to be strong and work together to make encryption work. We believe this is so important that we report on the industry's adoption of TLS on our [safe email site](#).

Google has also upgraded all our RSA certificates to 2048-bit keys, making our encryption in transit for Google Apps and all other Google services even stronger.

Perfect forward secrecy (PFS) minimizes the impact of a compromised key, or a cryptographic breakthrough. It protects network data by using a short-term key that lasts only a couple of days and is only held in memory, rather than a key that's used for years and kept on durable storage. Google was the first major web player to enable perfect forward secrecy by default.

Google encrypts all Google Apps data as it moves between our data centers on our private network.

Low latency and highly available solution

Google designs the components of our platform to be highly redundant. This redundancy applies to our server design, how we store data, network and Internet connectivity, and the software services themselves. This "redundancy of everything" includes the handling of errors by design and creates a solution that is not dependant on a single server, data center, or network connection. Google's data centers are geographically distributed to minimize the effects of regional disruptions such as natural disasters and local outages. In the event of hardware, software, or network failure, data is automatically shifted from one facility to another so that Google Apps customers can continue working in most cases without interruption. Customers with global workforces can collaborate on documents, video conferencing and more without additional configuration or expense. Global teams share a highly performant and low latency experience as they work together on a single global network.

Google's highly redundant infrastructure also helps protect our customers from data loss. For Google Apps, our recovery point objective (RPO) target is zero, and our recovery time objective (RTO) design target is also zero. We aim to achieve these targets through live or synchronous replication:

Google's data centers are geographically distributed to minimize the effects of regional disruptions such as natural disasters and local outages.



actions you take in Google Apps Products are simultaneously replicated in two data centers at once, so that if one data center fails, we transfer your data over to the other one that's also been reflecting your actions. Customer data is divided into digital pieces with random file names. Neither their content nor their file names are stored in readily human-readable format, and stored customer data cannot be traced to a particular customer or application just by inspecting it in storage. Each piece is then replicated in near-real time over multiple disks, multiple servers, and multiple data centers to avoid a single point of failure. To further prepare for the worst, we conduct disaster recovery drills in which we assume that individual data centers—including our corporate headquarters—won't be available for 30 days. We regularly test our readiness for plausible scenarios as well as more imaginative crises, like alien and zombie invasions.

Our highly redundant design has allowed Google to achieve an uptime of 99.984% for Gmail for the last years with **no scheduled downtime**. Simply put, when Google needs to service or upgrade our platform, users do not experience downtime or maintenance windows.

Service availability

Some of Google's services may not be available in some jurisdictions. Often these interruptions are temporary due to network outages, but others are permanent due to government-mandated blocks. Google's Transparency Report also shows [recent and ongoing disruptions of traffic](#) to Google products. We provide this data to help the public analyze and understand the availability of online information.

Independent Third-Party Certifications

Google's customers and regulators expect independent verification of our security, privacy, and compliance controls. In order to provide this, we undergo several independent third-party audits on a regular basis. For each one, an independent auditor examines our data centers, infrastructure, and operations. Regular audits are conducted to certify our compliance with the auditing standards ISO 27001, SOC 2 and SOC 3, as well as with the U.S. Federal Information Security Modernization Act of 2014 (FISMA) for Google Apps for Government. When customers consider Google Apps, these certifications can help them confirm that the product suite meets their security, compliance and data processing needs.

ISO 27001

ISO 27001 is one of the most widely recognized and accepted independent security standards. Google has earned it for the systems, technology, processes, and data centers that run Google Apps. Our compliance with the international standard was certified by Ernst & Young CertifyPoint, an ISO certification body accredited by the Dutch Accreditation Council (a member of the International Accreditation Forum, or IAF). Our ISO 27001 certificate and scoping document are available in [our Trust Center](#).

SOC 2/3

In 2014, the American Institute of Certified Public Accountants (AICPA) Assurance Services Executive Committee (ASEC) released the revised version of the Trust Services Principles and Criteria (TSP). SOC (Service Organization Controls) is an audit framework for non-privacy principles that include security, availability, processing integrity, and confidentiality. Google has both SOC 2 and SOC 3 reports. Our SOC 3 report is available for [download](#) without a nondisclosure agreement. The SOC 3 confirms our compliance with the principles of security, availability, processing integrity and confidentiality.

FISMA

FISMA is a U.S. federal law pertaining to the information security of federal agencies' information systems. The law requires agencies to ensure that their systems and those operated by service providers such as Google meet minimum security requirements specified by the National Institute of Standards and Technology (NIST). Google, as a cloud provider, has one of the longest track records in meeting these requirements, and we maintain a current authorization to operate (ATO) for Google Apps for Government. The Federal Risk and Authorization Management Program (FedRAMP) adds a number of requirements on top of the FISMA requirements that cloud providers must meet. For more information on the FedRAMP program and vendor status, please visit [fedramp.gov](#).

Data Usage

Our philosophy

Google Apps customers own their data, not Google. The data that Google Apps organizations and users put into our systems is theirs, and we do not scan it for advertisements nor sell it to third parties. We offer our customers a detailed [data processing amendment](#) that describes our commitment to protecting customer data. It states that Google will not process data for any purpose other





than to fulfill our contractual obligations. Furthermore, if customers delete their data, we commit to deleting it from our systems within 180 days. Finally, we provide tools that make it easy for customer administrators to take their data with them if they choose to stop using our services, without penalty or additional cost imposed by Google.

No advertising in Google Apps

There is **no** advertising in the [Google Apps Core Services](#), and we have no plans to change this in the future. Google does not collect, scan or use data in Google Apps Core Services for advertising purposes. Customer administrators can restrict access to Non-Core Services from the Google Apps Admin console. Google indexes customer data to provide beneficial services, such as spam filtering, virus detection, spellcheck and the ability to search for emails and files within an individual account.

Data Access and Restrictions

Administrative access

To keep data private and secure, Google logically isolates each customer's Google Apps data from that of other customers and users, even when it's stored on the same physical server. Only a small group of Google employees have access to customer data. For Google employees, access rights and levels are based on their job function and role, using the concepts of least-privilege and need-to-know to match access privileges to defined responsibilities. Google employees are only granted a limited set of default permissions to access company resources, such as employee email and Google's internal employee portal. Requests for additional access follow a formal process that involves a request and an approval from a data or system owner, manager, or other executives, as dictated by Google's security policies. Approvals are managed by workflow tools that maintain audit records of all changes. These tools control both the modification of authorization settings and the approval process to ensure consistent application of the approval policies. An employee's authorization settings are used to control access to all resources, including data and systems for Google Apps products. Support services are only provided to authorized customer administrators whose identities have been verified in several ways. Googler access is monitored and audited by our dedicated security, privacy, and internal audit teams.

For customer administrators

Within customer organizations, administrative roles and privileges for Google Apps are configured and controlled by the customer. This means that individual team members can manage certain services or perform specific administrative functions without gaining access to all settings and data. Integrated audit logs offer a detailed history of administrative actions, helping customers monitor internal access to data and adherence to their own policies.

Law enforcement data requests

The customer, as the data owner, is primarily responsible for responding to law enforcement data requests; however, like other technology and communications companies, Google may receive direct requests from governments and courts around the world about how a person has used the company's services. We take measures to protect customers' privacy and limit excessive requests while also meeting our legal obligations. Respect for the privacy and security of data you store with Google remains our priority as we comply with these legal requests. When we receive such a request, our team reviews the request to make sure it satisfies legal requirements and Google's policies. Generally speaking, for us to comply, the request must be made in writing, signed by an authorized official of the requesting agency and issued under an appropriate law. If we believe a request is overly broad, we'll seek to narrow it, and we push back often and when necessary. For example, in 2006 Google was the only major search company that refused a U.S. government request to hand over two months of user search queries. We objected to the subpoena, and eventually a court denied the government's request. In some cases we receive a request for all information associated with a Google account, and we may ask the requesting agency to limit it to a specific product or service. We believe the public deserves to know the full extent to which governments request user information from Google. That's why we became the first company to start regularly publishing reports about government data requests. Detailed information about data requests and Google's response to them is available in our [Transparency Report](#). It is Google's policy to notify customers about requests for their data unless specifically prohibited by law or court order.

We believe the public deserves to know the full extent to which governments request user information from Google. That's why we became the first company to start regularly publishing reports about government data requests.

Third-party suppliers

Google directly conducts virtually all data processing activities to provide our services. However, Google may engage some [third-party suppliers](#) to provide services related to Google Apps, including customer and technical support. Prior to onboarding third-party suppliers, Google conducts an assessment of the security and privacy practices of third-party suppliers



to ensure they provide a level of security and privacy appropriate to their access to data and the scope of the services they are engaged to provide. Once Google has assessed the risks presented by the third-party supplier, the supplier is required to enter into appropriate security, confidentiality, and privacy contract terms.

Regulatory Compliance

Our customers have varying regulatory [compliance](#) needs. Our clients operate across regulated industries, including finance, pharmaceutical and manufacturing.

Google contractually commits to the following:

- Google will maintain adherence to ISO 27001 and SOC 2/3 audits during the term of the agreement.
- Defined Security Standards. Google will define how data is processed, stored, and protected through specific defined security standards.
- Access to our Data Privacy Officer. Customers may contact Google's Data Privacy Officer for questions or comments.
- Data Portability. Administrators can export customer data in [standard formats](#) at any time during the term of the agreement. Google does not charge a fee for exporting data.

Data processing amendment

Google takes a global approach to our commitments on data processing. Google and many of our customers operate in a global environment. We offer **all** of our users the same high level of protections via our [data processing amendment](#). The commitments in our data processing amendment are designed to facilitate compliance with jurisdictional-specific laws or regulations. Your organization can opt into our data processing amendment by following the instructions in our [Help Center](#).

EU Data Protection Directive

The Article 29 Working Party is an independent European advisory body focused on data protection and privacy. They have provided guidance on how to meet European data privacy requirements when engaging with cloud computing providers. Google provides capabilities and contractual commitments created to meet data protection recommendations provided by the Article 29 Working Party.

U.S.-EU and U.S.-Swiss Safe Harbor Frameworks

More than half of Google's business customers are based outside of the United States, with many of them operating in Europe. These businesses must comply with the European Commission's Data Protection Directive, which regulates the transfer of personal data within the European Union. The [U.S.-EU Safe Harbor Framework](#) provides a method for European companies to transfer personal data outside the European Union in a manner that's consistent with the Directive. Google [certifies](#) that it adheres to its principles and to those of the U.S.-Swiss Safe Harbor Framework.

EU model contract clauses

In 2010, the European Commission approved model contract clauses as a means of compliance with the requirements of the Directive. The effect of this decision is that by incorporating certain provisions into a contract, personal data can flow from those subject to the Directive to providers outside the EU or the European Economic Area. Google has a broad customer base in Europe. By adopting [EU model contract clauses](#), we're offering customers an additional option for compliance with the Directive.

U.S. Health Insurance Portability and Accountability Act (HIPAA)

Google Apps supports our customers' compliance with the U.S. Health Insurance Portability and Accountability Act (HIPAA), which governs the confidentiality and privacy of protected health information (PHI). Customers who are subject to HIPAA and wish to use Google Apps with PHI must sign a [business associate agreement \(BAA\)](#) with Google. The BAA covers Gmail, Google Calendar, Google Drive, Google Sites and Google Apps Vault.

U.S. Family Educational Rights and Privacy Act (FERPA)

More than 30 million students rely on Google Apps for Education. Google Apps for Education services comply with FERPA (Family Educational Rights and Privacy Act) and our commitment to do so is included in our agreements.

Children's Online Privacy Protection Act of 1998 (COPPA)

Protecting children online is important to us. We contractually require Google Apps for Education schools to obtain parental consent that COPPA calls for to use our services, and our services can be used in compliance with COPPA.



Empowering Users and Administrators to Improve Security and Compliance

Google builds security into its structure, technology, operations and approach to customer data. Our robust security infrastructure and systems become the default for each and every Google Apps customer. But beyond these levels, users are actively empowered to enhance and customize their individual security settings to meet their business needs through dashboards and account security wizards. Google Apps also offers administrators full control to configure infrastructure, applications and system integrations in a single dashboard via our Admin console — regardless of the size of the organization. This approach simplifies administration and configuration. Consider deployment of DKIM (a phishing prevention feature) in an on-premise email system. Administrators would need to patch and configure every server separately, and any misconfiguration would cause a service outage. Using our Admin console, DKIM is configured in minutes across thousands or hundreds of thousands of accounts with peace of mind and no outage or maintenance window required. Administrators have many powerful tools at their disposal, such as authentication features like 2-step verification and single sign-on, and email security policies like secure transport (TLS) enforcement, which can be configured by organizations to meet security and system integration requirements. Below are some key features that can help customize Google Apps for your security and compliance needs:

User authentication/authorization features

2-step verification

[2-step verification](#) adds an extra layer of security to Google Apps accounts by requiring users to enter a verification code in addition to their username and password when they sign in. This can greatly reduce the risk of unauthorized access if a user's password is compromised. Verification codes are delivered on a one-time basis to a user's Android, BlackBerry, iPhone, or other mobile phone. Administrators can choose to turn on 2-step verification for their domain at any time.

Security Key

[Security Key](#) is an enhancement for 2-step verification. Google, working with the FIDO Alliance standards organization, developed the Security Key — an actual physical key used to access your Google Account. It sends an encrypted signature rather than a code, and helps ensure that your login cannot be phished. Google for Work admins will be able to easily deploy, monitor and manage the Security Key at scale with new controls in the Admin console with no additional software to install. IT admins will see where and when employees last used their keys with usage tracking and reports. If Security Keys are lost, admins can easily revoke access to those keys and provide backup codes so employees can still sign-in and get work done.

Single sign-on (SAML 2.0)

Google Apps offers customers a [single sign-on \(SSO\) service](#) that lets users access multiple services using the same sign-in page and authentication credentials. It is based on SAML 2.0, an XML standard that allows secure web domains to exchange user authentication and authorization data. For additional security, SSO accepts public keys and certificates generated with either the RSA or DSA algorithm. Customer organizations can use the SSO service to integrate single sign-on for Google Apps into their LDAP or other SSO system.

OAuth 2.0 and OpenID Connect

Google Apps supports [OAuth 2.0 and OpenID Connect](#), an open protocol for authentication and authorization. This allows customers to configure one single sign-on service (SSO) for multiple cloud solutions. Users can log on to third-party applications through Google Apps—and vice versa—without re-entering their credentials or sharing sensitive password information.

Data management features

Information Rights Management (IRM)

With [Information Rights Management](#) ("IRM") you can disable downloading, printing and copying from the advanced sharing menu — perfect for when the file you're sharing is only meant for a few select people. This new option is available for any file stored in Google Drive, including documents, spreadsheets and presentations created in Google Docs.

Google Apps also offers administrators full control to configure infrastructure, applications and system integrations in a single dashboard via our Admin console — regardless of the size of the organization.

Drive audit log

The [Drive audit log](#) lists every time your domain's users view, create, update, delete or share Drive content. This includes content you create in Google Docs, Sheets, Slides and other Google Apps, as well as content created elsewhere that you upload to Drive, such as PDFs and Word files.

Drive content compliance / alerting

Google Apps for Work has [an additional feature](#) that allows Administrators to keep track of when specific actions are taken in Drive and can set up [custom Drive alerts](#). So if you want to know when a file containing the word "confidential" in the title is shared outside the company, now you'll know. And there are more events coming to Drive audit, including download, print and preview alerts.

Trusted domains for Drive sharing

Google Apps for Work and Education administrators will allow for domain [whitelisting](#). End users can share to those trusted domains, but can't share to other external domains. Great for partnerships, subsidiaries or other arrangements where certain domains are trusted and users are allowed to share to them.

Email Security features

Secure transport (TLS) enforcement

Google Apps administrators can require that email to or from specific domains or email addresses be encrypted with [Transport Layer Security \(TLS\)](#). For instance, a customer organization may choose to transmit all messages to its outside legal counsel via a secure connection. If TLS is not available at a specified domain, inbound mail will be rejected and outbound mail will not be transmitted.

Phishing prevention

Spammers can sometimes forge the "From" address on an email message so that it appears to come from a reputable organization's domain. Known as [phishing](#), this practice is often an attempt to collect sensitive data. To help prevent phishing, Google participates in the [DMARC program](#), which lets domain owners tell email providers how to handle unauthenticated messages from their domain. Google Apps customers can implement DMARC by creating a DMARC record within their admin settings and implementing an SPF record and DKIM keys on all outbound mail streams.

Email content compliance

Administrators can choose to scan Google Apps email messages for [predefined sets of words, phrases, text patterns or numerical patterns](#). They can create rules that either reject matching emails before they reach their intended recipients or deliver them with modifications. Customers have used this setting to monitor sensitive or restricted data, such as credit card information, internal project code names, URLs, telephone numbers, employee identification numbers, and social security numbers.

Objectionable content

The [objectionable content](#) setting enables administrators to specify what action to perform for messages based on custom word lists. With objectionable content policies, administrators choose whether messages containing certain words (such as obscenities) are rejected or delivered with modifications; for example, to notify others when the content of a message matches the rules that you set. Administrators can also configure this setting to reject outbound emails that may contain sensitive company information; for example, by setting up an outbound filter for the word *confidential*.

Restricted email delivery

By default, users with Gmail accounts at your domain can send mail to and receive mail from any email address. However, in some cases, administrators may want to [restrict the email addresses](#) your users can exchange mail with. For example, a school might want to allow its students to exchange mail with the faculty and other students, but not with people outside of the school. Use the Restrict delivery setting to allow the sending or receiving of email messages only from addresses or domains that administrators specify. When administrators add a Restrict delivery setting, users cannot communicate with anyone, except those authorized. Users who attempt to send mail to a domain not listed will see a message that specifies a policy prohibiting mail to that address, confirming that the mail is unsent. Users receive only authenticated messages from listed domains. Messages sent from unlisted domains—or messages from listed domains that can't be verified using DKIM or SPF records—are returned to the sender with a message about the policy.

Google Apps administrators can require that email to or from specific domains or email addresses be encrypted with Transport Layer Security (TLS).

eDiscovery features

eDiscovery allows organizations to stay prepared in case of lawsuits and other legal matters. [Google Vault](#) is the eDiscovery solution for Google Apps that lets customers retain, archive, search and export their business Gmail. Administrators can also search and export files stored in Google Drive.

Email retention policy

[Retention rules](#) control how long certain messages in your domain are retained before they are removed from user mailboxes and expunged from all Google systems. Google Apps allows you to set a default retention rule for your entire domain. For more advanced implementations, [Google Vault](#) allows administrators to create custom retention rules to retain specific content. This advanced configuration allows administrators to specify the number of days to retain messages, whether to delete them permanently after their retention periods, whether to retain messages with specific labels, and whether to let users manage email deletion themselves.

Legal holds

[Google Vault](#) allows administrators to place [legal holds](#) on users to preserve all their emails and on-the-record chats indefinitely in order to meet legal or other retention obligations. You can place legal holds on all content in a user's account, or target specific content based on dates and terms. If a user deletes messages that are on hold, the messages are removed from the user's view, but they are not deleted from Google servers until the hold is removed.

Search/discovery

[Google Vault](#) allows administrators to [search Gmail and Drive accounts](#) by user account, organizational unit, date or keyword. Search results include email, on-the-record chats, Google file types and non-Google file types such as PDF, DOCX and JPG.

Evidence export

[Google Vault](#) allows administrators to have the ability to [export](#) specific email, on-the-record chats and files to standard formats for additional processing and review in a manner that supports legal matters while respecting chain of custody guidelines.

Support for third-party email platforms

The [comprehensive mail storage setting](#) ensures that a copy of all sent or received mail in your domain—including mail sent or received by non-Gmail mailboxes—is stored in the associated users' Gmail mailboxes. For organizations that reroute mail to non-Gmail mail servers, this setting also ensures storage of mail in Gmail mailboxes for archiving and eDiscovery purposes.

Securing endpoints

Mobile device management (MDM)

[Mobile device management](#) in Google Apps eliminates the need for on-premises device or third-party management solutions.

Administrators can enforce policies over mobile devices in their organization, encrypt data on devices, and perform actions like remotely wiping or locking lost or stolen devices. This type of control helps ensure the security of business data, even if employees choose to work on their personal phones and tablets. Mobile device management in Google Apps works with Android, iOS, Windows Phone, and smartphones and tablets using Microsoft Exchange ActiveSync, such as BlackBerry 10.

Policy-based Chrome browser security

All of the tools and features in Google Apps are best supported by Google Chrome. **Administrators can apply [security and usage policies](#) across Windows, OSX, Linux, iOS, and Android.**

Chrome's standard security features include Safe Browsing, sandboxing, and managed updates that protect users from malicious sites, viruses, malware, and phishing attacks. There are also measures in place to prevent cross-site scripting, which attackers can use to steal private data. Google Apps administrators can deploy Chrome for Work across their organization and customize it to meet their needs.

Over [280 policies](#) help administrators control how employees use Chrome across devices. For example, administrators can enable automatic updates to get the latest security fixes, block or allow specific apps, and configure support for legacy browsers.

Chrome device management

The Google Apps Admin Console applies policy to [Chrome devices](#) such as Chromebooks, Chromeboxes, and Chromebox for meetings, which are fast, secure, and cost-effective computers that run Chrome as an operating system. Administrators can easily manage security and other settings for their organization's Chrome devices from a single place. They can configure Chrome features for their users, set up access to VPNs and WiFi networks, pre-install apps and extensions, restrict sign-in to certain users, and more.

Administrators can enforce policies over mobile devices in their organization, encrypt data on devices, and perform actions like remotely wiping or locking lost or stolen devices.

Data Recovery

Restore a recently deleted user

An administrator can [restore a deleted user account](#) for up to five days after date of deletion. After five days, the Admin console permanently deletes the user account, and it can't be restored, even if you contact Google technical support. Please note that only customer Administrators can delete accounts.

Restore a user's Drive or Gmail data

An administrator can restore a user's Drive or Gmail data for up to 25 days after the data is removed from the user's trash. Google will delete all Customer-deleted data from its systems as soon as reasonably practicable and within a maximum period of 180 days.

An administrator can restore a user's Drive or Gmail data for up to 25 days after date of deletion.

Security reports

Google Apps administrators have access to [security reports](#) that provide vital information on their organization's exposure to data compromise. They can quickly discover which particular users pose security risks by eschewing 2-step verification, installing external apps, or sharing documents indiscriminately. Administrators can also choose to [receive alerts](#) when suspicious login activity occurs, indicating a possible security threat.

Conclusion

The protection of user data is a primary design consideration for all of Google's infrastructure, applications and personnel operations. Protection of user data is far from being an afterthought or the focus of occasional initiatives, it's an integral part of what we do. We believe that Google can offer a level of protection that very few can match. Because protecting your data is part of our core business, Google can develop security innovations such as 2-step authentication and stronger encryption methods. We are able to make extensive investments in security, resources and expertise at a scale that few can afford. Our scale of operations and collaboration with the security research community enable Google to address vulnerabilities quickly or prevent them entirely. Google's security and operational procedures are verified by independent third-party auditors.

Data protection is more than just security. Google's strong contractual commitments make sure our customers maintain control over the data and how it is processed, including the assurance that your data is not used for advertising or any other purpose than to deliver Google Apps services.

For these reasons and more, over five million organizations across the globe, including 64 percent of the Fortune 500, trust Google with their most valuable asset: their information. Google will continue to invest in our platform to allow our users to benefit from our services in a secure and transparent manner.

Data protection is more than just security. Google's strong contractual commitments make sure our customers maintain control over the data and how it is processed, including the assurance that your data is not used for advertising or any other purpose than to deliver Google Apps services.



Google Apps Service Level Agreement

Google Apps SLA. During the Term of the applicable Google Apps Agreement (the "Agreement"), the Google Apps Covered Services web interface will be operational and available to Customer at least 99.9% of the time in any calendar month (the "Google Apps SLA"). If Google does not meet the Google Apps SLA, and if Customer meets its obligations under this Google Apps SLA, Customer will be eligible to receive the Service Credits described below. This Google Apps SLA states Customer's sole and exclusive remedy for any failure by Google to meet the Google Apps SLA.

Definitions. The following definitions shall apply to the Google Apps SLA.

"Downtime" means, for a domain, if there is more than a five per cent user error rate. Downtime is measured based on server-side error rate.

"Google Apps Covered Services" means the Gmail, Google Calendar, Google Talk, Google Docs and Drive, Google Groups, Google Sites and Google Apps Vault components of the Service. This does not include the Gmail Labs functionality, Google Apps – Postini Services, Gmail Voice or Video Chat components of the Service.

"Monthly Uptime Percentage" means total number of minutes in a calendar month minus the number of minutes of Downtime suffered in a calendar month, divided by the total number of minutes in a calendar month.

"Service" means the Google Apps for Business service (also known as Google Apps Premier Edition), Google Apps for Government service, Google Apps for ISPs service (also known as Google Apps Partner Edition), Google Apps for Education service (also known as Google Apps Education Edition) or Google Apps Vault (as applicable) provided by Google to Customer under the Agreement.

"Service Credit" means the following:

Monthly Uptime Percentage	Days of Service added to the end of the Service term (or monetary credit equal to the value of days of service for monthly postpay billing customers), at no charge to Customer
< 99.9% - >= 99.0%	3
< 99.0% - >= 95.0%	7
< 95.0%	15

Customer Must Request Service Credit. In order to receive any of the Service Credits described above, Customer must notify Google within thirty days from the time Customer becomes eligible to receive a Service Credit. Failure to comply with this requirement will forfeit Customer's right to receive a Service Credit.

Maximum Service Credit. The aggregate maximum number of Service Credits to be issued by Google to Customer for all Downtime that occurs in a single calendar month shall not exceed fifteen days of Service added to the end of Customer's term for the Service (or the value of 15 days of service in the form of a monetary credit to a monthly billing customer's account). Service Credits may not be exchanged for, or converted to, monetary amounts, except for customers who are on Google's monthly billing plan.

Google Apps SLA Exclusions. The Google Apps SLA does not apply to any services that expressly exclude this Google Apps SLA (as stated in the documentation for such services) or any performance issues: (i) caused by factors described in the "Force Majeure" section of the Agreement; or (ii) that resulted from Customer's equipment or third-party equipment, or both (not within the primary control of Google).

State of Utah NASPO Cloud Solutions

Technical Proposal

Solicitation No: CH16012



UNISYS

TABLE OF CONTENTS

1.0	(M) Signature page (rfp Section 5.1)	1- ~~32~~
2.0	(M) Executive Summary (RFP Section 5.4)	2-1
3.0	Mandatory Minimums (RFP Section 5).....	3-4
4.0	Business Profile (RFP Section 6).....	4-1
5.0	Organization Profile (Rfp section 7).....	5-1
6.0	Technical response – aws – ms included (RFP Section 8).....	6-1
7.0	Confidential, protected or proprietary information.....	7-1
8.0	Exceptions and/or Additions to the Standard Terms and Conditions	8-1
Appendix 3 – AWS CAIQ		
Appendix 4 – Microsoft O365		
Appendix 5 – Microsoft Azure		
Appendix 6 – Microsoft Dynamics		
Appendix 7 – Google CAIQ		
Appendix 8 – AWS SLA		
Appendix 9 – Microsoft SLA		
Appendix 10 – Google Apps SLA		
Appendix 11 – Unisys-1		
Appendix 12 – AWS-1		
Appendix 13 – Google-1		
Appendix 14 – AWS Risk and Compliance White Paper		
Appendix 15 – Windows Azure Security and Compliance		
Appendix 16 – Google DPA		
Appendix 17 – Microsoft Online Services Terms		
Appendix 18 – Google Compute (Cloud) SLA		
Appendix 19 _ Getting Started With VMs on Windows Azure		
Appendix 20 – AWS+Access+Policy+(State)		
Appendix 21 –AWS+Access+Policy+(State)		
Appendix 22 – Google Apps For work		
Appendix 23 – Google Apps Security		
Appendix 24 – Windows Azure-Secure Privacy Compliance		

1.0 (M) SIGNATURE PAGE (RFP SECTION 5.1)


Section Title: RFP Signature Page. The Lead State's Request for Proposal Signature Page completed and signed. See Section 5.1 of the RFP.

5.1 (M) SIGNATURE PAGE Proposals must be submitted with a vendor information form, located on Bidsync as an attachment to the RFP, which must contain an ORIGINAL HANDWRITTEN signature executed in INK OR AN ELECTRONIC SIGNATURE, and be returned with the Offeror's proposal.

Unisys Response: Our vendor information form is below.



State of Utah Vendor Information Form

Lead Company Name (include division if applicable) Unisys Corporation		Federal Tax Identification Number 38-0387840		State of Utah Sales Tax ID Number 12278699-002-STC	
Outgoing Address 801 Lakeview Drive Ste. 100		City Bluebell	State PA	Zip Code 19422	
Remittance Address (if different from ordering address) 810 Lakeview Drive Ste. 100		City Bluebell	State PA	Zip Code 19422	
Type: ☐ Proprietorship ☐ Partnership ☐ Government ☒ For-Profit Corporation ☐ Non-Profit Corporation		Company Contact Person Mr. Harsh Bajpai			
Telephone Number (include area code) 703-439-6200		Fax Number (include area code)			
Company's Internet Web Address http://www.unisys.com		Email Address Harsh.Bajpai@unisys.com			
Offeror's Authorized Representative's Signature 					
Type or Print Name Mr. Harsh Bajpai					
Position or Title of Authorized Representative Director, Cloud Solutions and Services					
Date MARCH 10, 2016					

07/15/2014

2.0 (M) EXECUTIVE SUMMARY (RFP SECTION 5.4)

- **Section Title: Executive Summary.** The one or two page executive summary is to briefly describe the Offeror's Proposal. This summary should highlight the major features of the Proposal. It must indicate any requirements that cannot be met by the Offeror. The Lead State should be able to determine the essence of the Proposal by reading the executive summary. See Section 5.4 of the RFP.
- **(M) EXECUTIVE SUMMARY**
- **Offerors must provide an Executive Summary of its proposal. An Executive Summary should highlight the major features of an Offeror's proposal. Briefly describe the proposal in no more than three (3) pages. The evaluation committee should be able to determine the essence of the proposal by reading the Executive Summary. Any requirements that cannot be met by the Offeror must be included.**

Unisys Response:

Unisys appreciates the importance of this innovative and flexible solicitation. We are pleased to offer the cloud deployment models (private, community, public, and hybrid) and the underlying cloud services models (IaaS, PaaS, and SaaS) requested in this solicitation. We will also provide architecture, design, implementation, and value-add services that will support the various cloud solution options. The following table illustrates the depth and breadth of the offerings in our proposal.

	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Private Cloud	Unisys	Unisys	Unisys
Hybrid Cloud	Unisys	Unisys	Unisys
Cloud Professional Services	Unisys	Unisys	Unisys
Cloud Value Add Services	Unisys	Unisys	Unisys


Our cloud partners in this proposal are Amazon Web Services (AWS), Microsoft and Google. Unisys will resell some of the offerings through our own partner agreements and fill orders for the others directly, as required by the Participating Entity.

Unisys is offering the State of Utah and the Participating Entities a set of world-class, best-of-breed cloud providers. Proven solution options will provide the Participating States with a diverse set of products, solutions and services that are typically needed for its evolving requirements. Our ability to deliver all types of cloud solutions across all partners; seamlessly and expeditiously with our own certified and experienced staff will provide the Participating Entities with a quick reach-back vehicle that they can leverage as needed. Our proven, past performance in supporting cloud transformations and cloud implementations for clients in state, local, and Federal governments will provide the States with a qualified level of confidence in Unisys. Our strong, formal partnerships allow us direct access to the latest resources of our partners, which we can leverage expeditiously. Finally, Unisys can provide the Participating Entities with a one-stop shop for all types of cloud solutions and services. This will minimize risk to the State and establish transparency and vendor accountability, thereby improving the State's overall effectiveness and efficiency in all cloud-related matters.

The Unisys Solution Approach for NASPO – Transforming Today's Enterprise to Tomorrow's Digital Government

As a trusted agent and partner, Unisys has provided value-added information technology services for more than 50 years to clients in state, local, U.S. Federal, and foreign governments in 180 countries. As the prime holder of the WSCA Public Cloud Hosting Service contract, Unisys exceeds current contract

requirements by providing best-in-class experience, quality, and timely service delivery. We pride ourselves as being an honest broker because we understand our clients' requirements, including those of foreign governments throughout the world.



Our ability to rapidly deploy needed resources and specialized services is unparalleled. Our efficient and effective approach focuses on providing precisely the services required at the best possible cost. Our subject matter experts power a portfolio of capabilities aligned with our clients' changing priorities with a focus on delivering next-generation Digital Government Solutions and Services.



UNISYS Digital Government Transformational Business Benefits

It is critical in today's digital environment that the State partners with the right organization that can bring the knowledge, skills, and domain expertise to the partnership. As the incumbent, our goal is to build on the quality service the State receives today. Unisys has a deep appreciation and understanding of providing mission-critical and innovative solutions to the public sector. Our combined cloud hosting and services capabilities in our proposal will exceed the State's key objectives as follows:

- Service enablement:** The pace of change is accelerating faster than public organizations can adjust their plans, investments, and—more important—their employee base. The pressure to adopt cloud services is intense; adopting cloud services haphazardly could be detrimental. There are persistent concerns for security, compliance, regulations, and cost. Unisys brings staff to deliver comprehensive consulting, integration, and support services that enable clients to migrate to cloud services while maintaining operational compatibility and transforming their existing infrastructure for the delivery of an agile IT-as-a-Service model. CloudBuild Services is a multivendor offering for the planning, design, and implementation of an enterprise cloud. These services will help the State to avoid the “cloud-in-a-corner” syndrome, decrease risk in migrating to cloud services, and decrease potential project failures and budget overruns.

- **Improved security and comfort level:** As more and more of an organization’s mission-critical and day-to-day operations move to the cloud, security becomes much more important. Unisys brings our award-winning Stealth security products, which provide a smarter way to manage risk in enterprises. Stealth is now fully integrated with the Amazon Web Services cloud infrastructure. Stealth on AWS will provide the State with a trusted path to extend its security protections to the cloud, unlocking huge cost savings and providing needed agility to respond to constituents’ concerns.
- **Cost efficiencies:** Unisys understands the financial benefits available to users of cloud computing services. In particular, the very competitive pay-as-you-go cloud computing model in our proposal, paired with full control over fine-grained characteristics of the State’s cloud-based infrastructures, will enable users to achieve significant savings in ongoing operating expenses.
- **Flexibility and scalability:** Orchestration in a cloud computing environment empowers Unisys clients to have complete control over the timing and availability of IT resources. We combine tasks into workflows so the provisioning and management of various IT components and their associated resources can be automated. The productivity gained from the ability to provision services in minutes (not days or weeks, as in many IT environments) enables users to respond to changing conditions and new requirements quickly without sacrificing service delivery quality. Our proposal exceeds the RFP requirements by enabling advanced automation capabilities not available in most cloud computing environments.
- **Significant reduction of support staff time:** By freeing resources currently performing easily automatable tasks, and adopting a “factory” model for service delivery, IT organizations can reallocate staff to more productive uses that improve the ability of states and other entities to meet their goals sooner or more effectively. Our proposal fully supports the needs of Participating Entities to offload activities to automated environments in the cloud computing ecosystem.

“Of the IT vendors, Unisys continues to have one of the clearest sets of offerings for secure private, public and hybrid models.” – 451 Research

3.0 MANDATORY MINIMUMS (RFP SECTION 5)

5.2 (M) COVER LETTER

Proposals must include a cover letter on official letterhead of the Offeror. The cover letter must identify the RFP Title and number, and must be signed by an individual authorized to commit the Offeror to the work proposed. In addition, the cover letter must include:

March 10, 2016

Christopher Hughes, Assistant Director
State of Utah, Division of Purchasing
3150 State Office Building, Capitol Hill
Salt Lake City, UT 84114-1061

Subject: Unisys Response to Utah Solicitation Number CH16012 – NASPO ValuePoint Cloud Hosting Services

Dear Mr. Hughes:


Unisys is pleased to provide our response to Solicitation Number CH16012 – NASPO ValuePoint Master Agreement for Public Cloud Hosting Services. We sincerely appreciate the importance of this solicitation.

Unisys is pleased to offer the requested cloud deployment models (private, community, public, and hybrid) and the underlying cloud services models (IaaS, PaaS, and SaaS) as part of our proposal. We will also provide architecture, design, implementation, and value-add services that will support the various cloud solution options. Our cloud service provider partners in this proposal are Amazon Web Services (AWS), Microsoft, and Google.

Unisys believes that offering the State of Utah and the Participating Entities a set of world-class, best-of-breed cloud providers is a win-win proposition for the State. Proven solution options will provide the State with a full set of services that are typically needed for its diverse IT ecosystems. Our ability to deliver cloud solutions seamlessly and expeditiously with our own certified and experienced staff will provide Participating Entities with a quick reach-back vehicle that they can leverage as needed. Our proven past performance in supporting cloud transformations and implementations for clients in state, local, and Federal governments will provide the State of Utah with a qualified level of confidence in Unisys. Our strong, formal partnerships with AWS, Microsoft and Google allow us direct access to the latest resources of our partners, which we can leverage expeditiously. Finally, Unisys can provide the Participating Entities with a one-stop shop for cloud solutions and services. This will minimize risk to the State and establish transparency and vendor accountability, thereby improving the State's overall effectiveness and efficiency in cloud-related matters.

Unisys, the prime contractor, is a public Fortune 500 corporation (NYSE: UIS) that operates in more than 100 countries and has more than 20,000 associates. For more than 50 years, we have provided the U.S. Federal and state governments with mission-critical information technology (IT) solutions and related consulting services, including small- and large-scale Cloud Services. More than 45 percent of our revenue of \$4 billion comes from our valued long-term government clients. Understanding our client needs by working closely with clients with our consultants is a cornerstone of our heritage and success in government and business in the United States and around the world.

Sincerely,



Harsh Bajpai, Unisys Corporation
Director, Cloud Solutions & Services

Email: Harsh.Bajpai@unisys.com; Telephone: 703-439-6200

5.2.1 A statement indicating the Offeror's understanding that they may be required to negotiate additional terms and conditions, including additional administrative fees, with Participating Entities when executing a Participating Addendum.

Unisys understands that we may be required to negotiate additional terms and conditions, including additional administrative fees, with Participating Entities when executing a Participating Addendum.

5.2.2 A statement naming the firms and/or staff responsible for writing the proposal.

Unisys is responsible for the writing of this proposal with information provided from our partners.

5.2.3 A statement that Offeror is not currently suspended, debarred or otherwise excluded from federal or state procurement and non-procurement programs.

Unisys and our partners for this proposal are not currently suspended, debarred, or otherwise excluded from Federal or state procurement and non-procurement programs.

5.2.4 A statement acknowledging that a 0.25% NASPO ValuePoint Administrative Fee and any Participating Entity Administrative fee will apply to total sales for the Master Agreement(s) awarded from the RFP.

Unisys acknowledges the statement above.

5.2.5 A statement identifying the service model(s) (SaaS, IaaS, and/or PaaS) and deployment model(s) that it is capable of providing under the terms of the RFP. See Attachment C for a determination of each service model subcategory. The services models, deployment models and risk categories can be found in the Scope of Services, Attachment D. Note: Multiple service and/or deployment model selection is permitted, and at least one service model must be identified. See Attachment H.

5.2.6 A statement identifying the data risk categories that the Offeror is capable of storing and securing. See Attachment D and Attachment H.

Unisys is pleased to offer all four cloud deployment models: public cloud, private cloud, community cloud, and hybrid cloud. Under these deployment models, we are also pleased to offer all three cloud service delivery methods: IaaS, PaaS, and SaaS.

Attainment of certifications is a lengthy ongoing process for vendors. Additionally, certification standards are constantly evolving. Our partners in this proposal are at a FISMA Moderate level. Therefore, across our cloud deployment models and delivery options, we are positioned to secure and store FISMA Low and Moderate risk data types. For FISMA High, Unisys Stealth for data in motion, when combined with encryption technologies for data at rest, can meet the requirements of FISMA High risk data. For the private cloud model, we can design a cloud environment based on the required standards and controls defined by the client.

In accordance with the requirements of this RFP, we completed Attachment H – Identification of Service Models Matrix and provided it in **Attachment H**.

Attachment H – Identification Service Models

Service Model:	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered:
SaaS	Unisys AWS Microsoft Google	Unisys AWS Microsoft Google	See response above	All
IaaS	Unisys AWS Microsoft	Unisys AWS Microsoft	See response above	All
PaaS	Unisys AWS Microsoft	Unisys AWS Microsoft	See response above	All

5.3 (M) ACKNOWLEDGEMENT OF AMENDMENTS

If the RFP is amended, the Offeror must acknowledge each amendment with a signature on the acknowledgement form provided with each amendment. Failure to return a signed copy of each amendment acknowledgement form with the proposal may result in the proposal being found non-responsive.

Unisys acknowledges the amendments released under this RFP. Unisys provides our signed copy of the amendment acknowledgement form below.

AMENDMENT ACKNOWLEDGEMENT FORM**ACKNOWLEDGEMENT OF AMENDMENTS TO RFP (SOLICITATION CH16012)**

This attachment represents that the Offeror has read, reviewed, and understands the totality of Solicitation CH16012, including the final RFP document posted on February 10, 2016.

By signing below, the Offeror attest to reviewing the documents listed above.

UNISYS CORPORATION
Offeror
[Signature]
Representative Signature

5.4 (M) EXECUTIVE SUMMARY

Offerors must provide an Executive Summary of its proposal. An Executive Summary should highlight the major features of an Offeror's proposal. Briefly describe the proposal in no more than three (3) pages. The evaluation committee should be able to determine the essence of the proposal by reading the Executive Summary. Any requirements that cannot be met by the Offeror must be included.

Unisys provides our executive summary in Section 2.0, Executive Summary.

5.5 (M) GENERAL REQUIREMENTS

5.5.1 *Offeror must agree that if awarded a contract it will provide a Usage Report Administrator responsible for the quarterly sales reporting described the Master Agreement Terms and Conditions, and if applicable Participating Addendums.*

If Unisys is awarded this contract, we will provide a Usage Report Administrator responsible for the quarterly sales reporting described in the Master Agreement Terms and Conditions, and if applicable, Participating Addendums.

5.5.2 *Offeror must provide a statement that it agrees to cooperate with NASPO ValuePoint and SciQuest (and any authorized agent or successor entity to SciQuest) with uploading an Offeror's ordering instructions, if awarded a contract.*

Unisys agrees with the statement above.

5.5.3 *Offeror must at a minimum complete, provide, and maintain a completed CSA STAR Registry Self-Assessment. Offeror must either submit a completed The Consensus Assessments Initiative Questionnaire (CAIQ), Exhibit 1 to Attachment B, or submit a report documenting compliance with Cloud Controls Matrix (CCM), Exhibit 2 to Attachment B. Offeror must also represent and warrant the accuracy and currency of the information on the completed. Offerors are encouraged to complete and submit both exhibits to Attachment B.*

5.5.4 *Offeror, as part of its proposal, must provide a sample of its Service Level Agreement, which should define the performance and other operating parameters within which the infrastructure must operate to meet IT System and Purchasing Entity's requirements.*

In the private cloud space, government clients contract Unisys to design and implement a cloud platform on their premises. We build this platform to order in accordance with the client's requirements for security, governance, compliance, certification, and accreditation. Unisys then can manage a custom environment for the client or return it to the client to manage it. In the hybrid cloud space, government clients run a private environment that mirrors an external environment, which could be public, community, or private. Therefore, we are providing responses to RFP requirements 5.5.3 and 5.5.4 as they apply to our cloud partners in this proposal.

Appendix 3 – AWS CAIQ

Appendix 4 – Microsoft Office365 CCM

Appendix 5 – Microsoft Azure CCM

Appendix 6 – Microsoft Dynamics CCM

Appendix 7 – Google CAIQ

The appendix to this proposal provides our collective responses to RFP requirement 5.5.3 here:

Appendix 8 – AWS SLA

Appendix 9 – Microsoft SLA

Appendix 10 – Google Apps SLA

5.7 RECERTIFICATION OF MANDATORY MINIMUMS AND TECHNICAL SPECIFICATIONS

Offeror must acknowledge that if it is awarded a contract under the RFP that it will annually certify to the Lead State that it still meets or exceeds the technical capabilities discussed in its proposal.

Unisys acknowledges that if we are awarded a contract under the RFP, we will annually certify to the Lead State that we still meet or exceed the technical capabilities discussed in our proposal.

4.0 BUSINESS PROFILE (RFP SECTION 6)

Section Title: Business Profile: This section should constitute the Offeror's response to the items described in Section 6 of the RFP. An Offeror's response must be a specific point-by-point response, in the order listed, to each requirement in the Section 6 of the RFP.

6 BUSINESS INFORMATION

6.1 (M)(E) BUSINESS PROFILE

Provide a profile of your business including: year started, organizational structure, client base (including any focus by region, market sector, etc.), growth over the last three (3) years, number of employees, employee retention rates (specific for employees that may be associated with the services related to the RFP) over the last two (2) years, etc. Businesses must demonstrate a minimum of three (3) years of experience providing cloud solutions for large scale projects, including government experience, to be eligible for award.

Unisys Response:

Unisys Corporation was founded in 1886 and is based in Blue Bell, Pennsylvania. We provide information technology services and products worldwide. Approximately 88 percent of our revenue is derived from services. The remaining 12 percent of our revenue is derived from the sale of technology products and solutions. Cloud and infrastructure services represent 52 percent of our total services revenue. Sales of commercial products and services to various agencies of the U.S. Federal government represented about 19 percent of our total consolidated revenue in 2015. Our market capitalization is more than \$500 million. Our revenue for the year ending in 2014 was \$3.4 billion. We employ more than 20,000 associates worldwide. Our services organization in North America alone is 3400 persons strong. Our attrition rate for cloud and data center infrastructure services within North America is less than 23%.

Section 6.2 provides our experience in working on large-scale projects to provide government agencies with cloud solutions.

6.2 (M)(E) Scope of Experience

Describe in detail the business' experience with government or large consortium contracts similar to the Master Agreements sought through this RFP. Provide the approximate dollar value of the business' five (5) largest contracts in the last two (2) years, under which the Offeror provided Solutions identical or very similar to those required by this RFP. Government experience is preferred.

Unisys Response:

Scope of Experience: 1 Contract Term: Current Segment: U.S. Federal Government Solution/Offering: Private Cloud Storage	U.S. Department of the Treasury, Internal Revenue Service (IRS) GSA's Alliant Government-wide Acquisition Contract
On July 26, 2012, Unisys was awarded the Enterprise Storage Acquisition task order from the Internal Revenue Service (IRS) to provide private cloud-based storage for the agency's information records. The task order, awarded against GSA's Alliant Government-wide Acquisition Contract, had a 1-year base period and nine 1-year options worth up to \$139 million over 10 years. Unisys is fulfilling our third contract option year. Under the task order, Unisys acquired storage assets owned and managed by IRS and now manages more than 6 PB of storage space in several IRS data centers and facilities. We planned, built, deployed, and now maintain a new IRS storage environment, transitioning those assets to a private cloud-based SaaS model. Under this model, IRS pays for storage capacity as needed instead of purchasing storage equipment necessary to meet its peak demands. Using virtualization technology, Unisys mitigated	

Scope of Experience: 1 Contract Term: Current Segment: U.S. Federal Government Solution/Offering: Private Cloud Storage	U.S. Department of the Treasury, Internal Revenue Service (IRS) GSA's Alliant Government-wide Acquisition Contract
overcapacity of storage without the “stovepiped” effect and was able to manage storage as a single entity. We migrated more than 90 percent of relevant IRS data to the private storage cloud we operate for the IRS. The current contract is supported by an on-site Program Manager who started his career in the U.S. Navy. He has 25 years of technology experience in the public sector. There is also a Deputy Program Manager with 25 years of experience. Six storage engineers, a security manager, a business analyst, and an engagement manager are also on site. They collectively have about 10 years of industry experience. Two storage architects serve as subject matter experts who have more than 20 years of experience in designing large-scale, complex storage systems.	
Scope of Experience: 2 Contract Term: Current Segment: U.S. Federal Government Solution/Offering: Cloud Hosting and Migration Services	U.S. Department of the Interior (DOI) DOI Foundation Cloud Hosting Service Contract
On August 15, 2013, the U.S. Department of the Interior (DOI) tapped 10 companies to compete for a series of cloud computing projects that could involve a total investment of up to \$10 billion. DOI structured the deal among 10 systems integrators—Unisys, IBM, Verizon, AT&T, Lockheed Martin, Aquilent, Autonomic Resources, CGI, GTRI, and Smartronix—to compete for cloud migration; cloud management; cloud hosting; and a series of projects surrounding public, private, and hybrid cloud. Under this master contract, Unisys won and delivered the first task order to move DOI's Financial and Business Management System (FBMS) to the cloud, a move that provides significant flexibility while meeting DOI's stringent requirements for the high availability of data and applications. This move was announced as completed on January 5, 2016. The project enables DOI to dynamically reallocate cloud computing resources overnight and evaluate service level agreement performance alongside monthly invoices. Unisys worked with a partner, Virtustream, to provide and Infrastructure as Service (IaaS) solution, including Virtustream's SAP in the cloud hosting services. The Unisys Team collaborated closely with DOI to orchestrate a coordinated migration to the cloud and to manage the services once live. The successful migration made DOI the first Federal organization to move its SAP-based financial management application to the cloud. The project also complies with the criteria established in the Office of Management and Budget's Cloud First policy, which mandates that federal agencies take advantage of the cost savings and efficiencies of cloud computing.	

Scope of Experience: 3 Contract Term: Completed in 2015 Segment: U.S. Federal Government Solution/Offering: Public Cloud	General Services Administration (GSA) Enterprise Email and Collaborative Services Contract
In 2010, Unisys won the Enterprise Email and Collaborative Services Contract from the GSA to upgrade and migrate the agency's email systems from Lotus Notes and Domino to a public cloud-based email service provided by Google. This \$9.8 million, 5-year contract ended in 2015. In the first migration of a Federal agency's employees to a cloud-based email environment, Unisys supported the GSA's migration from its legacy Lotus Notes and Domino software to Google Apps for Government Software as a Service (SaaS) based on the FISMA-accredited Google Apps Premier platform. Google Apps for Government includes features that support collaboration, office productivity, and email delivered as SaaS over the Internet. Unisys recognized the need to support a growing GSA workforce; therefore, we proposed a customizable solution to accommodate increases in email volume and the number of user mailboxes. Unisys provides incremental email backup with flexible retention capabilities. Our solution also provides the GSA with the ability to add, change, and delete change mailbox accounts and directory data. With the migration completed in July 2011, Unisys supports 18,000 GSA employee mailboxes, each containing 30 GB of email storage.	
Scope of Experience: 4 Contract Term: Current Segment: U.S. State Government Solution/Offering: Public Cloud Hosting and Services (IaaS, PaaS, and SaaS)	State of Utah Western States Contracting Alliance (WSCA) / National Association of State Procurement Officials (NASPO) / ValuePoint Contract
In January 2013, Unisys was one of the four vendors awarded the WSCA Public Cloud Hosting and Services Contract. Unisys, in partnership with Amazon Web Services (AWS), delivers public cloud offerings from Amazon, Salesforce, and Google to the states participating in the WSCA/NASPO/ValuePoint contract. Unisys also provides assessment, design, and implementation services through our own workforce. In 2015, we booked an order totaling approximately \$1.5 million through this contract vehicle. We are providing the following list of sample Unisys clients under the current WSCA/NASPO/ValuePoint contract with a brief description of each solution offering. 1. State of Washington a) Department of Natural Resources: IaaS via AWS. SAP upgrade project. Amazon S3 Storage. b) Department of Natural Resources: IaaS & PaaS via AWS. Application DevOps infrastructure in Amazon.	

Scope of Experience: 4 Contract Term: Current Segment: U.S. State Government Solution/Offering: Public Cloud Hosting and Services (IaaS, PaaS, and SaaS)	State of Utah Western States Contracting Alliance (WSCA) / National Association of State Procurement Officials (NASPO) / ValuePoint Contract
c) Department of Ecology: IaaS via AWS. Application Development and Testing. d) Washington State Attorney General's Office: IaaS. e) Department of Health: IaaS via AWS. f) OCIO Office: SaaS via Salesforce. g) Department of Transportation: IaaS via AWS. h) Department of Social and Health Services: IaaS via AWS. 1. State of Wyoming a) Office of the Chief Information Officer (OCIO) Application Development and Testing in Amazon, AWS 2. State of Arizona a) Department of Transportation: IaaS via AWS. b) CIO Office: Cloud Broker via Pantheon. c) Arizona Strategic Enterprise Technology Office: IaaS via AWS d) Department of Health Services: IaaS via AWS 3. State of Hawaii a) Cloud Broker between State of Hawaii and Pacxa, one of the largest local IT firms in Hawaii. 4. State of Utah a) Department of Technology Services: SaaS. 5. State of Iowa a) State of Iowa: IaaS via AWS. b) Iowa Workforce Development: IaaS via AWS. 6. State of Missouri a) Intelligent Application Alignment (IAA) via Unisys: This is a Cloud Readiness assessment service that does discovery, assessment, and cloud portability of workloads in a given environment. 7. State of Montana a) State of Montana Department of Justice: Public Cloud Hosting via AWS. 8. State of Wisconsin	

Scope of Experience: 4 Contract Term: Current Segment: U.S. State Government Solution/Offering: Public Cloud Hosting and Services (IaaS, PaaS, and SaaS)	State of Utah Western States Contracting Alliance (WSCA) / National Association of State Procurement Officials (NASPO) / ValuePoint Contract
a) Department of Administration: Intelligent Application Alignment (IaaS) via Unisys. This is a Cloud Readiness assessment service that does discovery, assessment, and cloud portability of workloads in a given environment.	

Scope of Experience: 5 Contract Term: Current Sector: U.S. State Government Solution/Offering: Hybrid (Public/Private) Cloud and Data Center Consolidation	Commonwealth of Pennsylvania (CoPA) Data PowerHouse (DPH) / Pennsylvania Compute Services (PACS)
In 1999, Unisys was contracted to consolidate CoPA’s data center environments to one main data center (the Data PowerHouse) with an associated disaster recovery center. We were also enlisted to provide ongoing data center services, including server management; system engineering; systems operations, maintenance, and support; data center operations; data storage and archiving; information security engineering; application hosting; and disaster recovery. We also provide hardware and software provisioning; data center infrastructure maintenance; physical and logical security; and data security services that include data loss protection, malware protection, intrusion prevention and scanning. Our support model was based on ITIL V3 processes. In 2014, the term of this engagement was extended to 2022, with the scope to include the design, build, implementation, and hosting of a hybrid cloud environment for CoPA and its more than 50 Commonwealth wide agencies. This follow-on contract is called PACS (PA Compute Services). The contract terms for this engagement also were extended to 2022. The total contract value is \$1.61 billion. Every year, this contract produces \$100 million in revenue. On this engagement, Unisys provides the full breadth of cloud brokering solutions, including design, implementation, hosting, managed services, and technology integration. As part of the extension, Unisys deployed the ServiceNow platform for Service Management including governance, policy and process deployment and the full ITIL complement of service delivery. As part of the effort, Unisys built a robust service catalog, B2B interfaces and full hybrid cloud infrastructure management and operations. The benefits to the state includes increasing the operational efficiency by 35%, reductions in labor to support the infrastructure by allowing operations administration for all cloud models in one virtual location and reductions in operating costs. Throughout the term, Unisys implemented numerous projects that gained efficiencies worth \$240 million for CoPA, including mainframe and server consolidations, virtualization, storage area network (SAN) consolidations, and an option for variable service levels. We also implemented technologies such as network core switches, mass storage devices, mainframes with virtual hosts, SAN switches, and tape devices that reduced operating costs and improved functionality. Unisys continues to support CoPA’s data center needs and is helping CoPA on an agency wide data center modernization initiative that includes a new data center service delivery model that creates a flexible, reliable, secure, and robust IT structure. This model’s foundation is built on the concept of on-demand computing consumption for efficiency and cost savings as well as a broad service catalog.	

6.3 (M) Financials

Offeror must provide audited financial statements, of the last two years, to the State that demonstrate that an Offeror meets at a minimum Dun and Bradstreet (D&B) credit rating of 3A2 or better, or a recognized equivalent rating. Please provide the Respondent’s D&B Number and the composite credit rating. The State reserves the right to verify this information. If a branch or wholly owned subsidiary is bidding on this RFP, please provide the D&B Number and score for the parent company that will be financially responsible for performance of the agreement.

Unisys Response:

Unisys financial statements for the past 2 years are available at our web URL:
<http://www.unisys.com/investor-relations/financials-filings>.

Our D&B number is: 00-535-8932.

Unisys is providing the following as-recognized equivalent ratings:

- Moody's rating of B1
- S&P rating of B+.

6.4 (E) GENERAL INFORMATION

6.4.1 Provide any pertinent general information about the depth and breadth of your Solutions and their overall use and acceptance in the cloud marketplace.

Unisys Response:

Unisys and its partners Amazon, Microsoft and Google are the industry leaders in the cloud marketplace. Our combined strength and worldwide presence across all sectors of the public and the private industry is well known and accepted. We offer all types of cloud delivery models and service options. We are constantly expanding and enhancing our cloud models and delivery options. Section 6.2 showcases Unisys strong experience in the public sector cloud space. Additionally, the graphic 2015 Gartner Magic Quadrant for Cloud demonstrates the strength of our partners in this RFP.

Exhibit 6.4.1.1: 2015 Gartner Magic Quadrant for Cloud



6.4.2 Offeror must describe whether or not its auditing capabilities and reports are consistent with SAS 70 or later versions including, SSAE 16 6/2011, or greater.

Unisys Response:

Unisys is subject to client audit requirements according to SSAE 16, SOC 1, Type 2. Our existing staff have more than 10 years of experience, including the previous SAS 70 auditing guidelines, and work very closely with our external auditors at KPMG to create three SSAE 16 Type 2 reports covering at least 12 service centers. For auditing our client-facing IT environments, we follow ISO 27007 and ISO 27008. For auditing our financial aspects (related to SSAE 16 and otherwise), we follow GAAP under the AICPA's Code of Professional Ethics under Rule 203, Accounting Principles.

6.5 (E) BILLING AND PRICING PRACTICES

DO NOT INCLUDE YOUR PRICING CATALOG, as part of your response to this question.

6.5.1 Describe your billing and pricing practices, including how your billing practices are transparent and easy to understand for Purchasing Entity's.

Unisys Response:

Unisys follows a reliable and transparent approach with billing, pricing, and collections to enable our clients to run their businesses effectively. We will provide the Purchasing Entity with invoices for services provided in a timely way as and when the work is completed. Unisys is always open for support; billing-related queries can be forwarded directly to our representatives, who will work toward resolving issues in an appropriate method. We will send invoices to the Purchasing Entity at its designated address without fail. Unisys can also email a soft copy to a designated email address when required.

Unisys will provide access to a Consumption Reporting Solution that will enable the Purchasing Entities to track their cloud usage and charges. The Consumption Reporting Solution will provide a tabular view to monitor cloud usage in real time. A graphical view provides entities a comparison of their past usage with the current month that will help in forecasting their costs in their cloud consumption initiatives. The entities can download usage and billing reports in appropriate formats such as excel and pdf as and when required through this solution.

6.5.2 Identify any typical cost impacts that a Purchasing Entity might need to consider, if any, to implement your cloud solutions.

Unisys Response:

For the costs pertaining to billing, see our Cost Proposal.

6.5.3 Offeror must describe how its Solutions are NIST compliant, as defined in NIST Special Publication 800-145, with the service models it offers.

Unisys Response:

Unisys and our providers adhere to the NIST essential characteristics for Cloud Computing. For more information, refer to our response to Section 8.1.2.

6.6 (E) SCOPE AND VARIETY OF CLOUD SOLUTIONS

Specify the scope and variety of the Solutions you offer under this solicitation. You may provide a list of the different SaaS, IaaS, and/or PaaS services and deployment models that you offer.

Unisys Response:

Unisys is pleased to offer the requested cloud deployment models (private, community, public, and hybrid) and the underlying cloud services models (IaaS, PaaS, and SaaS) as part of our proposal. We will also provide architecture, design, implementation, and value-add services to support the various cloud solution options. Our cloud service provider partners in this proposal are Amazon Web Service (AWS), Microsoft, and Google. Exhibit 6.6.1 captures our offerings for this proposal.

Exhibit 6.6.1 Unisys Cloud Solution Offerings for NASPO.

	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Private Cloud	Unisys	Unisys	Unisys
Hybrid Cloud	Unisys	Unisys	Unisys
Cloud Professional Services	Unisys	Unisys	Unisys
Cloud Value-Add Services	Unisys	Unisys	Unisys

6.7 (E) BEST PRACTICES

Specify your policies and procedures in ensuring visibility, compliance, data security and threat protection for cloud-delivered services; include any implementations of encryption or tokenization to control access to sensitive data.

Unisys Response:

Section 8.6 and Section 8.8 of our proposal outlines specific details on compliance, data security and threat protection, including certifications and accreditations as they apply to this solicitation.

Unisys is pleased to provide the following, additional relevant examples of our additional past performance on access security and securing the cloud.

Brief Background Unisys Stealth

The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. Stealth can substitute traditional hardware topology for software-based cryptography. Stealth microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users (including database administrators).

- Unisys Stealth is an innovative, software-based security solution that does the following:
- Conceals end points, making them undetectable to unauthorized parties inside and outside the enterprise
- Tightens access control by focusing on user identity instead of physical devices so that security moves with the user and is easier to manage
- Protects sensitive data in motion from potential compromise through encryption
- Reduces costs by allowing clients to consolidate and virtualize networks, servers, and cloud architectures
- Supports regulatory compliance requirements (PCI) that draw resources away from a client's core business.

Unisys Stealth software has a legacy of being ready for government. Stealth received the National Information Assurance Partnership's (NIAP's) coveted Evaluation Assurance Level 4+ (EAL4+) certification, a Common Criteria international standard in effect since 1999. Unisys is pursuing the new standard of certification called Protection Profiles, as acknowledged by the U.S. Federal Government.

Relevant past performance of implementations of encryption or tokenization to control access to sensitive data:

State of New York: Unisys is implementing Stealth for Enterprise for the State of New York internal private cloud environment. Unisys has been engaged on the contract for the past 24 months, and the work is under way. We have helped New York to migrate and onboard about 35 of the state's 53 agencies to the multi-tenant data center in Albany running on Stealth.

Amazon: Unisys Stealth Cloud is an offering on Amazon Web Services (AWS) available directly at the AWS Marketplace. Stealth Cloud protects data in motion and isolates end-point workloads in a multitenant environment.

Commonwealth of Pennsylvania: as already referenced above, Unisys designed, integrated, and implemented several of Pennsylvania's own security and access control technologies into the Commonwealth's centralized data center.

5.0 ORGANIZATION PROFILE (RFP SECTION 7)

Section Title: Organization Profile: This section should constitute the Offeror's response to the items described in Section 7 of the RFP. An Offeror's response must be a specific point-by-point response, in the order listed, to each requirement in the Section 7 of the RFP.

7 ORGANIZATION AND STAFFING

7.1 (ME) Contract Manager

The Offeror must provide a Contract Manager as the single point of contact for management of the NASPO ValuePoint Master Agreement, administered by the State of Utah. The Contract Manager must have experience managing contracts for cloud solutions.

Harsh Bajpai, Unisys Director of Cloud Solutions and Services, is the current contract manager for the WSCA/NASPO/ValuePoint Cloud Hosting Contract with the State of Utah. If Unisys is awarded this contract, he will continue to manage it. For information on his qualifications, see Section 7.1.2.

7.1.1 Provide the name, phone number, email address, and work hours of the person who will act as Contract Manager if you are awarded a Master Agreement.

Mr. Harsh Bajpai

Unisys Corporation

Director, Cloud Solutions and Services

Phone Number: 703-439-6200

Email Address: Harsh.Bajpai@unisys.com

Work Hours: 8:00 a.m. to 8:00 p.m. Eastern Standard Time

7.1.2 Describe in detail the Contract Manager's experience managing contracts of similar size and scope to the one that will be awarded from this RFP. Provide a detailed resume for the Contract Manager.

Mr. Bajpai has supported the following government contracts for Unisys:

- WSCA/NASPO/ValuePoint: 2012 to present; expires in 2017
- SSA ITSSC: 2012 to 2014
- NASA SWEP IV&V: 2010 to 2015
- FEMA COMMIT: 2010 to 2011
- DHS EAGLE: 2009 to 2011.

Unisys provides the following resume for Mr. Bajpai.

Profile Summary:

Mr. Bajpai is an information technology professional with 15 years of experience in the following market segments:

- U.S. Federal and state-level public sector enterprises
- Global news and information service providers as well as media organizations
- Global investment banking.

Skills Summary:

Corporate Functions: Information Technology, Business Development, Sales, Sales Operations, Client Engagements and Implementations, and Marketing

Business: Transformation and Alignment, Strategy Formulation and Execution, Professional Services Management, Program Management, Outsourcing and Offshoring, Process Improvement and Automation, Governance, Practice Development, and Product and Services Management

Technology: Technology Infrastructure and Data Center Transformation and Optimization (DCTO), Virtualization, Cloud & Managed Services, Program Management, Information Security and Compliance, Software Development Lifecycle (SDLC), and End User Computing Platforms

Professional Experience:

08/2008 – Present	Director	Unisys Corporation
Market Space: USDHS (HQ, USCIS, USCG, FEMA, NPPD and CBP), GSA, U.S. Treasury, USPTO, USDA, SSA, FDIC, DHA (MHS & VA), and USSOCOM		
Role Summary - Develop and manage delivery of technology product solutions to the U.S. Federal markets. Support an annual revenue target of \$44 million. - Manage the Unisys portion of high-value services for existing strategic accounts. Provide SME-based technology consultation and integration services. Average annual and total task order value is \$120 million. - Provide SME-based technology consultation and integration services for U.S. Federal, state, local, and key global accounts. - Establish and maintain channel partnerships, vendor relationships, and client relationships across the GS, SES, and executive levels.		
Key Achievements - Key active member of the team that won the following contracts: COMMIT, ITSSC, TASPO, SWEP IV and 5, BASICS, Eagle 1, Networks, GSA EEAS, and WSCA/NASPO/ValuePoint. Post award of these contracts, program managed and executed task orders under these large federal contract vehicles. - Developed the Data Center Transformation and Outsourcing Practice (DCTO) for the Unisys Federal and Commercial markets. This framework today is the Unisys DCTO methodology. - Developed and delivered the Data Center Migration strategy for the Social Security Administration (SSA) under the ITSSC contract. - Developed and delivered a secure telework solution framework for the U.S. Coast Guard. This engagement realized the first reference implementation for the Unisys Stealth platform in the U.S. Federal market. - Developed and delivered the <i>Forward!</i> by Unisys solution to FDIC. This engagement realized the first reference implementation for the Unisys Stealth solution in the U.S. Federal market. - Initiated and won a \$30 million, 3-year contract with a large northeastern state for data center consolidation, modernization, security, and managed security services. - Won the USSOCOM medallion from GDIT for leading and representing Unisys on a very large IDIQ procurement. - Developed and delivered the first Unisys private cloud solution to HSBC Bank in Hong Kong. This solution realized the first reference implementation for the Unisys Private Cloud offering. - Strategic team member behind the \$100 million SITA win for AirCore application and baggage claims systems. - As the manager of the WSCA/NASPO/ValuePoint contract, generated \$5 million in Unisys sales revenue under the public cloud hosting and services vertical in the State and Local		

08/2008 – Present	Director	Unisys Corporation
market segment.		
04/2006 – 08/2008	Senior Director, VP	Thomson Reuters
Position Summary & Functional Responsibilities • Management of the global infrastructure, application support, and I-Dev teams • End-to-end responsibility for the primary and secondary data centers of Reuters Media • Outsourcing and offshoring • Global infrastructure availability • Infrastructure platform and service delivery as well as project management and compliance functions for Reuters Media.		
Key Achievements • Outsourced the Reuters Media primary data center to a managed service provider on a utility computing platform. • Program managed the migration of the Reuters Media platform from .NET to the LAMJ stack. • Outsourced the QA and development infrastructure to a hosting provider. • Built the software development facility of Reuters Media in Beijing, China. • Outsourced the software development production support function of Reuters Media to Ukraine. • Built the globally load-balanced, live/live DR/COOP facility for Reuters Media in the United Kingdom with a managed service provider.		
01/2001 – 03/2006	Various roles – see below	Dow Jones & Co. Inc.
Position Summary & Functional Responsibilities 2004 – 2006: E-Commerce Product Manager for WSJ.com and Barrons.com 2003 – 2004: Business Information Security and Compliance Officer 2001 – 2003: Information Security Architect		
Key Achievements E-Commerce Product Manager for WSJ.com and Barrons.com • Delivered an e-commerce platform for the \$10 million WSJ.com and LexisNexis partnership. • Delivered the e-commerce platform for WSJ.com and <i>The Wall Street Journal</i> combined product offering in Asia and Europe. • Owned the WSJ.com and Barrons.com e-commerce platform. Business Information Security and Compliance Officer • Was accountable for the Sarbanes-Oxley program for all external business units of Dow Jones & Co. • As a technology SME for the core M&A team of Dow Jones & Co., executed M&A activities for ONI, Factiva.com, and VentureWire. Information Security Architect • Oversaw information security and assurance of Factiva.com, a Reuters and Dow Jones		

01/2001 – 03/2006	Various roles – see below	Dow Jones & Co. Inc.
Company. - Delivered forensic analysis for the legal department and external law enforcement agencies. - Introduced and deployed Open Source production-class tools such as intrusion detection, vulnerability analysis, and configuration management.		

10/1998 – 01/2001	Technical Architect	SnS Medical Systems
- Assumed end-to-end architectural and operational ownership of the hosted software billing system. - Assumed service delivery ownership of client-owned and -managed systems. - Owned the software service delivery and security for the enterprise. - Managed deployment and field service consultants.		

08/1995 – 10/1998	General Manager, Operations	Accor Inc.
- Canvassed and presented quarterly competitive market analysis reports for the regional district and for the 15 properties in the district. - Successfully turned around operationally challenged and P&L-stagnant locations across NJ, PA, and CT. - Designed and Implemented operational metrics and reporting for the regional district. - Oversaw annual revenue of \$2 million.		

Education:

MBA, Business Administration, University of Maryland University College, Adelphi, MD, 2005

M.S., e-Commerce, University of Maryland University College, Adelphi, MD, 2003

B.A., Marketing & Finance, California State University, Fullerton, Fullerton, CA, 1995

A.A., Social & Behavioral Sciences, Citrus College, Glendora, CA, 1993

A.S., Business Administration, Citrus College, Glendora, CA, 1993

Certifications, Affiliations, and other Professional Work Experience

CISSP, MCSE, MSCE+I, and MCT

Middlesex County Court, Juvenile Conference Committee Member, 2004 – 2006

Middlesex County College, Adjunct Professor, Department of Computer Science, 2004 – 2006

Princeton University, Guest Lecturer, 03/2013

Baltimore National Aquarium, Volunteer Diver, 06/2013 – Present

7.1.3 Describe in detail the roles and responsibilities of the Contract Manager as they apply to the NASPO ValuePoint Master Agreement that will be awarded from this RFP.

Harsh Bajpai is our current contract manager for the WSCA/NASPO/ValuePoint Cloud Hosting Contract with the State of Utah. If Unisys is awarded this contract, he will continue to manage it and serve as the single point of contact to Utah and the Participating States.

Mr. Bajpai will continue to facilitate onboarding of new customers and respond to questions and concerns about new and existing procurements. He also will continue to coordinate customer enablement through our partners. He will provide the collaboration and support required from our partners in this proposal.

The following functions are required to support a contract of this scope and size. The current WSCA/NASPO/ValuePoint contract has talented individuals who support the functions stated below. If Unisys is awarded this contract, we are committed to provide this support structure.

Contracts Desk: A team of talented individuals supports the day-to-day operations of current and new contracts. The Contract Operations Specialist supports Mr. Bajpai in writing statements of work (SOWs); resolving billing-related issues; managing the calculation and disbursement of fees, credits, etc.; and overseeing terms and conditions as required by this solicitation. If Unisys is awarded this contract, we will assign a dedicated Contract Operations Specialist to it.

Attorney: If Unisys is awarded this contract, we will dedicate an attorney to support Mr. Bajpai on required legal matters.

Enablement, Billing, and Reporting: If Unisys is awarded this contract, we will dedicate a team of billing and reporting staff to Mr. Bajpai. This team will provide customer enablement with our various partners on this contract. This team will also provide the billing details and reports required to manage and execute this contract.

Certified Architects: Unisys has geographically positioned certified architects in and near the states that are our current clients. These architects can deliver full end-to-end solutions. They also engage frequently with our clients as consultants and advisors. If Unisys is awarded this contract, we will continue to support and expand this function. These architects will support Mr. Bajpai.

Client Executives: Unisys has geographically positioned Client Executives in and near the states that are our current clients. Our Client Executives are account managers dedicated to a given state. A Client Executive's main function is to provide our clients with onsite presence and support. If Unisys is awarded this contract, we will continue to support and expand this function. These Client Executives will support Mr. Bajpai.

Subject Matter Experts: Unisys maintains a very strong partnership with our partners in this proposal. Formal partnerships require us to maintain a certain number of certified personnel to fulfill this partnership's terms and conditions. Unisys will continue to strengthen our bench based on the required demands. Additionally, our partnerships give us direct access to technical and nontechnical resources in our partner organizations. Unisys bench staff and our partners' resources support Mr. Bajpai under the current WSCA/NASPO/ValuePoint contract. If Unisys is awarded this contract, we will continue to support and expand this resource base to support Mr. Bajpai.

6.0 TECHNICAL RESPONSE – AWS – MS INCLUDED (RFP SECTION 8)

Section Title: Technical Response. This section should constitute the Technical response of the proposal and must contain at least the following information:

A complete narrative of the Offeror's assessment of the Cloud Solutions to be provided, the Offerors ability and approach, and the resources necessary to fulfill the requirements. This should demonstrate the Offeror's understanding of the desired overall performance expectations and clearly indicate any options or alternatives proposed.

A specific point-by-point response, in the order listed, to each requirement in the Section 8 of the RFP. Offerors should not provide links to a website as part of its response.

Offeror's should focus their proposals on the technical qualifications and capabilities described in the RFP. Offerors should not include sales brochures as part of their response.

8 TECHNICAL REQUIREMENTS

If applicable to an Offeror's Solution, an Offeror must provide a point by point responses to each technical requirement demonstrating its technical capabilities. If a technical requirement is not applicable to an Offeror's Solution then the Offeror must explain why the technical requirement is not applicable.

If an Offeror's proposal contains more than one Solution (i.e., SaaS and PaaS) then the Offeror must provide a response for each Solution. However, Offerors do not need to submit a proposal for each Solution.

8.1 (M)(E) TECHNICAL REQUIREMENTS

8.1.1 Offeror must identify the cloud service model(s) and deployment model(s) it intends to provide to Eligible Users. See Attachment D.

Unisys Response:

Unisys has the ability and the qualifications to provide the four types of cloud deployment methods (Public, Private, Community, and Hybrid) defined in Attachment D. Unisys also has the ability and the qualifications to provide the three types of cloud service models (IaaS, PaaS, and SaaS) defined in Attachment D.

To provide this wide and deep cloud offering catalog, Unisys partnered with Amazon Web Services (AWS), Microsoft, and Google for this proposal. Unisys believes that offering a selection of world-class, best-of-breed cloud providers will provide the Participating Entities with the needed flexibility, agility, and technological options in a competitive environment. We also believe that offering these solutions and services from a cloud-neutral systems integrator like Unisys minimizes risk, promotes accountability, and improves efficiency. With this approach, the Participating Entities can choose their cloud solutions from a competitive set of vendors that are backed by a single, experienced cloud-neutral systems integrator. Exhibit 1 summarizes the Unisys offering for this proposal.

Exhibit 1: Unisys Cloud Solutions Offerings to NASPO.

	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft Google
Private Cloud	Unisys	Unisys	Unisys
Hybrid Cloud	Unisys	Unisys	Unisys
Cloud Professional Services	Unisys	Unisys	Unisys
Cloud Value Add Services	Unisys	Unisys	Unisys

Our partners in this proposal provide qualified offerings in the public and the community (government-specific) cloud deployment space. Our partners also provide the three cloud service (IaaS, PaaS, and SaaS) models for the U.S. public sector. Unisys is a recognized leader in providing hybrid and private cloud deployment methods in the public sector. We continue to build more private and hybrid cloud environments for federal, state, and local governments. Unisys is also a recognized leader in SaaS-based applications services platforms. Our main presence is in the Health and Human Services, Public Safety and Justice, Law Enforcement, Border Security, Digital Government, and Citizen Services markets.

As the incumbent, Unisys is fully aware of what is needed to fulfill the requirements and the overall performance expectations necessary for this contract vehicle. We believe that our end to end cloud expertise and experience along with strong partnerships will provide unique capabilities for our government clients.

8.1.2 For the purposes of the RFP, meeting the NIST essential characteristics is a primary concern. As such, describe how your proposed solution(s) meet the following characteristics, as defined in NIST Special Publication 800-145:

Refer to the point-by-point responses below.

8.1.2.1 NIST Characteristic - On-Demand Self-Service: Provide a brief written description of how the cloud solutions proposed satisfies this individual essential NIST Characteristic. Attest capability and briefly describe how self-service technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of On-demand Self Service, one of the 5 essential characteristics of the NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Examples of these technologies include those from vendors such as BMC, Microsoft, VMware, HP, ServiceNow, Open Source Foundation, IBM, and Dell. Unisys has leveraged combinations of these technologies to create customized, automated self-service portals that can manage a workload's entire life cycle in an internal, private cloud environment or in a hybrid environment. Unisys has relevant past performance in this space in the U.S. Federal, state, and local markets. We provide some of these past performance details in Section 4.0, Business Profile.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Amazon Web Services (AWS) provides customers with on-demand access to a wide range of cloud infrastructure services, charging them only for the resources they actually use. AWS will enable the State to eliminate the need for costly hardware and the administrative pain that goes along with owning and operating it. Instead of the weeks and months it takes to plan, budget, procure, set up, deploy, operate, and hire for a new project, AWS customers can simply sign up for AWS and immediately begin deployment in the cloud with the equivalent of 1, 10, 100, or 1,000 servers. Whether an organization needs to prototype an application or host a production solution, AWS makes it simple for customers to get started and be productive.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Windows Azure is an internet-scale, high-availability cloud fabric operating on globally-distributed Microsoft data centers. Windows Azure and related tools support the development and deployment of applications into a hosted environment that extends the on-premises data center. On-demand self-service refers to the service provided by Microsoft Azure that enables the provision of resources on demand whenever required. In on-demand services can be enabled from the HTML Portal, using Azure API, using CLI for Mac, Linux, and Windows with Azure Service Management. Azure on-demand self-service resource sourcing programmatically is a prime feature allowing the user to scale the infrastructure. Microsoft provides a full-featured portal that enables customers to order the hundreds of cloud services available from Azure. Azure can scale from one to thousands of virtual machine instances with the benefit of built-in virtual networking and load balancing. Microsoft bills customers per minute, allowing them to pay only for what they use.

Google Public and Community Cloud for SaaS: Google Apps administrators have access to an Admin Panel that allows them to perform user administration, service configurations, reporting on demand. Additional tools that support this characteristic include the APIs in the Admin SDK that can automate many common tasks as well as an extensive set of third-party tools developed by trusted independent software vendors.

8.1.2.2 NIST Characteristic - Broad Network Access: Provide a brief written description of how the cloud solutions proposed satisfies this individual essential NIST Characteristic. Attest capability and briefly describe how network access is provided.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Broad Network Access, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Vendor-based examples of technologies include those of vendors such as Cisco, Juniper, Nortel, Foundry, and F5. Technology-based examples include leased lines, Internet, VPN (SSL and IPsec), web applications, thick clients, thin clients, and mobile clients. Using these technologies, Unisys has provided our Federal, state, and local clients with low-latency, redundant broad network access methods.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS provides a simple way to access servers, storage, databases, and a broad set of application services over the Internet. Cloud computing providers such as AWS own and maintain the network-connected hardware required for these application services; customers provision and use what they need from a web application, a mobile client, or programmatically at published and well-documented APIs.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Azure allows customers to access their Web, Mobile, IaaS, SaaS, and PaaS platforms easily. Network access to the cloud is available over VPN or dedicated high-speed connections between the customer and the Microsoft cloud network. Connectivity between Microsoft data centers is delivered by one of the largest fiber networks in the world, establishing high levels of performance and security. Microsoft Azure enables the enterprise to create an on-premises network route from on-premises VPN device and the Azure virtual network. Configure on-premises hardware or software VPN device to terminate the VPN tunnel, which uses Internet Protocol security (IPsec). Gateway connects on-premises to Microsoft Azure through many connection bandwidth using software only or hardware based connectivity. Basic VPN connection, Standard VPN, ExpressRoute gateway connections.

Google Public and Community Cloud SaaS: Google's network delivers speeds of 1 petabyte per second of total bidirectional bandwidth, has dozens of network points of presence, and offers a clear and mature peering/content delivery policy so that customers are connected to Google's infrastructure, quickly negating the need for high-cost dedicated access lines. Google Apps (SaaS) is device, platform, and network neutral. As a pure-play cloud-based solution, the services can be accessed from modern browsers from a network access point. Customers that have requirements to limit access to trusted devices or trusted access points have access to Device Management tools from the Admin Panel at no additional cost.

8.1.2.3 NIST Characteristic - Resource Pooling: Provide a brief written description of how the cloud solutions proposed satisfies this individual essential NIST Characteristic. Attest capability and briefly describe how resource pooling technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Resource Pooling, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Technology-based examples include Unisys Stealth, VMware, Hyper-V, and Xen-based virtualization; Unisys ClearPath Forward, a hardware partitioning

utility compute platform; storage systems from EMC and NetApp; etc. Using these technologies, we created shared services environments for various Federal and state entities. Examples include the Internal Revenue Service; the Commonwealth of Pennsylvania; and the States of New York, California, and Michigan. Section 4.0, Business Profile contains the solution sets that Unisys has created.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: The AWS environment is a virtualized, multitenant environment. AWS has implemented security management processes, PCI controls, and other security controls designed to isolate each customer from other customers. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture has been validated by an independent PCI Qualified Security Assessor (QSA) and was found to comply with the requirements of PCI DSS version 2.0, published in October 2010.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Azure resource pooling supports scalable systems involved in cloud computing and software as a service (SaaS) and enable near-infinite growth with immediate availability. The kinds of services that can apply to a resource pooling strategy include data storage services, processing services, bandwidth provided services, and other Azure-related compute elasticity.

Google Public and Community Cloud for SaaS: Google's global infrastructure is a shared pool of resources that dynamically serve each end user with a primary data center access point that may rotate throughout the session without the end user being aware and at the same time replicate data across at least two additional geographically dispersed data centers that also may rotate through the session.

8.1.2.4 NIST Characteristic - *Rapid Elasticity*: Provide a brief written description of how the cloud solutions proposed satisfies this NIST Characteristic. Attest capability and briefly describe how rapid elasticity technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Rapid Elasticity, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Examples of these technologies include those from vendors such as Microsoft, Amazon, VMware, Dell, Chef and Puppet. Unisys has leveraged combinations of these technologies to create customized deployment configurations with pre-defined business rules enabling auto-scaling that can seamlessly scale resources up or down by provisioning and de-provisioning resources automatically as the workloads change. The resources shall in turn be handled by a load balancer to distribute the workloads to different resources in an internal, private cloud environment or in a hybrid environment.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS provides a massive global cloud infrastructure that will allow the State to quickly innovate, experiment, and iterate. Instead of waiting weeks or months for hardware, the State can instantly deploy new applications, instantly scale up as its workload grows, and instantly scale down according to demand. Customers need to be confident that their existing infrastructure can handle a spike in traffic and that the spike will not interfere with normal business operations. Elastic Load Balancing and Auto Scaling can automatically scale a customer's AWS resources up to meet unexpected demand and scale those resources down as demand decreases.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Azure supports rapid elasticity allowing automated requests for additional resources (i.e. compute, disk space, connectivity and other types of services). The State can automatically scale resources according to a variety of metrics. The ability to scale resources up or down, combined with per minute billing, makes Azure the most efficient computing platform on the market. The automatic scaling feature at the Azure portal is very intuitive and

easy to use. This feature is used by thousands of customers needing the ability to turn compute resources up or down automatically.

Google Public and Community Cloud for SaaS: Google Apps (SaaS) can scale from one user to tens of thousands of users. The use of the services included is designed to scale horizontally in the per user storage allocation. Google Apps can be purchased with an allocation of 30 GB per user or with unlimited storage.

8.1.2.5 NIST Characteristic - Measured Service: Provide a brief written description of how the cloud solutions proposed satisfies this NIST Characteristic. Attest capability and briefly describe how measured service technical capability is met.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has custom solutions that meet the requirement of Measured Service, one of the 5 essential characteristics of NIST definition of Cloud. As a systems integrator, Unisys has integrated many hardware and software solution sets and has designed, implemented and managed various solutions to meet the requirements of our clients and the National Institute for Standards and Technology (NIST). Examples include VantagePoint, a solution offered by Unisys that we implemented at the Commonwealth of Pennsylvania. For the IRS, we provide a customized private storage environment that is based on a measurable SLA down to the data types, which maps to storage offerings categorized into Platinum, Gold, Silver, and Bronze Storage.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS uses automated monitoring systems to provide a high service performance and availability. Proactive monitoring is available from a variety of online tools for internal and external use. Systems in AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so that personnel are always available to respond to operational issues. This schedule includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft Azure supports NIST's principal areas of measured services such as measured service setup, which allows Azure to control a system, user, or tenant's usage of resources with metering capability. Azure supports automated remote services measurement tools to provide auditing and accountability for utilization. Azure-measured service confirms that even when there is no specific interaction for a service change, service change is still audited to support billing cycles.

Google Public and Community Cloud for SaaS: Google Apps (SaaS) is already optimized for unlimited use by end users. Google maintains high availability, low latency, and fault tolerance as a part of its contracted services. No metering or rate limiting will be required because Google does not charge for bandwidth use

8.1.3 Offeror must identify for each Solution the subcategories that it offers for each service model. For example if an Offeror provides a SaaS offering then it should be divided into education SaaS offerings, e-procurement SaaS offerings, information SaaS offering, etc.

Unisys Private and Hybrid Cloud for IaaS and PaaS: These are custom solutions that Unisys can design, implement, and manage for clients according to their needs. As a systems integrator, Unisys has integrated many hardware and software solution sets to meet the requirements of our clients and NIST. Our custom examples include private storage as a service for the IRS, hybrid storage as a service for the Commonwealth of Pennsylvania, and private cloud for the U.S. Department of the Interior.


Unisys Public and Community, Private, and Hybrid Cloud for SaaS: Unisys is a leader in providing SaaS for the Health and Human Services, Public Safety and Justice, Law Enforcement, Border Security, Digital Government, and Citizen Services markets. We provide custom integrated applications that can be hosted in the public and community cloud environments as an IaaS or a PaaS; or these solutions can be hosted in a hybrid or a private cloud. A recent example includes 311 Service for the City of Philadelphia. The 311 application is a Unisys application that is hosted on Salesforce.com. Unisys customized and

integrated the application for the City of Philadelphia and delivered it on SalesForce.com on a consumption basis. **Appendix 11 – Unisys-1** contains the definition summary of each cloud offering provided by Unisys.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Exhibit 2 illustrates each service models offered by Amazon in the public and community cloud deployment methods. In the point-by-point responses, below as required by this RFP, we have defined each service that is illustrated in Exhibit 2. **Appendix 12 – AWS-1** contains the definition summary of each one of the services in Exhibit 2 in a consolidated table.

Exhibit 2.AWS Offerings.

Aw



UT003

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

IaaS is most like traditional IT delivery. Customers provision their own virtual machines, define their own networks, and allocate their own virtual hard disks. IaaS shifts the burden of operating datacenters, virtualization hosts, and hypervisors. In addition, the business continuity and disaster recovery infrastructure is shifted from the enterprise to the service provider.

Deploying an application and managing an IaaS environment provides the most flexibility that Azure has to offer. With any deployment choice, there will be pros and cons that must be considered. The greatest benefit of an IaaS implementation is that it offers the greatest amount of control from the operating system to manage access to the application.

IaaS Services by Microsoft includes:

Systems Center Virtual Machine Manager: System Center Virtual Machine Manager provides centralized administration and management of a virtual environment.

Systems Center Operations Manager: System Center Operations Manager 2007 provides end-to-end monitoring for the enterprise IT environment.

Systems Center Configuration Manager: Systems Center Configuration Manager 2007 provides a comprehensive solution for change and configuration management for the Microsoft platform.

PaaS extends IaaS further by providing multitenant services that customers subscribe to. Platform services are a transformational computing model that can dramatically reduce the costs and increase the agility of delivering applications to end users internally and externally. PaaS users bring their own application code but leverage robust platforms, which they do not need to maintain.

With PaaS applications, many of the layers of management are removed and more flexibility is provided than an application running on IaaS instances. Specifically, there is no need to manage the operating system, including patching, which reduces some of the complexity of designing the deployment.

A significant benefit of deploying an application running in a PaaS environment is the ability to quickly and automatically scale up the application to meet the demand when traffic is high, and inversely scale down when the demand is less. Deploying an application in the PaaS model is very cost effective from a scalability and manageability perspective.

PaaS Services by Microsoft includes:

Microsoft SQL Azure Database: a fully relational cloud database solution built on SQL Server that offers highly available, scalable, multi-tenant database services.

Windows Azure: delivers on-demand compute and storage to host, scale, and manage web applications through Microsoft data centers.

SaaS is the real promise of cloud computing. By integrating applications from one or multiple vendors, customers need to bring only their data and configurations. They can eliminate the costs of building and maintaining applications and platform services and still deliver the secure, robust solutions to the end users.

Choosing an Azure SaaS offering provides the least amount of responsibility on the customer's side. At the same time, providing a lesser amount of flexibility in comparison with an IaaS or PaaS approach.

SaaS Services by Microsoft includes:

Office 365: a group of software plus services subscriptions that provides productivity software and related services to its subscribers. Office 365 offers plans that include e-mail and social networking services through hosted versions of Exchange Server, Skype for Business Server, SharePoint, and Office Online integration with Yammer, as well as access to Office software.

Microsoft Dynamics: with minimal configuration, offers constituent relationship management (CRM) and other extended CRM solutions to help automate workflow and centralize information.

Microsoft Exchange Online: delivers email with protection, plus calendar and contacts.


Microsoft Live Meeting: delivers hosted web conferencing.

Microsoft SharePoint: creates a highly secure, central location for collaboration, content, and workflow.

Microsoft Communications Online: provides real-time, person-to-person communication with text, voice, and video across agencies.


Many scenarios need to implement a blend of Azure offerings to meet the needs of their organization and application requirements. The following diagram (Exhibit 3) highlights the main differences from a manageability perspective, when using public cloud SaaS, PaaS, IaaS and On-Premises implementations

Exhibit 3: Microsoft IaaS/PaaS/SaaS differentiation



Google Public and Community Cloud for SaaS: Exhibit 4 and Exhibit 5 illustrate each service model offered Google in the public and community cloud deployment models. **Appendix 13 – Google-1** contains the definition summary of each services in these exhibits in a consolidated table.

Exhibit 4. Google SaaS.



8.1.4 As applicable to an Offeror's proposal, Offeror must describe its willingness to comply with, the requirements of Attachments C & D.

Unisys is willing to comply with the 'NIST Service Models' and 'Scope of Services'.

8.1.5 As applicable to an Offeror's proposal, Offeror must describe how its offerings adhere to the services, definitions, and deployment models identified in the Scope of Services, in Attachment D.

Unisys and our partners for this proposal have a proven track record in providing, building, deploying, and managing public, community, private, and hybrid cloud environments that deliver IaaS, PaaS, and SaaS services. We independently and jointly delivered and built these environments for U.S. Federal, state, and local governments that also follow the requirements defined in the State provided Attachment D. Our collective and individual past performances are stated in Section 4.0, Business Profile.

Additionally, our certifications and accreditations that are listed throughout this proposal attest to our understanding and our adherence to the services, definitions, and deployment models identified in the State provided Attachment D.

8.2 (E) SUBCONTRACTORS

8.2.1 Offerors must explain whether they intend to provide all cloud solutions directly or through the use of Subcontractors. Higher points may be earned by providing all services directly or by providing details of highly qualified Subcontractors; lower scores may be earned for failure to provide detailed plans for providing services or failure to provide detail regarding specific Subcontractors. Any Subcontractor that an Offeror chooses to use in fulfilling the requirements of the RFP must also meet all Administrative, Business and Technical Requirements of the RFP, as applicable to the Solutions provided. Subcontractors do not need to comply with Section 6.3.

8.2.2 Offeror must describe the extent to which it intends to use subcontractors to perform contract requirements. Include each position providing service and provide a detailed description of how the subcontractors are anticipated to be involved under the Master Agreement.

Unisys does not intend to use subcontractors. We intend to use our highly qualified, world-recognized, strategic, and strong industry partners (Amazon, Microsoft, and Google) to provide some of the cloud solutions detailed in our proposal. Unisys will resell our partners' cloud models and services using our agreements with our partners for the delivery of IaaS, PaaS, and SaaS from public and community clouds.

Unisys will directly provide private and hybrid cloud models that will be capable of delivering IaaS, PaaS, and SaaS. Additionally, we will offer free professional services and value-add services.

The Master Agreement will be an agreement between the State and Unisys.

8.2.3 If the subcontractor is known, provide the qualifications of the subcontractor to provide the services; if not, describe how you will guarantee selection of a subcontractor that meets the experience requirements of the RFP. Include a description of how the Offeror will ensure that all subcontractors and their employees will meet all Statement of Work requirements.

Not applicable based on the response above. Unisys is not using subcontractors. We are using our partners. Unisys has partnership agreements in place with the partners in this proposal.

8.3 (E) WORKING WITH PURCHASING ENTITIES

8.3.1 Offeror must describe how it will work with Purchasing Entities before, during, and after a Data Breach, as defined in the Attachments and Exhibits. Include information such as:

- **Personnel who will be involved at various stages, include detail on how the Contract Manager in Section 7 will be involved;**
- **Response times;**
- **Processes and timelines;**
- **Methods of communication and assistance; and**
- **Other information vital to understanding the service you provide.**

Unisys follows a Security Incident Response Process that details the routine handling of intentional or inadvertent information security events affecting the integrity, confidentiality, authentication, non-repudiation, and availability of information, and the information technology infrastructure of Unisys. The Unisys representative will be our Contract Manager, Harsh Bajpai. He will be supported by a core incident response team that includes representatives from our IT/IS, legal, communications, program management, and partner organizations as well as subject matter experts from various areas of technology as needed. This team from Unisys and our partners will closely coordinate with the State's incident response teams to contain, assess, mitigate, and prevent incidents. Once engaged for a confirmed breach, the members of our core incident response team will remain with the incident until it is resolved to the State's satisfaction.

The process of identification will be a joint responsibility between the State as well as Unisys and our partners because incidents can be detected by the parties involved. If the State detects an incident, the State can contact our Contract Manager using the information in Section 5.0, Organization Profile. If Unisys or our partners detect an incident, we will contact the State. Our Contract Manager will assemble the required incident response team, establish an internal and an external communications plan, and summon an investigation team of technology, legal, public relations, and communications experts. Unisys will work with the State to provide required response times based on the incident's severity. We will cooperate with the involved parties from Federal, State, and local authorities under the State's guidance and Federal law.

8.3.2 Offeror must describe how it will not engage in nor permit its agents to push adware, software, or marketing not explicitly authorized by the Participating Entity or the Master Agreement.

Unisys does not push adware, software, or marketing that is not explicitly authorized by the Participating Entity or the Master Agreement.

8.3.3 Offeror must describe whether its application-hosting environments support a user test/staging environment that is identical to production.

In the public and the community IaaS cloud space for application hosting, clients create, manage, and develop solutions in their own environment. These environments are created in accordance with their functional requirements and specifications. Vendors provide these environments that are bound by SLAs. Vendors will provide test and development environments for PaaS and SaaS delivery models. For those environments, the underlying infrastructure, features, functionality, and specifications are also bound by the governing SLAs. In the hybrid and private cloud models, IaaS, PaaS, and SaaS environments are built to client requirements and the governing SLAs.

8.3.4 Offeror must describe whether or not its computer applications and Web sites are be accessible to people with disabilities, and must comply with Participating entity accessibility policies and the Americans with Disability Act, as applicable.

The solutions in this proposal are accessible to people with disabilities and complies with Participating Entity accessibility policies and the Americans with Disabilities Act.

8.3.5 Offeror must describe whether or not its applications and content delivered through Web browsers are be accessible using current released versions of multiple browser platforms (such as Internet Explorer, Firefox, Chrome, and Safari) at minimum.

Unisys and our partners support current versions of Internet Explorer, Firefox, Safari, and Google Chrome. Unisys and our partners constantly track and move with the industry when it comes to browser compatibilities and browser technologies in the marketplace. This also includes browsers that run on computers, laptops, tablets, mobile devices, etc.

8.3.6 Offeror must describe how it will, prior to the execution of a Service Level Agreement, meet with the Purchasing Entity and cooperate and hold a meeting to determine whether any sensitive or personal information will be stored or used by the Offeror that is subject to any law, rule or regulation providing for specific compliance obligations.

As an incumbent that has fulfilled the requirements for nine Participating Entities under the current contract, Unisys takes our interaction with the Purchasing Entities seriously and continuously strives to enhance this experience. A purchasing entity usually contacts Unisys directly or through a third party. Upon mutual determination, our Contract Manager, Harsh Bajpai, and a Unisys Client Executive conduct an initial client call with the Purchasing Entity. This call focuses mainly on technical and functional requirements as well as regulatory compliance requirements and obligations. If a follow-up, deep-dive call is required, Unisys provides the appropriate certified subject matter experts (SMEs). At this stage, the Purchasing Entity brings in its SMEs. If further discussions are required, Unisys has designated staff, as described in Section 5.0, Organization Profile. Unisys and our SME resources understand the workload and the governing requirements and can suggest an appropriate cloud destination to the Participating Entity.

8.3.7 Offeror must describe any project schedule plans or work plans that Offerors use in implementing their Solutions with customers. Offerors should include timelines for developing, testing, and implementing Solutions for customers.

Public and Community Cloud Offerings for IaaS, PaaS, and SaaS from Amazon, Microsoft and Google: Unisys will provide a Statement of Work (SOW) within 5 business days of receiving the firm intent to purchase from the customer. Upon receiving the signed copy of the SOW, we will provision the account for the customer and provide login details within 5 business days.

Private and Hybrid Cloud Offerings for IaaS, PaaS, and SaaS from Unisys: Because these are customized solutions of varying sizes and requirements, Unisys is providing example transition, design, build, and migration times for our Federal and State government clients. The example client list for reference includes USDA, IRS, and Commonwealth of Pennsylvania. They each have a footprint of 1,000 or more servers. Our typical transition time for their environments is 30, 60, or 90 days. We can design and build a new environment in 3 to 6 months. Workload migration depends on the participation and the prioritization from the client, workload complexity, mission criticality, current redundancies, and migration provisions in place. Therefore, it is difficult to estimate these migrations. Although Unisys has

migrated some of our smaller state-level clients in 3 to 6 months, some of our engagements with larger Federal and state-level clients have lasted up to 24 months.

8.4 (E) CUSTOMER SERVICE

8.4.1 Offeror must describe how it ensure excellent customer service is provided to Purchasing Entities. Include:

- **Quality assurance measures;**
- **Escalation plan for addressing problems and/or complaints; and**
- **Service Level Agreement (SLA).**

Problem Reporting and Escalation

Unisys Technology and Support Services are available from Unisys to our clients 24x365. Our Technology and Support Services organization responds to a client's request for support. Usually, the client calls a central toll-free number or enters a service incident at our website. Before contacting Unisys, the client is required to consolidate and vet incidents and then document the suspected fault, specifying its severity and the impact the problem causes. This step is known as Level 1 Support, in which Unisys verifies the client's entitlement to support.

If a call or a website contact cannot be directly answered and resolved, our Level 1 Support team triages it and sends it to the appropriate Level Support 2 team that knows the specific target application. This team can be an engineering support team at Unisys or a third-party technology vendor with which Unisys partners for this solution. Our Level 2 Support team analyzes the incident and rejects problems not directly related to a fault in the application. Our Level 2 Support team may need to work with our Level 1 Support team to call the client several times to seek information to help isolate the problem or to develop a fix.

Once a fault is isolated, our Level 2 Support team submits a fault report to our development team (Unisys or the third-party technology partner). This is known as Support Level 3, which is responsible for the final fix. Our Level 3 development team then fixes and tests the fault. The fix is then made available to the client. To escalate an issue to the Unisys Support system at the same toll-free number, the client can ask to speak to a Duty Manager.

The SLAs are available in **Appendices 8 – 10**.

Appendix 8 – AWS SLA

See Appendix 8

Appendix 9 – Microsoft SLA

See Appendix 9

Appendix 10 – Google Apps SLA

See Appendix 10

8.4.2 Offeror must describe its ability to comply with the following customer service requirements:

a. You must have one lead representative for each entity that executes a Participating Addendum. Contact information shall be kept current.

If Unisys is awarded the contract, we will provide one lead representative for each entity that executes a Participating Addendum and keep the contact information current.

b. Customer Service Representative(s) must be available by phone or email at a minimum, from 7AM to 6PM on Monday through Sunday for the applicable time zones.

Unisys customer service and the customer service of our partners in this proposal operate a help desk 24x365.

c. Customer Service Representative will respond to inquiries within one business day.

The Unisys contract manager identified in Section 5.0, Organization Profile will respond to queries within 1 business day. Our standard customer service is available 24x365. Our technical support response times will be based on the support model purchased. For this proposal, the level of support that we propose comes with a response time of 60 minutes or shorter.

d. *You must provide design services for the applicable categories.*

Unisys is positioned to provide design services for the applicable categories. The Cloud Solutions section of our Cost Proposal lists the various services and their descriptions.

e. *You must provide Installation Services for the applicable categories.*

Unisys is positioned to installation services for the applicable categories. The Cloud Solutions section of our Cost Proposal lists the various services and their descriptions.

8.5 (E) SECURITY OF INFORMATION

8.5.1 *Offeror must describe the measures it takes to protect data. Include a description of the method by which you will hold, protect, and dispose of data following completion of any contract services.*

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

Unisys and our partners implemented a wide range of controls to protect the data. Each of us is certified and accredited to many industry and government standards. The responses to RFP requirements 8.6.1 through 8.6.6 list those standards. In general, data protection controls have been implemented at multiple layers across people, process, and technology. The responses to RFP requirements 8.6.1 through 8.6.6 provide these details.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS customers retain control and ownership of their data, and all data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers should consider the sensitivity of their data and decide if and how they will encrypt data while it is in transit and while it is at rest.

There are several options for encrypting data at rest, ranging from completely automated AWS encryption solutions to manual, client-side options. Choosing the right solutions depends on which AWS cloud services are being used and customer requirements for key management. Information on protecting data at rest using encryption can be found in the Protecting Data Using Encryption section of the Amazon Simple Storage Service (Amazon S3) Developer Guide available on the Internet.

For Securing Data in Transit services from AWS provide support for both Internet Protocol Security (IPSec) and Secure Sockets Layer/Transport Layer Security (SSL/TLS) for protection of data in transit. IPSec is a protocol that extends the IP protocol stack, often in network infrastructure, and allows applications on upper layers to communicate securely without modification. SSL/TLS, on the other hand, operates at the session layer, and while there are third-party SSL/TLS wrappers, it often requires support at the application layer as well. Unisys Stealth available as an AWS service provides end-to-end, IPSec based protection for data in-transit as well as workload isolation within the cloud and/or the private enterprise without the need for application modification.

It is important that customers understand some important basics regarding data ownership and management in the cloud shared responsibility model:

1. Customers continue to own their data.
2. Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
3. Customers can download or delete their data whenever they like.
4. Customers should consider the sensitivity of their data, and decide if and how to encrypt the data while it is in transit and at rest.

AWS provides customers with the ability to delete their data. However, AWS customers retain control and ownership of their data, and it is the customer's responsibility to manage their data

Data Recovery/Transfer

AWS allows customers to move data as needed on and off AWS storage using the public Internet or AWS Direct Connect (which lets customers establish a dedicated network connection between their network and AWS).

AWS Import/Export accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS transfers customer data directly onto and off of storage devices using Amazon's high-speed internal network and bypassing the Internet. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than customers upgrading their connectivity. With Import/Export encryption is mandatory, and AWS will encrypt customer data using the password they specified and transfer it onto the device

Deleting Data

Customers can use Multi-Object Delete to delete large numbers of objects from Amazon S3. This feature allows customers to send multiple object keys in a single request to speed up their deletes. Amazon does not charge customers for using Multi-Object Delete.

Customers can use the Object Expiration feature to remove objects from their buckets after a specified number of days. With Object Expiration customers can define the expiration rules for a set of objects in their bucket through the Lifecycle Configuration policy that they apply to the bucket. Each Object Expiration rule allows customers to specify a prefix and an expiration period.

Archiving Data

With Amazon S3's lifecycle policies, customers can configure their objects to be archived to Amazon Glacier or deleted after a specific period of time. Customers can use this policy-driven automation to quickly and easily reduce storage costs as well as save time. In each rule customers can specify a prefix, a time period, a transition to Amazon Glacier, and/or an expiration. For example, customers could create a rule that archives all objects with the common prefix "logs/" 30 days from creation, and expires these objects after 365 days from creation. Customers can also create a separate rule that only expires all objects with the prefix "backups/" 90 days from creation. Lifecycle policies apply to both existing and new S3 objects, ensuring that customers can optimize storage and maximize cost savings for all current data and any new data placed in S3 without time-consuming manual data review and migration.

AWS Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft data centers are in nondescript buildings that are physically constructed, managed, and monitored 24 hours a day to protect data and services from unauthorized access as well as environmental threats. The data centers are surrounded by a fence with access restricted by badge controlled gates. Preapproved deliveries are received in a secure loading bay and monitored by authorized personnel. Loading bays are physically isolated from information processing facilities.

CCTV is used to monitor physical access to data centers and the information systems. Cameras are positioned to monitor perimeter doors, facility entrances and exits, interior aisles, caged areas, high-

security areas, shipping and receiving, and facility external areas such as parking lots and other areas of the facilities.

Microsoft data centers receive SSAE16/ISAE 3402 Attestation and are certified to ISO 27001. Microsoft GFS and Azure maintain a current, documented, and audited inventory of equipment and network components for which they are responsible. GFS uses automated mechanisms to detect discrepancies of device configuration by comparing them with the defined policies. To prevent unauthorized access, GFS turns the unused ports off by default.

Windows Azure Fabric Controlled Hardware Device Authentication maintains a set of credentials (keys, passwords, or both) used to authenticate itself to various Windows Azure hardware devices under its control. The system used for transporting, persisting, and using these credentials is designed to make it unnecessary for Windows Azure developers, administrators, and backup services and personnel to be exposed to secret information. The Azure platform is designed and architected specifically to prevent the possibility of moving or replicating production data outside the Azure cloud environment. Controls for protecting production data include the following:

- Physical and logical network boundaries with strictly enforced change control policies
- Segregation of duty requiring a business need to access an environment
- Highly restricted physical and logical access to the cloud environment
- Strict controls based on SDL and OSA that define coding practices, quality testing, and code promotion
- Ongoing awareness and training on security, privacy, and secure coding practices
- Continuous logging and audit of system access
- Regular compliance audits to maintain control effectiveness.

Microsoft Azure customers must define policies and establish controls for the replication or high availability of production data and the demarcation of their production environment.

Internally, Microsoft tracks data flows and network connectivity among its facilities worldwide. Microsoft will not transfer Customer Data outside the geographies the customer specifies (for example, from Europe to the United States or from the United States to Asia) except when necessary for Microsoft to provide customer support, troubleshoot the service, or comply with legal requirements; or when the customer configures the account to enable the transfer of customer data, including through the use of the following features and functionality:

- Features that do not enable geographic selection such as Content Delivery Network (CDN) that provides a global caching service
- Web and worker roles, which back up software deployment packages to the United States regardless of deployment geography
- Preview, beta, or other prerelease features that may store or transfer customer data to the United States regardless of deployment geography
- Azure Active Directory (except for Access Control), which may store Active Directory data globally except for the United States (when Active Directory data remains in the United States) and Europe (when Active Directory Data is in Europe and the United States)
- Azure Multi-Factor Authentication, which stores authentication data in the United States
- Azure RemoteApp, which may store end user names and device IP addresses globally, depending on where the end user accesses the service.

Microsoft established policies, procedures, and mechanisms for effective key management to support encryption of data in storage and in transmission for the Azure service's key components. Azure provides each subscription with an associated logical certificate store that enables automatic deployment of service-specific certificates, and to which customers can upload their own.

Certificates used in Azure are x.509 v3 certificates and can be signed by another trusted certificate, or they can be self-signed. The certificate store is independent of hosted services, so it can store certificates regardless of whether those services currently use them. These certificates and other credentials uploaded to Azure are stored in encrypted form.

Microsoft Azure does not encrypt tenant data in storage. However, tools in Azure and third-party tools allow encryption of data in Azure storage. To implement encryption at rest, customers use .NET cryptographic services.

For customers using virtual machines, additional options are available, including the Encrypting File System in Windows Server 2008 R2 (and newer), Azure Rights Management Services, and Transparent Data Encryption (TDE) in SQL Server 2008 R2 (and newer).

When using Azure SQL Database, externally encrypted records cannot be queried by T-SQL (other than “retrieve all”) and may require a schema change such as the introduction of surrogate keys to enable retrieval of specific records or ranges of records.

Microsoft uses best practice procedures and a wiping solution that complies with NIST 800-88, Guidelines on Media Sanitization. For hard drives that cannot be wiped, Microsoft uses a destruction process that destroys it (shredding) and renders the recovery of information impossible (e.g., disintegrate, shred, pulverize, or incinerate). The appropriate means of disposal is determined by the asset type. Records of the destruction are retained. Microsoft Azure services use approved media storage and disposal management services. Paper documents are destroyed by approved means at the predetermined end-of-life cycle.

Data destruction techniques vary depending on the type of data object being destroyed, whether it be subscriptions, storage, virtual machines, or databases. In Azure’s multitenant environment, careful attention is taken to prevent one customer’s data from leaking into another customer’s data; or when a customer deletes data, no other customer (usually including the customer that once owned the data) from gaining access to that deleted data.

Azure follows NIST 800-88 guidelines, which address the principal concern of guarding data from unintentional release. These guidelines encompass electronic and physical sanitization.

Google Public and Community Cloud for SaaS: Google has deployed encryption services for data at rest and data in flight as one aspect of data protection. Data protection is a primary design consideration for all of Google’s infrastructure, applications and personnel operations. These factors are audited against several Industry Security Standards such as SOC2, ISO27001 and FedRAMP. The key security controls include multi-point perimeter protection, real time scanning for malware, systems designs that block or disable unnecessary services and protocols, least privilege access models, extensive logging and log review cycles. As data is written to Google’s infrastructure it is obfuscated, sharded, encrypted and replicated in a unique manner that not only ensure data availability but also security. As no single file for one single user at one single customer account is stored all in one place, any effort to access the data through unauthorized channels is extremely complex and highly unlikely.

Data disposal at the completion of any contract services is largely the responsibility of the customer. They would exercise an exit strategy that begins with migrating the data to a different service provider, deleting the data from Google, deleting the User accounts and closing or cancelling the service. Google agrees, via the Data Processing Agreement, that all customer data shards will be off all servers and backup tapes within 180 days of deletion.

8.5.2 Offeror must describe how it intends to comply with all applicable laws and related to data privacy and security.

In accordance with mandated policy, Unisys and our partners in this proposal comply with regulatory requirements for data privacy and security. Designated staff such as the contract manager, security

officers, and subject matter experts coordinate with the legal offices and external enforcement agencies to facilitate the requirements for data privacy and security.

8.5.3 *Offeror must describe how it will not access a Purchasing Entity's user accounts or data, except in the course of data center operations, response to service or technical issues, as required by the express terms of the Master Agreement, the applicable Participating Addendum, and/or the applicable Service Level Agreement.*

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

In accordance with established security policy, standards, and procedures, Unisys and our partners limit access to data relevant to job responsibilities, segregation of duties, and access controls. Technology controls are established to monitor, enforce, and report on violations. Unisys also provides provisions for a customer to enforce available security policies around their data as and when applicable. Unisys understands and respects the sensitivity of customer data and will not access a Purchasing Entity's user accounts or data.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS does not access customer data, and customers are given the choice as to how they store, manage and protect their data. AWS customers retain control and ownership of their data, and all data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers should consider the sensitivity of their data and decide if and how they will encrypt data while it is in transit and while it is at rest.

AWS Regions and Availability Zones: The AWS cloud infrastructure is built around regions and Availability Zones. A region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity and housed in separate facilities. These Availability Zones offer you the ability to operate production applications and databases that are more highly available, fault tolerant and scalable than would be possible from a single data center. AWS currently has 12 regions and 32 Availability Zones throughout the world: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), South America (Sao Paulo), and China (Beijing).. ~~Figure 1~~ depicts the current AWS regions and Edge Locations, along with new regions that are coming soon.

Formatte
Not Bold




Figure 1 – Global Map of AWS Regions and Edge Locations

Figure 2 illustrates the relationship between regions and Availability Zones.




Figure 2 – Regions and Availability Zones

Shared Responsibility: As cloud computing customers are building systems on top of cloud infrastructure, the security and compliance responsibilities are shared between the Cloud Service Providers (CSP) and cloud customers. In an Infrastructure as a Service (IaaS) model, customers control how they architect and secure their applications and data put on the infrastructure, while CSPs are responsible for providing services on a highly secure and controlled platform and providing a wide array of additional security features. The level of CSP and customer responsibilities in this shared responsibility model depends on the cloud deployment. Customers should be clear as to their responsibilities in each model. AWS's shared responsibility/security model is depicted in **Figure 3**.




Figure 34 – AWS Shared Responsibility Model

- **AWS Responsibility:** AWS operates, manages, and controls the infrastructure components, from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate.
- **Customer/Partner Responsibility:** Customers/partners assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, configuration of the AWS-provided security group firewalls, and other security, change management, and logging features.

AWS does not access customer data, and customers are given the choice as to how they store, manage and protect their data. There are four important basics regarding data ownership and management in the shared responsibility model:

- 1) Customers continue to own their data.
- 2) Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
- 3) Customers can download or delete their data whenever they like.
- 4) Customers should consider the sensitivity of their data and decide if and how to encrypt the data while it is in transit and at rest.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Privacy: Customers own and control their data.

For more than 20 years, Microsoft has been a leader in creating robust online solutions designed to protect its customers' privacy. Microsoft's time-tested approach to privacy and data protection is grounded in its commitment to organizations' ownership of and control over the collection, use, and distribution of their information.

Microsoft strives to be transparent in its privacy practices, offer customers meaningful privacy choices, and responsibly manage the data Microsoft stores and processes. One measure of Microsoft's commitment to the privacy of customer data is its adoption of the world's first code of practice for cloud privacy, ISO/IEC 27018.

Customers own their own data. With Azure, they have ownership of customer data, including text, sound, video, or image files and software, that are provided to Microsoft provides to them, or on their behalf, through the use of Azure. Customers can access their any time and for any reason without assistance from Microsoft. Microsoft does not use customer data or derive information from it for advertising or data mining.

Customers are in control of their data. Because the data customers host on Azure belongs to them, they have control over where their data is stored and how it is securely accessed and deleted.

Microsoft responds as follows to government and law enforcement requests to access data. When a government wants customer data—including for national security purposes—it must follow the applicable legal process, serving Microsoft with a court order for content or a subpoena for account information. If compelled to disclose customer data, Microsoft will promptly notify customers and provide a copy of the demand, unless legally prohibited from doing so. Microsoft does not provide governments with direct or unfettered access to customer data except as a customer directs or where required by law.

Security: Microsoft keeps customer data safe.

Microsoft leveraged its decades-long experience building enterprise software and running some of the world's largest online services to create a robust set of security technologies and practices. These help to confirm that Azure infrastructure is resilient to attack, safeguard user access to the Azure environment, and keep customer data secure through encrypted communications as well as threat management and mitigation practices, including regular penetration testing.

To manage and control identity and user access to customers' environments, data, and applications, Microsoft federates user identities to Azure Active Directory and enables multifactor authentication for more secure sign-in.

Microsoft encrypts communications and operation processes as follows. For data in transit, Azure uses industry-standard transport protocols between user devices and Microsoft data centers and in data centers themselves. For data at rest, Azure offers a wide range of encryption capabilities up to AES-256, giving customers the flexibility to choose the solution that best meets their needs.

To secure networks, Azure provides the infrastructure necessary to securely connect virtual machines to one another and to connect on-premises data centers with Azure virtual machines. To block unauthorized traffic to and in Microsoft data centers, Azure uses a variety of technologies. Azure Virtual Network extends a customer's on-premises network to the cloud over a site-to-site VPN.

To protect against online threats, Azure offers Microsoft Antimalware for cloud services and virtual machines. To help to mitigate threats to the Azure platform, Microsoft uses intrusion detection, denial-of-service (DDoS) attack prevention, regular penetration testing, and data analytics and machine learning tools.

To help organizations comply with national, regional, and industry-specific requirements governing the collection and use of individuals' data, Microsoft offers the most comprehensive set of certifications and attestations. Microsoft accomplishes this breadth of compliance offerings with a two-pronged approach:

- First, a team of Microsoft experts works with its engineering and operations teams, as well as external regulatory bodies, to track existing standards and regulations, developing hundreds of controls for the product teams to build into Microsoft's cloud services.
- Second, because regulations and standards are always evolving, Microsoft compliance experts anticipate upcoming changes to help maintain continuous compliance—researching draft regulations, assessing potential new requirements, and developing corresponding controls.

To demonstrate that these controls deliver compliance customers can rely on, Microsoft enterprise cloud services are independently validated through certifications and attestations as well as third-party audits.

In-scope services in the Microsoft Cloud meet key international and industry-specific compliance standards such as ISO/IEC 27001 and ISO/IEC 27018, FedRAMP, and SOC 1 and SOC 2. They also meet regional and country-specific standards and contractual commitments, including the EU Model Clauses, UK G-Cloud, Singapore MTCS, and Australia CCSL (IRAP). Additionally, rigorous third-party audits, such as by the British Standards Institution and Deloitte, validate the adherence of Microsoft cloud services to the strict requirements these standards mandate.

Google Public and Community Cloud for SaaS: Google will take and implement appropriate technical and organizational measures to protect Customer Data against accidental or unlawful destruction or accidental loss or alteration or unauthorized disclosure or access or other unauthorized processing. Google may update or modify the Security Measures from time to time provided that such updates and modifications do not result in the degradation of the overall security of the Services. Customer agrees that it is solely responsible for its use of the Services, including securing its account authentication credentials, and that Google has no obligation to protect Customer Data that Customer elects to store or transfer outside of Google's and its Subprocessors' systems (e.g., offline or on-premise storage). Google commits via terms of service to abide by all applicable laws and to hold all of customer's data as confidential and solely owned by the customer.

8.6 (E) PRIVACY AND SECURITY

8.6.1 Offeror must describe its commitment for its Solutions to comply with NIST, as defined in NIST Special Publication 800-145, and any other relevant industry standards, as it relates to the Scope of Services described in Attachment D, including supporting the different types of data that you may receive.

8.6.2 Offeror must list all government or standards organization security certifications it currently holds that apply specifically to the Offeror's proposal, as well as those in process at time of response. Specifically include HIPAA, FERPA, CJIS Security Policy, PCI Data Security Standards (DSS), IRS Publication 1075, FISMA, NIST 800-53, NIST SP 800-171, and FIPS 200 if they apply.

8.6.3 Offeror must describe its security practices in place to secure data and applications, including threats from outside the service center as well as other customers co-located within the same service center.

8.6.4 Offeror must describe its data confidentiality standards and practices that are in place to ensure data confidentiality. This must include not only prevention of exposure to unauthorized personnel, but also managing and reviewing access that administrators have to stored data. Include information on your hardware policies (laptops, mobile etc).

8.6.5 Offeror must provide a detailed list of the third-party attestations, reports, security credentials (e.g., FedRamp), and certifications relating to data security, integrity, and other controls.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4 and 8.6.5:

Unisys Cloud Services align with the NIST definition of cloud computing (NIST 800-145). Additionally, Unisys is certified to ISO 27001 for information security, which meets the security requirements for cloud computing. Along with ISO 27001, Unisys also aligned our security organization with other standards and frameworks such as HIPAA, HITECH, GLBA, ISAE 3402/SSAE 16, SOGP, NIST, Data Protection Acts, COBIT, PCI-DSS, FedRAMP, FISMA, CSA, FFIEC, ISO (27000, 9000, 20000, and 22301), ENISA, and various other country- and region-specific standards and requirements.

Unisys delivery centers are certified to ISO 27001 (Information Security), ISO 20000 (Service Management), ISO 9001 (Quality Management), and ISO 22301 (Business Continuity). Several of our client environments are certified to PCI-DSS standards. Additionally, our delivery centers are assessed annually for SOC 1 reporting. Independent third-party agencies audit these certifications every year.

Where required, physical segregation of the work space is implemented, and the workspace is secured through access controls. Networks are segregated by VLANs and protected by firewalls. Access control mechanisms are implemented to access data and applications.

Unisys adopted a data protection policy and standards for securing data. These standards define controls such as data classification, system maintenance, access controls, data storage, external data transfer,

media handling, and media disposal. These controls are followed across Unisys and audited regularly for compliance. To complement these controls, Unisys also has policies for Internet and email use, mobile devices, and acceptable use, supplemented by various security standards.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4, and 8.6.5:

The AWS cloud infrastructure is designed and managed in alignment with the following regulations, standards, and best practices:

- Federal Risk and Authorization Management Program (FedRAMP)
- Service Organization Controls (SOC) 1/American Institute of Certified Public Accountants (AICPA): AT 801 (formerly Statement on Standards for Attestation Engagements [aud] No. 16)/International Standard on Assurance Engagements (ISAE) 3402 (formerly Statement on Auditing Standards [SAS] No. 70)
- SOC 2
- SOC 3
- Payment Card Industry Data Security Standard (PCI DSS)
- International Organization for Standardization (ISO) 27001
- ISO 27017
- ISO 27018
- ISO 9001
- Department of Defense (DoD) Security Requirements Guide (SRG) security impact levels 2 and 4
- Federal Information Security Management Act (FISMA)
- US Health Insurance Portability and Accountability Act (HIPAA)
- FBI Criminal Justice Information Services (CJIS)
- National Institute of Standards and Technology (NIST) 800-171
- International Traffic in Arms Regulations (ITAR)
- Federal Information Processing Standard (FIPS) 140-2
- Family Educational Rights and Privacy Act (FERPA)
- Information Security Registered Assessors Program (IRAP) (Australia)
- IT-Grundschutz (Germany).

The AWS white paper in **Appendix 14 AWS Risk and Compliance Whitepaper** contains further details:

See Appendix 14 AWS Risk and Compliance Whitepaper

Additionally, the AWS white paper in **Appendix 15 AWS Security Whitepaper** also contains relevant information surrounding privacy and security.

See Appendix 15 AWS Security Whitepaper

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4 and 8.6.5:

Microsoft agrees that, during the term of a Purchasing Entity's subscription for its Government Community Cloud Services (as defined in Microsoft's terms and conditions), those services will be operated in accordance with a written data security policy and control framework that is consistent with the requirements of NIST 800-53 Revision 4, or successor standards and guidelines (if any), established to support Federal Risk and Authorization Management Program (FedRAMP) accreditation at a Moderate Impact level. Microsoft intends for Government Community Cloud Services to support FedRAMP Authority to Operate (ATO). Microsoft will use commercially reasonable efforts to obtain an ATO from a Federal agency, and to maintain this ATO through continuous monitoring processes and by conducting

regular FedRAMP audits. The Microsoft Trust Center available on line contains detailed information on Microsoft's compliance and adherence to other standards, such as CJIS, IRS 1075, HIPAA, FERPA, ISO/IEC 27001 and 27018, SOC 1, and SOC 2. The **Windows Azure Security and Compliance** white paper in **Appendix 24** contains further details:

See Appendix 24 **Windows Azure Security and Compliance White Paper**

Microsoft Dynamics CRM Online Government has received an Agency ATO from HUD; and Microsoft Office 365 U.S. Government has received an Agency ATO from DHHS.

The US Food and Drug Administration (FDA) Code of Federal Regulations (CFR) Title 21 Part 11 lists requirements for the security of electronic records of companies that sell food and drugs manufactured or consumed in the United States

The Defense Information Systems Agency (DISA) Cloud Service Support has granted a DISA Impact Level 2 Provisional Authorization (PA) to Microsoft Azure, Microsoft Azure Government, Microsoft Office 365 MT, and Microsoft Office 365 U.S. Government, based on their FedRAMP authorizations.

Microsoft contractual commitments, customers that are subject to FERPA can use Microsoft Azure, Microsoft Dynamics CRM Online, and Microsoft Office 365 and comply with FERPA.

NIST publishes a list of vendors and their cryptographic modules validated for FIPS 140-2. Rather than validate individual components and products, Microsoft certifies the underlying cryptographic modules used in Microsoft products, including Microsoft enterprise cloud services.

Microsoft engaged outside assessors to validate that Microsoft Azure and Microsoft Office 365 meet the FISC Version 8 requirements.

Microsoft enterprise cloud services offer customers a HIPAA Business Associate Agreement (BAA) that stipulates adherence to HIPAA's security and privacy provisions.

Microsoft Azure and Microsoft Office 365 were among the first cloud services to achieve this certification for the storage and processing of unclassified (DLM) data.

Microsoft Azure Government and Microsoft Office 365 U.S. Government cloud services provide a contractual commitment that they have the appropriate controls in place, and the security capabilities necessary for customers to meet the substantive requirements of IRS 1075.

The ISO/IEC 27001 certificate validates that Microsoft enterprise cloud services have implemented the internationally recognized information security controls defined in this standard, including guidelines and general principles for initiating, implementing, maintaining, and improving information security management within an organization.

Microsoft was the first cloud provider to adhere to the ISO/IEC 27018 code of practice, covering the processing of personal information by cloud service providers.

Azure complies with Payment Card Industry (PCI) Data Security Standards (DSS) Level 1 version 3.0, the global certification standard for organizations that accept most payment cards and store, process, or transmit cardholder data.

A Voluntary Product Accessibility Template, or VPAT, is a standardized form developed by the Information Technology Industry Council to document whether a product meets key regulations of Section 508, an amendment to the Rehabilitation Act of 1973. Microsoft offers detailed VPATs for many of its core cloud services, describing the accessibility features of those services.

Service Organization Controls (SOC) are a series of accounting standards that measure the control of financial information for a service organization. Azure's SOC 1 and SOC 2 Type 2 audit reports attest to the effectiveness of the design and operation of its security controls.

Other country-specific standards and contractual commitments, including the EU Model Clauses, UK G-Cloud, Singapore MTCS, and Australia CCSL (IRAP).

As shown in Exhibit 5, Microsoft cloud services have the largest compliance portfolio in the industry, the United States, and the world.

Exhibit 5. Regulatory Compliance of Microsoft Cloud Services.

Regulatory Compliance

Microsoft cloud services have the largest compliance portfolio in the industry

Industry	 ISO 27001	 SOC 1 Type 2	 SOC 2 Type 2	 PCI DSS Level 1	 Cloud Controls Matrix	 ISO 27018	 Content Delivery and Security Association	 Shared Assessments			
United States	 FedRAMP JAB P-ATO	 HIPAA / HITECH	 FIPS 140-2	 21 CFR Part 11	 FERPA	 DISA Level 2	 CJIS	 IRS 1075	 ITAR ready	 Section 508 VPA	
Regional	 European Union Model Clauses	 EU Safe Harbor	 GOV.UK	 China Multi Layer Protection Scheme	 China GB 18030	 China CCCPF	 Singapore MTCS Level 3	 Australian Signals Directorate	 New Zealand GCIO	 Japan Financial Services	 ENISA IAF

Data in Government Clouds are maintained in Azure Datacenters in the United States. Azure Government is a government-community cloud that you can trust, enabling US government agencies and their partners to transform their mission-critical workloads to the cloud. A physical and logical network-isolated instance of Azure, operated by screened US persons, Azure Government has been specifically designed to meet the needs of US government agencies and their partners.

Security Practices in Place to Secure Data and Applications:

The Microsoft Azure trustworthy foundation concept maintains application security through a process of continuous security improvement with its Security Development Lifecycle (SDL) and Operational Security Assurance (OSA) programs using Prevent Breach and Assume Breach security postures.

Prevent Breach works through the use of ongoing threat modeling, code review, and security testing; Assume Breach uses Red Team Blue Team exercises, live site penetration testing, and centralized security logging and monitoring to identify and address potential gaps; test security response plans; reduce exposure to attack; and reduce access from a compromised system, periodic postbreach assessment, and clean state.

To validate services, Azure uses third-party penetration testing based on the OWASP (Open Web Application Security Project) top 10 and CREST-certified testers. The outputs of testing are tracked on the risk register, which is audited and reviewed regularly to confirm compliance with Microsoft security practices.

Microsoft maintains and regularly updates the Azure Information Security Management Policy and information security guidelines, standard operating procedures for data security, and contractual commitments to international data protection directives that apply across Azure services.

Additionally, Microsoft Azure software updates are reviewed for unauthorized changes through the Security Development Lifecycle (SDL) change and release management processes. Automated mechanisms are used to perform integrity scans at least every hour to detect system anomalies or unauthorized changes. Microsoft applies SDL to design, develop, and implement Microsoft Azure services. SDL helps to confirm that communication and collaboration services are highly secure, even at the foundation level, and align with other industry standards, including FedRAMP, ISO, and NIST.

Data Confidentiality Standards and Practices That Are in Place to Maintain Data Confidentiality

Microsoft Azure adopted applicable corporate and organizational security policies, including an Information Security Policy. The policies were approved, published, and communicated across Azure teams. The Information Security Policy requires that access to Microsoft Azure assets to be granted based on business justification, with the asset owner's authorization and limits based on "need-to-know" and "least privilege" principles. Additionally, the policy also addresses requirements for the access management life cycle, including access provisioning, authentication, access authorization, removal of access rights, and periodic access reviews.

Access to Microsoft buildings is controlled, and access is restricted to those with a card reader (swiping the card reader with an authorized ID badge) or biometrics for entry into data centers. Front-desk personnel are required to positively identify full-time employees or authorized contractors without ID cards. Staff must always wear identity badges and are required to challenge or report individuals without badges. Authorized Microsoft personnel must escort guests.

Password policies for corporate domain accounts are managed through Microsoft's corporate Active Directory policy that specifies minimum requirements for password length, complexity, and expiration. Temporary passwords are communicated to users using Microsoft's established processes. Azure services and infrastructure must at least meet Microsoft corporate requirements, but an internal organization can increase the strength beyond this standard, at their own discretion and to meet their security needs. Technological safeguards, such as encrypted communications and operational processes, enhance the security of our customers' data. For data in transit, the Microsoft Cloud uses industry-standard encrypted transport protocols between user devices and Microsoft datacenters, and within datacenters themselves. For data at rest, the Microsoft Cloud offers a wide range of encryption capabilities up to AES-256, giving you the flexibility to choose the solution that best meets your needs.

Microsoft Azure uses Active Directory (AD) to manage user accounts. Azure Active Directory is a comprehensive identity and access management cloud solution that helps secure access to your data and on-premises and cloud applications, and simplifies the management of users and groups. It combines core directory services, advanced identity governance, security, and application access management, and is a key component of Microsoft Cloud services, including Microsoft Azure, Office 365, Microsoft Dynamics CRM Online, and Intune, as well as thousands of third-party SaaS apps. Azure Active Directory also makes it easy for developers to build policy-based identity management into their applications. The designated security group owners must approve security group membership in Microsoft Azure. Automated procedures are in place to disable AD accounts on the user's leave date. Domain-level user accounts are disabled after 90 days of inactivity.

Mobile and wireless devices are not permitted in Microsoft data centers where customer data is stored. Wireless access to Azure production environments, customer environments, or both is prohibited.

Employees, contractors, and third-party users are formally notified to destroy or return, as applicable, physical material that Microsoft provided to them during their employment or period of the contractor agreement; electronic media must be removed from contractor or third-party infrastructure. Microsoft may also conduct an audit to verify that data is removed in an appropriate way.

Google Public and Community Cloud for SaaS Collective Response to RFP requirements 8.6.1, 8.6.2, 8.6.3, 8.6.4 and 8.6.5:

Google has an FedRAMP ATO at the Moderate impact baseline. FedRAMP incorporates many NIST SPs and FIPS including 800-53, FIPS 199, FIPS 200), and has a specific offering. Google Apps for Education that is FERPA and COPPA compliant. Other compliance standards such as HIPAA and CJIS don't offer certification per se, but are commonly accommodated (i.e. Google will sign a BAA to meet HiTECH/HIPAA requirements, and has numerous customers who bear responsibility for meeting CJIS processing requirements). PCI DSS is generally not applicable to SaaS systems (though we can do email

hygiene processing to protect against incidental usage), but Google IaaS/PaaS does meet PCI DSS v3 standards. Google also holds and is committed to maintaining SOC2 and ISO27001 certifications.

At Google, members of the information security team review security plans for networks, systems, and services. They provide project-specific consulting services to Google's product and engineering teams. They monitor for suspicious activity on Google's networks, address information security threats, perform routine security evaluations and audits, and engage outside experts to conduct regular security assessments. Google specifically built a full-time team, known as Project Zero, which aims to prevent targeted attacks by reporting bugs to software vendors and filing them in an external database. Google has automated mechanisms that scan the infrastructure continuously to detect and correct deviations from the desired security configuration of its infrastructure. Google's security monitoring program is focused on information gathered from internal network traffic, employee actions on systems, and outside knowledge of vulnerabilities. At many points across Google's global network, internal traffic is inspected for suspicious behavior, such as the presence of traffic that might indicate botnet connections. This analysis is performed using a combination of Open Source and commercial tools for traffic capture and parsing. A proprietary correlation system built on top of Google technology also supports this analysis. Network analysis is supplemented by examining system logs to identify unusual behavior, such as attempted access of customer data. Google security engineers issue standing search alerts on public data repositories to look for security incidents that might affect Google's infrastructure. They actively review inbound security reports and monitor public mailing lists, blog posts, and wikis. Automated network analysis helps determine when an unknown threat may exist and escalates the threat to Google security staff; network analysis is supplemented by automated analysis of system logs. Google uses a proprietary storage and processing mechanism that isolates processing of data in change-rooted (chrooted) jails in a physical server. Access permissions restrict the ability for processes to interact between jails. Additionally, tenant data is stripped in chunks across many different drives; each chunk has its own access control list. This helps confirm that data is logically isolated between customers in storage and during processing. In addition to the processing and storage mechanisms already described, Google's security controls, including least privilege rights, logging, and auditing, were implemented consistently with FedRAMP requirements.

A copy of Google's Data Processing Agreement (DPA) is included in this proposal as **Appendix 16-Google-DPA**. See Appendix 16

In this DPA, the obligations of Google to hold customer data as confidential and wholly owned by the customer is detailed. Google's internal data access processes and policies are designed to prevent unauthorized persons, systems, or both from gaining access to systems used to process personal data. Google aims to design its systems to: (1) only allow authorized persons to access data they are authorized to access; and (2) confirm that personal data cannot be read, copied, altered, or removed without authorization during processing, use, and after recording. The systems are designed to detect inappropriate access. Google uses a centralized access management system to control personnel access to production servers and provides access only to a limited number of authorized personnel. LDAP, Kerberos, and a proprietary system using RSA keys are designed to provide Google with secure and flexible access mechanisms. These mechanisms are designed to grant only approved access rights to site hosts, logs, data, and configuration information. Google requires the use of unique user IDs, strong passwords; two-factor authentication, and carefully monitored access lists to minimize the potential for unauthorized account use. The granting or modification of access rights is based on the authorized personnel's job responsibilities, job duty requirements necessary to perform authorized tasks, and the need to know; they also must be in accordance with Google's internal data access policies and training. Approvals are managed by workflow tools that maintain audit records of changes. Access to systems is logged to create an audit trail for accountability. Where passwords are employed for authentication (e.g., login to workstations), password policies that follow at least industry-standard practices are implemented.

These standards include password expiration, restrictions on password reuse, and sufficient password strength. For access to extremely sensitive information (e.g., credit card data), Google uses hardware tokens. **Appendix 23 contains the Google Apps Security and Compliance** whitepaper contains additional details surrounding privacy and Security. See Appendix 23.

8.6.6 Offeror must describe its logging process including the types of services and devices logged; the event types logged; and the information fields. You should include detailed response on how you plan to maintain security certifications.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

Unisys has a Unisys Cloud Provisioning System logs all actions by both administrators and users. This logging system will be accessible by authorized Unisys personnel who can manage and monitor data for specific log types. Administrators can build templates and generate alerts on custom correlation events. All logs are text-based and depending on the log types can be customized between a syslog or log4j. This gives Unisys Cloud Provisioning System the ability to centralize log data and use specific tools such as splunk, q1 labs or SIEM tools to access logs.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS CloudTrail

AWS CloudTrail is a web service that records API calls to supported AWS services in an AWS account, delivering a log file to an Amazon Simple Storage Service (Amazon S3) bucket. To alleviate common challenges experienced in an on-premises environment, AWS CloudTrail makes it easier for customers to enhance security and operational processes while demonstrating compliance with policies or regulatory standards.

With AWS CloudTrail, customers can get a history of AWS API calls for their account, including API calls made from the AWS Management Console, AWS SDKs, command line tools, and higher level AWS services (such as AWS CloudFormation). The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

- For information on the services and features supported by AWS CloudTrail, visit the AWS CloudTrail FAQs at the AWS website.
- The AWS white paper Security at Scale: Logging In AWS provides an overview of common compliance requirements related to logging, detailing how AWS CloudTrail features can help satisfy these requirements.
- The AWS white paper Auditing Security Checklist for Use of AWS provides customers with a checklist to assist in evaluating AWS for an internal review or external audit.

AWS CloudTrail includes the following features:

- **Increased visibility:** AWS CloudTrail provides increased visibility into user activity by recording AWS API calls. Customers can answer questions such as: What actions did a given user take over a given period? For a given resource, which user took actions on it over a given period? What is the source IP address of a given activity? Which activities failed because of inadequate permissions?
- **Durable and inexpensive log file storage:** AWS CloudTrail uses Amazon S3 for log file storage and delivery, so that log files are stored durably and inexpensively. Customers can use Amazon S3 life cycle configuration rules to further reduce storage costs. For example, customers can define rules to automatically delete old log files or archive them to Amazon Glacier for additional savings.
- **Easy administration:** AWS CloudTrail is a fully managed service; customers simply turn on AWS CloudTrail for their account at the AWS Management Console, the Command Line Interface, or the AWS CloudTrail SDK to start receiving AWS CloudTrail log files in the specified Amazon S3 bucket.

- **Notifications for log file delivery:** AWS CloudTrail can be configured to publish a notification for each log file delivered, thereby enabling customers to automatically take action upon log file delivery. AWS CloudTrail uses the Amazon Simple Notification Service (Amazon SNS) for notifications.
- **Choice of partner solutions:** Many partners, including AlertLogic, Boundary, Loggly, Splunk, and Sumologic, offer integrated solutions to analyze AWS CloudTrail log files. These solutions include features such as change tracking, troubleshooting, and security analysis. For more information, see the AWS CloudTrail partners section of Amazon's home page.
- **Log file aggregation:** AWS CloudTrail can be configured to aggregate log files across multiple accounts and regions so that log files are delivered to a single bucket. For detailed instructions, refer to the Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket section of the user guide.

Amazon CloudWatch

Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications that run on AWS. Customers can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by customer applications and services, and log files that application generate. Customers can use Amazon CloudWatch to gain systemwide visibility into resource utilization, application performance, and operational health, using these insights to react and keep their application running smoothly.

Customers also can use CloudWatch Logs to monitor and troubleshoot systems and applications using their existing system, application, and custom log files. They can send their existing system, application, and custom log files to CloudWatch Logs and monitor these logs in near real time. This helps customers to better understand and operate their systems and applications; they can move their logs to highly durable, low-cost storage for later access.

LogAnalyzer for Amazon CloudFront

LogAnalyzer allows customers to analyze their Amazon CloudFront Logs using Amazon Elastic MapReduce (Amazon EMR). Using Amazon EMR and the LogAnalyzer application, customers can generate usage reports containing total traffic volume, object popularity, a breakdown of traffic by client IPs, and edge location. Reports are formatted as tab delimited text files and delivered to the Amazon S3 bucket that customers specify.

Amazon CloudFront's Access Logs provide detailed information on requests made for content delivered through Amazon CloudFront, the AWS content delivery service. The LogAnalyzer for Amazon CloudFront analyzes the service's raw log files to produce a series of reports that answer business questions commonly asked by content owners.

Reports Generated

This LogAnalyzer application produces four sets of reports based on Amazon CloudFront access logs. The Overall Volume Report displays the total amount of traffic delivered by CloudFront during the specified period. The Object Popularity Report shows how many times each customer object is requested. The Client IP report shows the traffic from each different client IP that made a request for content. The Edge Location Report shows the total number of traffic delivered through each edge location. Each report measures traffic in three ways: the total number of requests, the total number of bytes transferred, and the number of request broken down by HTTP response code. The LogAnalyzer is implemented using Cascading and is an example of how to construct an Amazon Elastic MapReduce application. Customers can also customize reports generated by the LogAnalyzer.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Azure enables customers to perform security event generation and collection from Azure Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) roles to central storage in their subscriptions. Customers

can then use HDInsight to aggregate and analyze the collected events. Additionally, these collected events can be exported to on-premises security information and event management (SIEM) systems for ongoing monitoring.

The Azure security logging, analysis, and monitoring life cycle includes the following tasks:

- Generation: Instrument applications and the infrastructure to raise events.
- Collection: Configure Azure to collect the various security logs in a storage account.
- Analysis: Use Azure tools such as HDInsight and on-premises SIEM systems to analyze the logs and generate security insights.
- Monitoring and reporting: Azure offers centralized monitoring and analysis systems that provide continuous visibility and timely alerts.

Security events are flagged on the Windows Event Log for the System, Security, and Application channels on virtual machines. To prevent events from being logged without potential data loss, it is important to appropriately configure the size of the event log. The size of this log must be based on the number of events that auditing policy settings generate and the event collection policies defined.

For Cloud Services applications that are deployed in Azure and virtual machines created from the Azure Virtual Machines Marketplace, a set of operating system security events are enabled by default. Customers can add, modify, or delete events to be audited by customizing the operating system audit policy. For more information, see Security Policy Settings Reference.

Customers can use the following methods to generate additional logs from operating system (such as audit policy changes) and Windows components (such as IIS):

- Group Policy to roll out policy settings for virtual machines in Azure that are domain joined
- Desired State Configuration (DSC) to push and manage policy settings; for more information, see Azure PowerShell DSC
- Service Deployment role startup code to roll out settings for Cloud Services (PaaS scenario).

Collection of security events and logs from Cloud Services or virtual machines in Azure occurs through two main methods:

- Azure Diagnostics, which collects events in a customer's Azure storage account
- Windows Event Forwarding (WEF), a technology on computers that run Windows.

Exhibit 6 lists several key differences between these two technologies. In accordance with a customer's requirements and these key differences, the appropriate method must to be chosen to implement log collection.

Exhibit 6. Key Differences between Azure Diagnostics and Windows Event Forwarding.

Azure Diagnostics	Windows Event Forwarding
Supports Azure Virtual Machines and Azure Cloud Services	Supports domain-joined Azure Virtual Machines only
Supports a variety of log formats, such as Windows event logs, Event Tracing for Windows (ETW) traces, and IIS logs. For more information, see Azure Diagnostics supported data sources	Supports Windows event logs only
Pushes collected data to Azure Storage	Moves collected data to central collector servers

Security Event Data Collection with Windows Event Forwarding

For domain-joined Azure Virtual Machines, customers can configure WEF by using Group Policy settings in the same way as for on-premises domain-joined computers. Using this approach, an organization could purchase an IaaS subscription, connect it to its corporate network by using ExpressRoute or site-to-site VPN, and then join the virtual machines that a customer has in Azure to the corporate domain. Afterwards, the customer can configure WEF from the domain-joined machines.

Event forwarding is broken into two parts: the source and the collector. The source is the computer in which the security logs are generated. The collector is the centralized server that collects and consolidates the event logs. IT administrators can subscribe to events so that they can receive and store events that are forwarded from remote computers (the event source). Collected Windows events can be sent to on-premises analysis tools, such as a SIEM, for further analysis.

Security Data Collection with Azure Diagnostics

Azure Diagnostics enables customers to collect diagnostic data from a cloud service worker role or web role, or from virtual machines running in Azure. It is a predefined guest agent extension that needs to be enabled and configured for data collection. A customer's subscription can include pushing the data to Azure Storage.

The data is encrypted in transit (by using HTTPS). The examples provided in this response use Azure Diagnostics 1.2. Microsoft recommends that customers upgrade to version 1.2 or higher for security data collection.

Exhibit 7 shows a high-level dataflow for security data collection that uses Azure Diagnostics and further analysis and monitoring.

Exhibit 7. High-level Dataflow for Collection of Security Data.

Google Public and Community Cloud for SaaS: Exhibit 8 lists Google's logging and reporting capability for the various Google offerings.

Exhibit 8. Google Reports.

Category	Report	Information Fields
Highlights	Apps Usage Activity	Total Email Files Owned Video Hangouts By Week, Month, Quarter, Half Year
	User Status	Total Number of Users Blocked Users Suspended Users Active Users
	Security	External Apps Installed Users not enrolled in Two-Step Verification Use of Less Secure (Third-Party) Apps
	Document Link Shared Status	Count of the following Shared Files: • Visible to Public • Visible to Anyone with Link • Visible to people at Domain • Visible to people at Domain with link • Private
Aggregate Reports 6 months of rolling data	GMail	Inbound Email Delivery Count of Delivered, Rerouted, Rejected Inbound Email Encryption Outbound Email Encryption Inbound Email Spam Outbound Email Delivery Total Emails
	Accounts	Two-step verification enrollment Two-step verification enforcement User account status Less secure apps access Admin status shows the number of users assigned to each Admin role defined
	Drive	Counts by day of the following files: • External Link Shared Files • Internal Link Shared Files • Total Files Owned
	Chrome Devices	Count of Weekly Chrome Device logins Count of devices by release channel Count of devices by boot mode Count of devices by version
	Mobile Managed Devices requires end	Count of 7-day-active Managed Devices by mobile OS (Android, iOS, and Active Sync)

Category	Report	Information Fields
	user configuration of mobile device management App. Management of Android, iOS and Active-Sync supported devices is included.	Count of 30-day active Managed Devices 7- and 30-day active managed Users Managed Android Devices by Android version Managed iOS Devices by iOS version
	Google+	Count of number of Video Hangouts Activity report on Chromebox for Meeting Devices
Security	General	Same as Accounts reports listed above.
	GMail	(POP) Last time used (IMAP) Last time used (Web) Last time used
	Drive	Same as Drive reports listed above.
Apps Usage Activity	General	Count by User of the following storage: • Total Storage used • GMail Storage used • Drive Storage used • Photos Storage used
	GMail	Same as GMail reports listed in Security and Aggregate Reports section except this is a count by user versus count by domain.
	Files Owned	Count by User of the following files: • Files Uploaded • Google Docs • Google Sheets • Google Slides • Google Forms • Google Drawings.
Account Activity	User Account Status	Record by user instead of by domain that covers the same attributes as the Security/General Report data, the Security/GMail data, and the Security and Apps Usage Drive data.
Audit Reports	Admin Activity	Log-level details of authorized administrator actions, including the following information: • Event Name • Event Description • Administrator Name • Data • IP Address.
	Login Report	End User Login Report that lists the following information: • User Name • Event Name - Failed Login - Login Challenge - Logout - Successful Login

Category	Report	Information Fields
		Suspicious Login • IP Address • Date • Login Type - Google Password - SAML
	Drive Audit Log ¹	A Drive activity report that lists the following information: • Document Title • Document ID • Document Type • Event Name • Event Description - Add to Folder - Create - Delete - Download - Edit - Move - Preview - Print - Remove from Folder - Rename - Restore - Trash - Upload - View - Editor Settings Change - Link Sharing Access Type Change • Date • Owner • User • IP Address.
	Calendar Audit Log	A Calendar activity report that lists the following information: • Activity Name - Access Level Change - Country Change - Calendar created - Calendar deleted - Description changed - Location changed - Time zone changed - Title changed - Subscription added or deleted - Event created - Event deleted - Guest added

¹ Available only with Google Apps Unlimited

Category	Report	Information Fields
		Guest removed Guest response change Event modified Event removed from trash Event restored Start time changed Title changed • User • Calendar ID • Date • Event Title • Event ID • API Kind • User Agent • IP Address.
	Token Report	A log of the following information on user activities tied to authorizing OAuth tokens for third-party apps: • Event Name Authorize Revoke • Application Name • User • Scope • Date • Client ID • IP Address.
	Groups Log	A log of the following information on user activities related to Google Groups for Business: • Event Name Group permission changed User accepted invitation to a group Owner approved request to join from Another user User joined a group User requested to join a group Basic settings changed Group created Group deleted Identity setting changed Information setting changed Information setting removed New member restriction changed Post replies setting changed Spam moderation setting changed Topic setting changed Message moderated Post from users will always be posted User added to group User banned from group User invitation revoked User invited

Category	Report	Information Fields
		Join request rejected User reinvited User removed from group - Event Description Includes the names of user and Group name - User - Date
	Email Log Search	A query engine to look up log detail on email delivery by selecting the following attributes: - Date - Sender - Recipient - Sender IP - Recipient IP - Message ID
Alerts	Predefined Alerts	Drive Settings Change TLS Failure User Deleted User Granted Admin Permission Exchange Journal Failure Suspended User made Active User Admin Privilege revoked Smarthost Failure GMail Settings Change User's Password Changed User Suspended Suspicious Login Activity Mobile Settings Change Calendar Settings Change Apps Outage Alert New User Added
	Custom Alerts	Custom alerts can be created against the following audit stream types: - Admin Audit - Token Audit - Calendar Audit - Login Audit - Drive Audit.

8.6.7 Offeror must describe whether it can restrict visibility of cloud hosted data and documents to specific users or groups.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys can restrict visibility of cloud hosted data and documents to specific users or groups. Unisys can accomplish this in multiple ways:

Unisys Stealth: The Unisys Stealth™ software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. By substituting traditional hardware topology for software-based cryptography, our Stealth Micro segmentation solutions prevent unauthorized access to sensitive information and reduce the

attack surface, thereby making endpoints invisible to certain users or user groups. Unisys Stealth tightens access control by focusing on user identity rather than physical devices, so security moves with the user and is easier to manage. It protects sensitive data in motion from potential compromise through encryption.

AD/LDAP: Unisys has in past achieved data access restriction to specific users or groups by synchronizing a user database with an external Lightweight Directory Access Protocol (LDAP) or Active Directory (AD). User roles and groups can be maintained in LDAP and AD which will authenticate and authorize each user based on its privileges. Access rights are determined by role. System administrators can create their own roles and map them to groups to suit a customer's business needs. System administrators can also define permissions for each role at a very granular level.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS Identity and Access Management (IAM) is a web service that enables AWS customers to manage users and user permissions in AWS. The service is targeted at organizations with many users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console. With AWS IAM, system administrators can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.

Permissions let system administrators specify who has access to AWS resources and which actions they can perform on those resources. Every AWS Identity and Access Management (IAM) user starts with no permissions. In other words, by default, users can do nothing—not even view their own access keys. To give a user permission to do something, system administrators can add the permission to the user (attach a policy to the user), or add the user to a group that has the desired permission.

Exhibit 9 shows the AWS IAM tab on the AWS Management Console.

Exhibit 9. AWS IAM Tab.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft Azure can restrict visibility of cloud-hosted data and documents to specific users or groups. Customer environments and data in Azure are isolated using numerous mechanisms, technologies, policies, processes, and architectural elements. Several of them are listed as follows:

- Virtual Networks – customer tenants and virtual machine deployments are kept logically separated through VNets that define DNS, security policies, and IP routing rules. Firewalls, ACLs, network security groups, IP filters, virtual appliances, load balancers, and network policies work together to prevent unauthorized traffic from entering or leaving a customer's tenant, either across network boundaries or between the virtualization host and guest.
- Encryption – customer data is encrypted in transit and at rest through configurable and standards-based providers using a variety of protocols such as BitLocker, AES-256 (in Azure Media Services), and IPsec (VNets).
- Access Control – Azure Storage, the Azure Portal, and other service components provide role-based access controls and key-based authentication to help confirm that only authorized entities can gain access to tenant data.

The concept of tenant containers is maintained in the Azure Active Directory service at several layers, from portals to persistent storage. These boundaries prevent a query scoped to a given tenant from returning directory data for another tenant. Front ends (Azure AD Sync, PowerShell, and Graph) store and retrieve data through an internal directory services API (DSAPI) that calls an authorization layer to confirm the data requested is allowed for the user requesting the data. These capabilities are available to customers for isolating and protecting their data, gaining access to only their data and no others.

Azure Resource Management RBAC roles have support for Azure Service Management API (Classic) resources using the following RBAC roles:

- Classic Network Contributor
- Classic Storage Contributor
- Classic Virtual Machine Contributor

Using these RBAC roles, it is possible to assign limited access to classic resources in the ARM Azure portal. The access is restricted to the abilities in the ARM portal for management of the resources.

RBAC is supported on classic Compute, Storage, and Networking objects. Compute includes IaaS VMs and PaaS Web/Worker roles. Networking includes vNets and subnets (NSGs are currently not supported). Storage includes storage accounts. Only classic resources in these three roles are supported.

Azure Resource Manager provides the ability to restrict operations on resources through resource management locks. Locks are policies which enforce a lock level at a particular scope. The scope can be a subscription, resource group or resource.

Google Public and Community Cloud for SaaS:

Restricting the visibility of cloud-hosted data and documents is mainly the customer's responsibility. As described in the collective response to RFP requirement 8.6.4, Google is obliged to confirm there is no unauthorized access of customer data. Customers must confirm that their end users use the service according to their acceptable use policies. During the onboarding process, customers will be assisted in the configuration controls that are included to enforce acceptable use and security policies and trained on how to oversee the ongoing use of the services in scope.

8.6.8 Offeror must describe its notification process in the event of a security incident, including relating to timing, incident levels. Offeror should take into consideration that Purchasing Entities may have different notification requirements based on applicable laws and the categorization type of the data being processed or stored.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

The Unisys proposed solution includes HP ArcSight to provide a robust SIEM security platform that the Unisys Security Officer in the Project Management Office (PMO) will use to monitor security incidents in the hosting services. To offer comprehensive cross-enterprise protection, our proposed solution converges monitoring, management, compliance, analytics, and response mechanisms for enhanced IT security. By unifying Security Event Management (SEM) and Security Information Management (SIM) capabilities, SIEM provides greater visibility to IT security; enhances protection of information, assets, and processes; and helps organizations to comply with GNU Privacy Guard (GPG)-13 (Protective Monitoring) and GPG-18 (Forensic Readiness). The Unisys SIEM solution leverages our best-in-class analytics engine based on ArcSight, the Unisys Noise Cancellation Advanced Analytics Platform (UNCAAP), and the PMO-dedicated Security Officer to minimize false alarm rates and offer advanced forensics and analytics with benefits of actionable intelligence, decision support, and security strategy insights.

With new application services and process improvement, the configuration of the rule base that operates the Unisys proposed solution will be aligned with the change management process, confirming the appropriate level of monitoring is used for each monitored system. Unisys will also work with customer proprietary network information (CPNI) to monitor current threat levels in the industry and gain early intelligence on potential upcoming attack vectors and changes to best practice as driven by security professionals who work in IT. Together with an effective communication strategy, technical and procedural services will provide defense in depth across the entire hosted service.

Unisys will use the Qualys Vulnerability Scanner to monitor the IP addresses in this engagement's scope at the mutually agreed intervals. The preferred method of scanning an environment is to perform a discovery scan before completing the vulnerability scan. This type of scan will provide an accounting of systems in a specific IP range and show what ports are open, what systems were discovered, and what services may be running on those ports. Typically, a nonintrusive scan will not have adverse effects on a system, though a slight risk does exist. This risk is higher when running an intrusive scan. System owners need to understand the risks that exist before the scan is run. The mutually agreed change control processes will be followed for systems that will be scanned.

After a scan is completed, analyzed, and published, the vulnerabilities must be remediated. This effort will be led by the Security Officer, who oversees the vulnerability scanning effort. The Security Officer should use the scan reports to create tickets for risks related to a system. The Security Officer will create tickets for vulnerabilities based on priority (Critical, High, Medium, etc.). The tickets will be routed to the Unisys Support team for remediation in the timelines agreed with the State.

A facility that allows for exemptions to remediate the identified vulnerability should be in place to account for situations to which one of the following factors apply:

- System stability
- Lack of an available patch for the operating system or application involved
- Any other business reason.

Unisys will also take over and manage the existing server antivirus infrastructure to reduce disruption to the State's transition to a new security provider model in the SIEM.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: AWS implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the AWS customer support team to be notified of operational issues that affect the customer experience. A Service Health Dashboard is available and maintained by the AWS customer support team to alert customers to issues that may be of broad impact. The AWS Security Center is available to provide customers with security and compliance details on AWS. Customers can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to issues that affect customers.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

If Microsoft becomes aware of unlawful access to customer data stored on its equipment or in its facilities, or unauthorized access to this equipment or facilities resulting in loss, disclosure, or alteration of customer data (each a security incident), Microsoft will promptly (1) notify the customer of the security incident; (2) investigate the security incident and provide the customer with detailed information on the security incident; and (3) take reasonable steps to mitigate the effects and to minimize damage resulting from the security incident.

Notifications of security incidents will be delivered to one or more of the customer's administrators by the means Microsoft selects, including email. It is the customer's sole responsibility to confirm that its administrators maintain accurate contact information on each applicable Online Services portal. Microsoft's obligation to report or respond to a security incident in this section is not an acknowledgement by Microsoft of fault or liability for the security incident.

The customer must notify Microsoft promptly of possible misuse of its accounts or authentication credentials or security incidents related to an Online Service.

Electronic Notices

Microsoft may provide the customer with information and notices about online services electronically, including by email, at the portal for the online service, or at a website that Microsoft identifies. Notice is given as of the date that Microsoft makes the information and notices available.

Incident Response Process

Microsoft maintains a record of security breaches with a description of the breach, the period, the consequences of the breach, the name of the reporter, to whom the breach was reported, and the procedure for recovering data.

For each security breach that is a security incident, notification by Microsoft (as described in "Electronic Notices") will be made without unreasonable delay and within 30 calendar days.

Microsoft tracks, or enables the customer to track, disclosures of customer data, including what data was disclosed, to whom, and at what time.

Microsoft security personnel verify logs at least every 6 months to propose remediation efforts if necessary.

Microsoft Security Development Lifecycle (SDL)

Microsoft SDL is supported in Azure; part of that development processes is notification. Azure has security controls in place to implement threat mitigation and to assist customers with mitigating potential threats in their environments. To detect threats and prevent exploits, Microsoft maintains continuous monitoring across servers, networks, and applications. Automated alerts notify administrators of anomalous behaviors, allowing them to take corrective action on internal and external threats.

Microsoft Operations Centers (MOCs) are globally distributed and work around the clock in a “follow-the-sun” model to enable Microsoft’s cloud services to be persistently available. Each MOC is staffed with a team of incident management professionals; collectively they monitor service health, process automation, infrastructure operations, event and crisis management, and communications across the business. They are responsible for more than 500 service components and monitor the servers and devices for the services Microsoft provides. Most critically, MOCs identify and resolve service incidents and outages when things go wrong.

Google Public and Community Cloud for SaaS: If Google becomes aware of a data incident, Google will promptly notify the customer of the incident and take reasonable steps to minimize harm and secure customer data. Notifications of data incidents will be sent to the notification email address provided by the customer in connection with the agreement or, at Google’s discretion, by direct communication (e.g., by a phone call or at an in-person meeting). The customer must confirm the contact information for the Notification Email Address is current and valid and must establish third-party notification obligations. Data incidents do not include unsuccessful access attempts or similar events that do not compromise the security or privacy of customer data, including pings, port scans, denial-of-service attacks, and other network attacks on firewalls or networked systems. Data incidents also do not include accidental loss or disclosure of customer data caused by the customer’s use of the services or the customer’s loss of account authentication credentials. Google’s obligation to report or respond to a data incident will not be construed as an acknowledgement by Google of faults or liabilities for the data incident.

8.6.9 Offeror must describe and identify whether or not it has any security controls, both physical and virtual Zones of Control Architectures (ZOCA), used to isolate hosted servers.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys provides security controls through its Stealth Offering:

The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures.. By substituting traditional hardware topology for software-based cryptography, our Stealth Micro segmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making endpoints invisible to unauthorized users. Unisys Stealth conceals endpoints making them undetectable to unauthorized parties inside and outside the enterprise. It tightens access control by focusing on user identity rather than physical devices, so security moves with the user and is easier to manage. Stealth protects sensitive data in motion from potential compromise through encryption and reduces costs by allowing you to consolidate and virtualize networks, servers, and cloud architectures.

The Unisys Stealth software has a legacy of being "Ready for Government" having received the National Information Assurance Partnership's (NIAP's) coveted Evaluation Assurance Level 4+ (EAL4+) certification, a Common Criteria international standard in effect since 1999. Stealth has been sold and

installed in US DoD and allied defense forces, as well as a variety of local and federal governments and commercial entities around the world.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Amazon Virtual Private Cloud (Amazon VPC) lets customers provision a logically isolated section of the Amazon Web Services (AWS) Cloud, where customers can launch AWS resources in a virtual network that they define. Customers have complete control over their virtual networking environment, including selection of their own IP address range, creation of subnets, and configuration of route tables and network gateways.

Customers can easily customize the network configuration for their Amazon Virtual Private Cloud. For example, they can create a public-facing subnet for their web servers that has access to the Internet, and move their backend systems such as databases or application servers to a private subnet with no Internet access. Customers can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

It is important for customers to realize the scale of engineering in the Microsoft Azure strategy. The designs are not for a single data center or a few dozen data centers, so the zoning must be completed at a hyperscale level. With more than 100 data centers worldwide, Microsoft has data centers in every region, connected by one of the largest cloud networks in the world. Only a few cloud providers are individually enterprise grade, hybrid, or hyperscale; Microsoft's cloud is the only one that offers all three.

The distributed and virtual networks in Azure help to isolate each customer's private network traffic logically from traffic belonging to other customers. A customer subscription can contain two types of isolated private network (and include firewall, load-balancing, and network address translation):

- **Deployment network:** Each deployment is isolated from other deployments at the network level. Several virtual machines in a deployment are allowed to communicate with each other at private IP addresses.
- **Virtual network:** Each virtual network is isolated from other virtual networks. Several deployments (inside the same subscription) can be moved to the same VNET and allowed to communicate at private IP addresses.

By default, virtual machines inside the private network do not receive inbound traffic from outside the deployment. The administrator defines an input end point that specifies which ports on which virtual machines should receive inbound traffic initiated from outside a deployment's isolated network—enabling traffic from the Internet and other deployments or customers inside Azure.

Microsoft Azure uses many safeguards to protect customer and enterprise data. These security practices and technologies include the following:

- **Identity and access management** – Azure Active Directory helps to enable only authorized users to access customers' environments, data, and applications; it provides multifactor authentication for highly secure sign-in.
- **Encryption** – Azure uses industry-standard protocols to encrypt data as it travels between devices and Microsoft data centers, and crosses in data centers
- **Secure networks** – Azure infrastructure relies on security practices and technologies to connect virtual machines to each other and to on-premises data centers while blocking unauthorized traffic. Azure Virtual Networks extend a customer's on-premises network to the cloud over a site-to-site virtual private network (VPN). Customers can also use ExpressRoute to create a cross-premises connection when they have to use the Internet.
- **Threat management** – Microsoft Antimalware protects Azure services and virtual machines. Microsoft also uses intrusion detection, denial-of-service (DoS) attack prevention, penetration testing, data analytics, and machine learning to constantly strengthen its defense and reduce risks.

- Compliance – Microsoft complies with international and industry-specific compliance standards and participates in rigorous third-party audits that verify Microsoft’s security controls.

Physical zoning and virtual zoning of host systems are supported through Microsoft Azure’s extensive use of software-defined networks that span across the data center from edge computing, not to the interval fiber backbones.

Physical access to facilities. Microsoft limits access to facilities that house information systems that process customer data to identified authorized individuals.

Physical access to components. Microsoft maintains records of incoming and outgoing media containing customer data, including the kind of media, the authorized sender and recipients, date and time, the number of media, and the types of customer data they contain.

Protection from disruption. Microsoft uses a variety of industry standard systems to protect against loss of data due to power supply failure or line interference.

Component disposal. Microsoft uses industry-standard processes to delete customer data when it is no longer needed.

With software enablement for restricting access, security is crucial to provide customer data and privacy. Microsoft goes beyond ISO and NIST standards to maintain physical and virtually safe environments. In the Azure infrastructure, for example, CPUs, storage, and SQL services are separated to support greater zoning and security requirements.

Software-defined Networks (to support zoning)

To improve flexibility and accelerate the adoption of advanced technologies into its network, Microsoft broadly adopted software-defined networking (SDN).

SDN provides the ability to dynamically modify Microsoft’s network using automated management tools to move data and resources to an area where they are best served. In an SDN environment, Microsoft can extract and separate the application, the control plane, and the transport of the data. This allows Microsoft to insert its own APIs to gain visibility of how the data flows and gain better control, and also allows Microsoft to upgrade network performance outside the hardware refresh cycle. Microsoft’s large, geographically distributed footprint of data centers and networks enables Microsoft to be located close to its customers to reduce network latency and allow for georedundant backup and failover.

Google Public and Community Cloud for SaaS:

Google’s focus on security and protection of data is among its primary design criteria. Google data center physical security features a layered security model, including safeguards like custom-designed electronic access cards, alarms, vehicle access barriers, perimeter fencing, metal detectors, and biometrics. The data center floor features laser-beam intrusion detection. Google’s data centers are monitored 24x7 by high-resolution interior and exterior cameras that can detect and track intruders. Access logs, activity records, and camera footage are available if an incident occurs. Data centers are also routinely patrolled by experienced security guards who underwent rigorous background checks and training. As visitors come closer to the data center floor, security measures increase. Access to the data center floor is possible only in a security corridor that implements multifactor access control using security badges and biometrics. Only approved employees with specific roles may enter. Less than one percent of visitors will ever step into a Google data center.

Google operates a global, multitenant environment running the world’s second largest IP data network providing customers with a low-latency, high-performing platform that runs 24x7.

Hard disks are assets that are tracked throughout their life cycle at Google from arrival to final destruction. These disks are components that Google uses to build its own servers from other component parts, including motherboards and a hardened, highly customized version of Linux.

Google uses a proprietary storage and processing mechanism that isolates processing of data in change-rooted (chrooted) jails in a physical server. Access permissions restrict the ability for processes to interact between jails. Additionally, tenant data is stripped in chunks across many different drives; each chunk has its own access control list. This helps to confirm that data is logically isolated between customers in storage and during processing.

In addition to the processing and storage mechanisms already described, Google security controls, including least privilege rights, logging, and auditing, was implemented to be consistent with FedRAMP requirements.

8.6.10 Provide Security Technical Reference Architectures that support Infrastructure as a Service (IaaS), Software as a Service (SaaS) & Platform as a Service (PaaS).

Illustrated below are several examples from all the providers in this RFP. The exhibit titles are self-explanatory.

Exhibit 10 is an example Unisys hybrid cloud model that is applicable across all three cloud services models.



Exhibit 11 is an example Unisys hybrid cloud model that is more specific to IaaS.

Establishing a Platform that enable Client Centric IaaS Solutions and a Cloud Ecosystem




Exhibit 12 below is an example view of the full-scale hybrid enterprise environment

Integrating Clients Cloud Environment within the context of the full IT Landscape




Exhibit 13 is an example of highly available, fault-tolerant reference architecture for IaaS on AWS.





Exhibit 13. AWS Fault Tolerance and High Availability.

Exhibit 14 is an example of reference architecture for SharePoint on AWS as PaaS.



ADVERTISEMENT SERVING



- When visitors load a web page, **serve** return a pointer to the ad resource to be displayed. These requests are handled by **Amazon EC2** instances (Amazon EC2 instances). They query a data set stored in an **Amazon DynamoDB** table to find relevant ads depending on the user profile.
- Ad files are downloaded from **Amazon CloudFront**, a content delivery service with low latency, high data-transfer speeds, and no commitments. Log information from **Amazon S3** is stored in **Amazon S3** using **Storage Service (Amazon S3)**, a highly available data store.
- The click-through servers are a group of **Amazon EC2** instances dedicated to collecting click-through data. This information is then stored in **Amazon S3** through web servers, which are periodically updated to **Amazon S3**.
- Ad impression and click-through data are retrieved and processed by **Amazon EMR**, an Elastic MapReduce, a managed Hadoop framework to process the data in a parallel job flow. The cluster's capacity can be dynamically scaled up or down to match the volume of data, reduce the processing time and the cost of running the job flow.
- Data processing results are pushed back into **Amazon DynamoDB**, a fully managed NoSQL database service that provides fast and predictable performance and scalability. **Amazon DynamoDB** tables can store and retrieve any amount of data, and serve any level of request traffic, with which are specific requirements for storing and quickly retrieving visitors' profile information.

Exhibit 16 is an example of reference architecture in the Windows Azure Cloud environment representing IaaS.




Exhibit 17 is an example of reference architecture for SharePoint in the Windows Azure Cloud environment representing PaaS.

SHAREPOINT 2013: APP OVERVIEW FOR IT PRO


APP AND SOLUTION COMPARISON

	CONTENTS	HOW TO GET	HOW TO USE/INSTALL	HOW TO MANAGE
APPS FOR SHAREPOINT Are stand-alone applications that provide specific information or functionality to a SharePoint site. Apps for SharePoint are easy for users to install, use, manage, upgrade, and remove. Apps can be SharePoint-hosted (reside and execute in SharePoint) or cloud-hosted (Windows Azure or other systems), or both.	APP (.app) 	Apps can be downloaded from: - Internal App Catalog (contains apps approved and uploaded by the organization) - Public SharePoint Store	Apps can be used: - In your hosted or on-premises SharePoint site App code is installed: - On a separate web site from your other sites in its own, isolated, domain - In the cloud (cloud-based apps)	Apps can be managed and monitored by: - Site administrators - SharePoint Online Service administrators for a tenancy (sandbox) - Farm administrators
SOLUTIONS Are small to large scale packages used to customize or enhance SharePoint sites. Full trust solutions need a farm or SharePoint Online Service administrator to deploy, manage, and remove. Partial trust solutions must be installed to a sandbox.	SOLUTION (.wsp) 	Solutions can be acquired from: - Third-party developers - Your own development team	Solutions can be used: - In your hosted (sandbox solutions only) or on-premises SharePoint site Solution code is installed: - As a full trust solution (in the global assembly cache) - As a sandboxed solution	Solutions can be installed and monitored by: - SharePoint Online Service administrators for a tenancy (sandbox) - Farm administrators (sandbox or full trust)

WHY USE APPS?


APP ISOLATION

Isolation prevents unauthorized access to users' data. The isolation level determines the limit to potential breaches in security. By default, apps are deployed to their own web site in a special, isolated domain, rather than to your farm or a sandbox. Processes run in that domain.



HOSTING OPTIONS

Apps for SharePoint can be hosted by a provider, autohosted by Windows Azure, hosted by SharePoint, or a combination of these. All can leverage SharePoint components. Custom code can only be run in provider-hosted or Windows Azure autohosted options. The following hosting options are supported:



MONITOR AND MANAGE APPS

Farm administrators and SharePoint Online Service Administrators can monitor apps for SharePoint and respond to errors and issues. Site owners can manage apps for their sites.

Task	Site level	SharePoint Online Tenancy level	Farm level
Add, delete and view app details Add or delete an app in a site View details about an app	✓	✓	✓
App Catalog Configure the App Catalog Manage the App Catalog		✓	✓
Monitor apps Specify apps to monitor View install locations, manage resources		✓	✓
Errors View and troubleshoot errors		✓	✓
Licenses View and manage licenses		✓	✓

To monitor and manage apps:

- Farm administrators: General Application Settings in Central Administration
- SharePoint Online Service Administrators: Apps pages in SharePoint Online Administration Center
- Site owners: All Site Content page

APP LIFECYCLE


INSTALL AND UPDATE

Users can add and update apps for SharePoint themselves. Installation and update/upgrade can happen at any time, initiated by the user. Users can get apps for SharePoint from the SharePoint Store, or from an App Catalog that is set up for their organization.

For on-premises deployments, IT can determine whether users can download and install apps for SharePoint and can restrict access to the SharePoint Store.


FIND

A user finds an app in the SharePoint Store or App Catalog.




INSTALL

A user adds an app to his or her site. Other users do the same.



UPDATE (DEV)

The app developer releases an update to an app.



UPGRADE

Users are notified of an update and decide whether to upgrade the app on their sites.




UNINSTALL AND DISABLE/REMOVE

Users can uninstall an app completely, which removes it from their sites.

The SharePoint Store administrators can disable an app that is unsafe and remove it from the Store.


UNINSTALL

A user decides not to continue using the app and uninstalls it. The app instance and all data is deleted.




DISABLE/REMOVE (STORE)

SharePoint Store administrators discover an issue with the app and want to remove it.



Store administrators disable and remove the app. Data from the app is retained for recovery purposes.



© 2012 Microsoft Corporation. All rights reserved. To send feedback about this documentation, please write to us at ITSPDocs@microsoft.com.



Exhibit 17. Reference Architecture for SharePoint in an Azure Cloud Environment Representing PaaS.

Exhibit 18 is an example of reference architecture for Office 365 on the Windows Azure Cloud Environment representing SaaS.





Exhibit 18. Reference Architecture for SharePoint in an Azure Cloud Environment Representing SaaS.

Exhibit 19 is an example of Google Apps reference architecture

8.6.11 Describe security procedures (background checks, foot printing logging, etc.) which are in place regarding Offoror's employees who have access to sensitive data.

Unisys Response:

Unisys and our partners mandate background verification for our associates to comply with laws of various countries. New associates are subject to a security check and verification of previous employment before they begin work, which includes a past employment check; an education check; Federal and state criminal checks; and, if required, drug testing for associates. For our cleared associates and associates working in the government space, additional screening extending to a Federal clearance is performed in accordance with the requirement descriptions.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS operates under a shared security responsibility model, where AWS is responsible for the security of the underlying cloud infrastructure and customers are responsible for securing the workloads they deploy in AWS. AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS conducts pre-employment criminal background checks, as permitted by law, for employees commensurate with their position and level of access. The AWS SOC reports provides additional details regarding the controls in place for background verification.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Pursuant to local laws, regulations, ethics and contractual constraints, Microsoft full-time employees based in the United States are required to successfully complete a standard background check as part of the hiring process. Background checks may include reviewing information on a candidate's education, employment, and criminal history.

Google Public and Community Cloud for SaaS:

Google personnel are required to conduct themselves in a manner consistent with the company's guidelines regarding confidentiality, business ethics, appropriate usage, and professional standards. Google conducts reasonably appropriate background checks to the extent legally permissible and in accordance with applicable local labor law and statutory regulations. Personnel are required to execute a confidentiality agreement and must acknowledge receipt of, and compliance with, Google's confidentiality and privacy policies. Personnel are provided with security training. Personnel handling Customer Data are required to complete additional requirements appropriate to their role (eg certifications). Google's personnel will not process Customer Data without authorization.

8.6.12 Describe the security measures and standards (i.e. NIST) which the Offeror has in place to secure the confidentiality of data at rest and in transit.

Unisys Response:

Unisys adequately secures database systems from security threats internally and externally. We established a workflow mechanism that allows specific queries to be run in the database. Only administrators have the control to specify these queries, thereby confirming that even trusted users have access to limited features of the database. Database systems are adequately hardened from threats such as SQL injections and cross-site scripting (XSS) attacks. Connections to databases are controlled using appropriate features in the database.

To achieve and maintain security across our networks, Unisys uses network controls to separate various client networks and the public network as appropriate through switch and firewall technologies. We established an access policy to govern associates who have access to networks and services. To encrypt network traffic, we use the latest cryptographic controls.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS offers you the ability to add a layer of security to your data at rest in the cloud, providing scalable and efficient encryption features. These include:

- Data encryption capabilities available in AWS storage and database services, such as EBS, S3, Glacier, Oracle RDS, SQL Server RDS, and Redshift
- Flexible key management options that allow you to choose whether to have AWS manage the encryption keys or maintain complete control over your keys
- Dedicated, hardware-based cryptographic key storage options for customers to help satisfy compliance requirements

In addition, AWS provides APIs for you to integrate encryption and data protection with any of the services you develop or deploy in an AWS environment.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft encrypts customer data at transit and at rest and provides a platform that complies with NIST standards. Industry and government regulations such as HIPAA and FedRAMP as well as international standards such as ISO 27001 lay out specific safeguards through processes and policies. Microsoft Azure and its customers share a responsibility to implement sufficient mechanisms to meet those obligations. Specifically, Microsoft provides a compliant platform for services, applications, and data; Azure customers must design and configure their cloud environment to maintain the confidentiality and integrity of their information assets.

Microsoft Azure contains many tools to safeguard data according to a customer's security and compliance needs. One of the keys to data protection in the cloud is accounting for the possible states in which data may occur, and what controls are available for that state. Specifically, this accounting is for the following states of customer data:

- At rest: This includes information storage objects, containers, and types that exist statically on physical media, be it magnetic or optical disk.
- In transit: When data is transferred between components, locations, or programs, such as over the network, across a service bus (from on-premises to cloud and vice versa, including hybrid connections such as ExpressRoute), or during an input/output process, it is thought of as being in motion. Data in transit does not always mean a communications process with a component outside a cloud service; data moves internally, also (e.g., between two virtual networks).
- In use: (or in process): Dynamic data use could be on a table kept in virtual memory, transactions in a message queue, or encryption keys in the CPU cache. Information acted on in some way by the host or guest during a process, such as real-time database queries running in active memory (instead of a page file sent out to disk), could be in different security states, depending on whether it is encrypted or decrypted, and the operator's security context.

Additionally, there are two fundamental types of data at rest:

- Data in production. Data is in a form of storage, (e.g., Azure SQL Database), and compute processes need to access that storage during production operations. For this data, encryption at rest aims to protect the data in that storage (whereas the compute aspect deals with data in use).
- Data not in production. Data is in a form of storage (e.g., a virtual hard disk [VHD], but that VHD is not in production use. For example, the data may be part of an upgrade operation, but the VHD is not yet loaded or mounted. Data encryption at rest is applicable here, but the compute aspect is not relevant to this scenario.

Google Public and Community Cloud for SaaS:

Google stores data in a multi-tenant environment on Google-owned servers. Data, the Services database and file system architecture are replicated between multiple geographically dispersed data centers. Google logically isolates data on a per End User basis at the application layer. Google logically isolates each Customer's data, and logically separates each End User's data from the data of other End Users, and data for an authenticated End User will not be displayed to another End User (unless the former End User or an Administrator allows the data to be shared). A central authentication system is used across all Services to increase uniform security of data. The Customer will be given control over specific data sharing policies. Those policies, in accordance with the functionality of the Services, will enable Customer to determine the product sharing settings applicable to End Users for specific purposes. Customer may choose to make use of certain logging capability that Google may make available via the Services, products and APIs. Customer agrees that its use of the APIs is subject to the API Terms of Use. Google agrees that changes to the APIs will not result in the degradation of the overall security of the Services.

Data centers are typically connected via high-speed private links to provide secure and fast data transfer between data centers. This is designed to prevent data from being read, copied, altered or removed without authorization during electronic transfer or transport or while being recorded onto data storage media. Google transfers data via Internet standard protocols.

8.6.13 Describe policies and procedures regarding notification to both the State and the Cardholders of a data breach, as defined in this RFP, and the mitigation of such a breach.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys maintains a well-defined incident management framework that confirms that security events are given the utmost priority. The framework's design allows it to receive manual and automated notification on various incidents, further classifying them into levels of criticality until they are resolved and a root cause analysis is done. We use technologies such as SIEM and ArcSight for incident management. A Unisys Security Operations Center (SOC) continuously monitors security events 24x7 and correlates those events with log management solutions. We notify subscribers to our SOC services of security events and breaches and their progress toward resolution.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

AWS Customers retain the responsibility to monitor their own environment for privacy breaches.

AWS has implemented a formal, documented incident response policy and program (including instructions on how to report internal and external security incidents). The policy addresses purpose, scope, roles, responsibilities, and management commitment. Systems within AWS are extensively instrumented to monitor key operational and security metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key metrics. When a threshold is crossed, the AWS incident response process is initiated. The Amazon Incident Response team employs industry - standard diagnostic procedures to drive resolution during business - impacting events. Staff operates 24x7x365 coverage to detect incidents and manage the impact to resolution.

AWS utilizes a three-phased approach to manage incidents:

1. Activation and Notification Phase: Incidents for AWS begin with the detection of an event. This can come from several sources including:
 - Metrics and alarms - AWS maintains an exceptional situational awareness capability, most issues are rapidly detected from 24x7x365 monitoring and alarming of real time metrics and service dashboards. The majority of incidents are detected in this manner. AWS utilizes early indicator alarms to proactively identify issues that may ultimately impact Customers.
 - Trouble ticket entered by an AWS employee
 - Calls to the 24X7X365 technical support hotline. If the event meets incident criteria, then the relevant on -call support engineer will start an engagement utilizing AWS Event Management Tool system to start the engagement and page relevant program resolvers (e.g. Security team). The resolvers will perform an analysis of the incident to determine if additional resolvers should be engaged and to determine the approximate root cause.
2. Recovery Phase - the relevant resolvers will perform break fix to address the incident. Once troubleshooting, break fix and affected components are addressed, the call leader will assign next steps in terms of follow -up documentation and follow - up actions and end the call engagement.
3. Reconstitution Phase - Once the relevant fix activities are complete the call leader will declare that the recovery phase is complete. Post mortem and deep root cause analysis of the incident will be assigned to the relevant team. The results of the post mortem will be reviewed by relevant senior management and relevant actions such as design changes etc. will be captured in a Correction of Errors (COE) document and tracked to completion.

In addition to the internal communication mechanisms detailed above, AWS has also implemented various methods of external communication to support its customer base and community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "Service Health Dashboard" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact.

AWS incident management program reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft Identity Manager and intrusion detection system tools are implemented within the Azure environment. Azure uses an early warning system to support real-time analysis of security events within its operational environment. Monitoring agents and the alert and incident management system generate near real-time alerts about events that could potentially compromise the system.

The Azure Incident Response process follows five main phases:

Identification – System and security alerts may be harvested, correlated, and analyzed. Events are investigated by Microsoft operational and security organizations. If an event indicates a security issue, the incident is assigned a severity classification and appropriately escalated within Microsoft. This escalation will include product, security, and engineering specialists.

Containment – The escalation team evaluates the scope and impact of an incident. The immediate priority of the escalation team is to ensure the incident is contained and data is safe. The escalation team forms the response, performs appropriate testing, and implements changes. In the case where in-depth investigation is required, content is collected from the subject systems using best-of-breed forensic software and industry best practices.

Eradication – After the situation is contained, the escalation team moves toward eradicating any damage caused by the security breach, and identifies the root cause for why the security issue occurred. If a vulnerability is determined, the escalation team reports the issue to product engineering.

Recovery – During recovery, software or configuration updates are applied to the system and services are returned to a full working capacity.

Lessons Learned – Each security incident is analyzed to ensure the appropriate mitigations applied to protect against future reoccurrence.

Google Public and Community Cloud for SaaS:

If Google becomes aware of a Data Incident, Google will promptly notify the customer of the Data Incident, and take reasonable steps to minimize harm and secure Customer Data. Notification(s) of any Data Incident(s) will be delivered to the Notification Email Address provided by Customer in connection with the Agreement or, at Google's discretion, by direct communication (e.g., by phone call or an in-person meeting). Customer acknowledges that it is solely responsible for ensuring the contact information given for purposes of the Notification Email Address is current and valid, and for fulfilling any third party notification obligations. Customer agrees that "Data Incidents" do not include: (i) unsuccessful access attempts or similar events that do not compromise the security or privacy of Customer Data, including pings, port scans, denial of service attacks and other network attacks on firewalls or networked systems; or (ii) accidental loss or disclosure of Customer Data caused by Customer's use of the Services or Customer's loss of account authentication credentials. Google's obligation to report or respond to a Data Incident under this Section will not be construed as an acknowledgement by Google of any fault or liability with respect to the Data Incident.

8.7 (E) MIGRATION AND REDEPLOYMENT PLAN

8.7.1 Offeror must describe how it manages the end of life activities of closing down a service to a Purchasing Entity and safely deprovisioning it before the Offeror is no longer contractually obligated to maintain the service, include planned and unplanned activities. An Offeror's response should include detail on how an Offeror maintains security of the data during this phase of an SLA, if the Offeror provides for redundancy during migration, and how portable the data is during migration.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

The client is primarily responsible for creation, usage, maintenance and destruction of its content when they are subscribed to services offered within the Cloud Account that Unisys would have provisioned.. In case when the subscription has ended, Unisys will provide a moratorium time period until the customer can safely backup or migrate their existing data, application stack and workloads as required. Once customer has completed with his backup or migration activities, Unisys will deprovision the Cloud Account.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

You (the client) are solely responsible for the development, content, operation, maintenance, and use of Your Content. For example, you are solely responsible for the technical operation of Your Content, including ensuring that calls you make to any Service are compatible with then-current APIs for that

Service; compliance of Your Content with the Acceptable Use Policy, the other Policies, and the law; any claims relating to Your Content; and properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Content violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act.

During the 30 days following termination (of the service) we will not erase any of Your Content as a result of the termination. You may retrieve Your Content from the Services only if you have paid any charges for any post-termination use of the Service Offerings and all other amounts due. We will provide you with the same post-termination data retrieval assistance that we generally make available to all customers.

Any additional post-termination assistance from us is subject to mutual agreement by you and us.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

During a customer's subscription, the customer will be able to access and extract customer data stored in each online service. Except for free trials, Microsoft retains customer data stored in the online service in a limited function account for 90 days after expiration or termination of the customer's subscription so that the customer can extract the data. After the 90-day retention period ends, Microsoft disables the customer's account and deletes the customer data.

Google Public and Community Cloud for SaaS:

During the Term, Google will provide Customer with the ability to correct, block, export and delete Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. Once Customer deletes Customer Data via the Services such that the Customer Data cannot be recovered by Customer (the "Customer-Deleted Data"), Google will delete the Customer-Deleted Data within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so. On the expiry or termination of the Agreement (or, if applicable on expiry of any post-termination period during which Google may agree to continue providing access to the Services), after a recovery period of up to 30 days following such expiry or termination, Google will thereafter delete the Customer-Deleted Data within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

8.7.2 Offeror must describe how it intends to provide an orderly return of data back to the Purchasing Entity, include any description in your SLA that describes the return of data to a customer.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

Unisys provides client the ability to download or migrate their data and workloads as and when a service has been terminated. Unisys also provides Cloud Application Migration Services that starts with a complete and analytical view of your application portfolio and cloud/data center capabilities. Using a set of criteria refined during customer engagements, Unisys can help you narrow the application pool of various cloud portable candidates and then provide the detailed assessment of costs and other factors that go into a cloud migration decision. Once the decision has been made, Unisys can help you with the planning and implementation of the application migration process including on-boarding of applications to AWS and other clouds.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

Amazon offers a variety of ways to transfer data in and out of AWS.

Direct Connect:

You can use a dedicated physical connection with AWS Direct Connect to accelerate network transfers between your datacenters and ours.

Using Direct Connect, you can establish a dedicated network connection between your network and one of the Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be

partitioned into multiple virtual interfaces. This approach means you can use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon Elastic Compute Cloud (EC2) instances running within an Amazon Virtual Private Cloud (Amazon VPC) using private IP space. At the same time, you can maintain network separation between the public and private environments. You can reconfigure virtual interfaces at any time to meet your changing needs.

AWS Import/Export Snowball:

Snowball is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of AWS. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, and secure, and can cost as little as one-fifth of the expense of using high-speed Internet for the same transfer.

AWS Import/Export Disk:

AWS Import/Export Disk accelerates moving large amounts of data into and out of the AWS cloud by using portable storage devices for transport. Disk transfers your data directly onto and off of your storage devices using the Amazon high-speed internal network and bypassing the Internet. For gigabyte-scale data sets that would travel at less than 10 Mbps, Disk is often faster than Internet transfer and more cost-effective than upgrading your connectivity.

Storage Gateways:

The AWS Storage Gateway service simplifies on-premises adoption of AWS storage. Your existing applications use industry-standard storage protocols to connect to a software appliance which stores your data in Amazon S3 and Amazon Glacier. With Storage Gateway, Data is compressed and securely transferred to AWS. SAN configurations offer stored or cached devices with point-in-time backups as Amazon Elastic Block Store (EBS) snapshots. Virtual tape library (VTL) configuration works with your existing backup software for cost-effective backup in Amazon S3 and long-term archival in Amazon Glacier.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: A customer's use of an Import/Export Service is conditioned on its compliance with the instructions provided by Microsoft for the preparation, treatment, and shipment of physical media containing its data (storage media). The customer is solely responsible for verifying the storage media and data are provided in compliance with laws and regulations. Microsoft has no duty toward the storage media and no liability for lost, damaged, or destroyed storage media. Storage media shipped to Microsoft must be shipped DAP Microsoft DCS Data Center (INCOTERMS 2010). Storage media shipped to the customer is shipped DAP Customer Dock (INCOTERMS 2010).

Google Public and Community Cloud for SaaS:

During the Term, Google will make available to Customer the Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. To the extent Customer, in its use and administration of the Services during the Term, does not have the ability to amend or delete Customer Data (as required by applicable law), or migrate Customer Data to another system or service provider, Google will, at Customer's reasonable expense, comply with any reasonable requests by Customer to assist in facilitating such actions to the extent Google is legally permitted to do so and has reasonable access to the relevant Customer Data.

8.8 (E) SERVICE OR DATA RECOVERY

8.8.1 Describe how you would respond to the following situations; include any contingency plan or policy.

- a. Extended downtime.**
- b. Suffers an unrecoverable loss of data.**

- c. *Offeror experiences a system failure.*
- d. *Ability to recover and restore data within 4 business hours in the event of a severe system outage.*
- e. *Describe your Recovery Point Objective (RPO) and Recovery Time Objective (RTO).*
- 8.8.2 *Describe your methodologies for the following backup and restore services:*
 - a. *Method of data backups*
 - b. *Method of server image backups*
 - c. *Digital location of backup storage (secondary storage, tape, etc.)*
 - d. *Alternate data center strategies for primary data centers in the continental United States.*

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS:

The advanced options of the Unisys Cloud Provisioning System allow users to gain significantly higher levels of default capability for backup and recovery operation.

Automation includes automatic detection of server failures. In a server failure, the Unisys Cloud Provisioning System automatically launches a new server based on the machine image configured for the server's group. When the replacement server comes on line, the Unisys Cloud Provisioning System reattaches volumes that were orphaned during the failure and restarts installed services.

For the Unisys Cloud Provisioning System, backups of participant applications and data are automated at two different levels. First, the Unisys Cloud Provisioning System calls scripts to run backups and uploads the backups to a cloud storage provider such as Amazon S3. For example, the State's scripts might create a compressed backup of its MySQL or SQL Server database. Second, in the Amazon cloud, the Unisys Cloud Provisioning System creates a complete file system backup by taking snapshots of attached EBS volumes. Both types of backups are run at custom intervals. Snapshots of images and data are captured on a user-defined schedule, can be moved to supported third-party cloud storage for increased disaster avoidance, and can be encrypted for increased security.

The Unisys Cloud Provisioning System also handles the creation and configuration of firewalls for the layers of a deployment. To create firewall rules, the Unisys Cloud Provisioning System opens access to database servers from application servers on the private IP address of the application servers. If load balancing is used, the application servers are reachable only through the load balancer to their respective private IP addresses.

The Unisys Cloud Provisioning System automatically scales the number of servers in a server group up or down according to a client's parameters. The Unisys Cloud Provisioning System periodically polls the health of the deployed services and reports the status inside the Unisys Cloud Provisioning System console.

Any time there is a state change (loss of server, backup failure, etc.), the Unisys Cloud Provisioning System sends alerts at the console and by email.

AWS Public and Community Cloud for IaaS, PaaS, and SaaS:

The AWS platform enables a lightweight approach to backup and recovery due in part to the following characteristics.

- Computers are now virtual abstract resources instantiated by code; they are not hardware based.
- Capacity is available at an incremental cost instead of an upfront cost.
- Resource provisioning takes place in minutes, lending itself to real-time configuration.
- Server images are available on demand, can be maintained by an organization, and can be activated immediately.

These characteristics offer customers opportunities to recover deleted or corrupted data with less infrastructure overhead.

Protecting Configurations Instead of Servers


The Amazon Elastic Compute Cloud (Amazon EC2) service enables the backup and recovery of a standard server, such as a web server or an application server, so that customers can focus on protecting their configuration and the state of data instead of the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, and temporary files. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an Amazon Machine Image (AMI) and can connect to existing storage volumes; for example, Amazon Elastic Block Store (Amazon EBS). Additionally, when launching a new instance, it is possible to pass user data to the instance that can be accessed internally as dynamic configuration parameters.


A sample workflow is as follows.

- Launch a new instance of a web server, passing to it the identity of the web server and the security credentials required for initial setup. The instance is based upon a prebuilt AMI that contains the operating system and relevant web server application (e.g., Apache or IIS).
- Upon startup, a boot script accesses a designated and secured Amazon Simple Storage Service (Amazon S3) bucket that contains the specified configuration files.
- The configuration file contains various instructions for setting up the server (e.g., web server parameters, locations of related servers, additional software to install, and patch updates).
- The server executes the specified configuration and is ready for service. An Open Source tool for performing this process called cloud-init is already installed on Amazon Linux AMIs and is available for several other Linux distributions.

Exhibit 20 depicts a traditional backup approach, and Exhibit 21 depicts an Amazon EC2 backup approach.

Exhibit 20. Traditional Backup Approach.

UT008

Exhibit 21 Amazon EC2 Backup Approach.

UT009

In the sample workflow, there is no need to back up the server itself. The relevant configuration is in the combination of the AMI and the configuration files. The only components that require backup and recovery are the AMI and configuration files.

Amazon Machine Image (AMI)

AMIs that customers register are automatically stored in their account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable. This means that the underlying storage mechanism for the AMIs is protected from multiple failure scenarios.

It is also possible to share AMIs between separate AWS accounts. Consequently, customers can create totally independent copies of the AMI by doing the following:

- Sharing the original AMI to another specified AWS account controlled by the customer
- Starting a new instance based upon the shared AMI
- Creating a new AMI from that running instance.

The new AMI is then stored in the second account and is an independent copy of the original AMI. Customers can also create multiple copies of the AMI in the same account.

Configuration Files

Customers use a variety of version management approaches for configuration files; they can follow the same regime for the files used to configure their Amazon EC2 instances. For example, a customer could store different versions of configuration files in designated locations and securely control them like other code. That customer could then use the appropriate backup cycle (daily, weekly, or monthly) to back up these code repositories and snapshots to protected locations. Additionally, customers can use Amazon S3 to store their configuration files, taking advantage of the service's durability, in addition to backing up the files regularly to an alternate location.

Database and File Servers

Backing up data for database and file servers differs from the web and application layers. Database and file servers contain larger amounts of business data (tens of gigabytes to multiple terabytes) that must be retained and always protected. For this data, customers can leverage efficient data movement techniques such as snapshots to create backups that are fast, reliable, and space efficient.

For databases that are built on RAID sets of Amazon EBS volumes (and have total storage of less than 1 TB), an alternate backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. Although the destination Amazon EBS volume will have slower performance, it is not being used for data access and can be easily snapshotted to Amazon S3 using the Amazon EBS snapshot capability.

Disaster Recovery

The AWS cloud supports many popular disaster recovery architectures from “pilot light” environments that are ready to scale up at a moment's notice to “hot standby” environments that enable rapid failover. With data centers in 12 regions around the world (4 in the United States), AWS provides a set of cloud-based disaster recovery services that enable rapid recovery of IT infrastructure and data.

General Disaster Recovery, COOP, and Backup Requirements and Issues

A traditional disaster recovery approach has some of the following needs and requirements:

- Facilities to house additional infrastructure, including power and cooling
- Security to maintain the physical protection of assets
- Suitable capacity to scale the environment
- Support for repairing, replacing, and refreshing the infrastructure

- Contractual agreements with an Internet Service Provider (ISP) to provide Internet connectivity that can sustain bandwidth utilization for the environment under a full load
- Network infrastructure such as firewalls, routers, switches, and load balancers
- Enough server capacity to run mission-critical services, including storage appliances for the supporting data, and servers to run applications and backend services such as user authentication, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), monitoring, and alerting.

AWS Capabilities for Disaster Recovery, COOP, and Backup Solutions

With AWS, customers can eliminate the need for additional physical infrastructure, offsite data replication, and upkeep of spare capacity. AWS uses distinct and geographically diverse Availability Zones (AZs) that are engineered to be isolated from failures in other AZs. This innovative and unique AWS feature enables customers to protect applications from the failure of a single location, resulting in significant cost savings and increased agility to change and optimize resources during a DR scenario.


AWS offers the following high-level disaster recovery capabilities:

- **Fast Performance:** Fast, disk-based storage and retrieval of files
- **No Tape:** Eliminated costs of transporting, storing, and retrieving tape media and associated tape backup software
- **Compliance:** Minimized downtime to avoid breaching Service Level Agreements (SLAs)
- **Elasticity:** The ability to add data quickly and delete expired data easily, without handling media
- **Security:** Secure and durable cloud disaster recovery platform with industry-recognized certifications and audits
- **Partners:** AWS solution providers and system integration partners to help with deployments.

Solution Use Cases

AWS can enable customers to cost-effectively operate several disaster recovery strategies. **Exhibit 22** shows a spectrum of scenarios—“backup & restore,” “pilot light,” “warm standby,” and “multi-site”—arranged by how quickly a system can be available to users after a disaster recovery event.

Exhibit 22. Spectrum of Disaster Recovery Options.



Each disaster recovery option is discussed in more detail as follows.

- **Backup and Restore:** In most traditional environments, data is backed up to tape and sent off site regularly. Recovery time will be the longest using this method; lack of automation leads to increased costs. Using Amazon Simple Storage Service (Amazon S3) is ideal for backup data because it is designed to provide 99.99999999 percent durability

Amazon S3 is designed to provide 99.99999999 percent durability of objects over a given year. This durability level corresponds to an average annual expected loss of 0.00000001 percent of objects.

of objects over a given year. Transferring data to and from Amazon S3 is typically done over the network; therefore, it is accessible from anywhere. Additionally, with AWS Storage Gateway, customers can automatically back up on-premises data to Amazon S3.

- **Pilot Light for Simple Recovery into AWS Warm Standby Solution:** The idea of the pilot light is an analogy that comes from the gas heater. In a gas heater, a small idle flame that is always on can quickly ignite the entire furnace to heat a house as needed. This scenario is analogous to a backup-and-restore scenario; however, customers must confirm that they have the most critical core elements of their system already configured and running in AWS (the pilot light). When the time comes for recovery, customers would rapidly provision a full-scale production environment around the critical core.
- **Warm Standby Solution in AWS:** The term “warm standby” is used to describe a disaster recovery scenario in which a scaled-down version of a fully functional environment is always running in the cloud. It further decreases recovery time because, in this scenario, some services are always running. By identifying business-critical systems, customers can fully duplicate these systems on AWS and have them always on.
- **Multi-Site Solution Deployed on AWS and On Site:** A multisite solution runs in AWS and on a customer’s existing on-premises infrastructure in an active/active configuration. During a disaster, an organization can simply send the traffic to AWS servers, which can scale to handle the full production load.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Azure Backup can now back up your on-premises application workloads, including Microsoft SQL Server, Hyper-V virtual machines, Microsoft SharePoint, and Microsoft Exchange. You can back up your applications to a local disk or to Azure, allowing you to eliminate local tape libraries and leverage the unlimited storage capability of Azure. You can also manage all your on-premises backups from a single user interface. Backup continues to support backups of your production IaaS virtual machines in Azure and to help protect your Windows client data and your shared files and folders.

System Center Data Protection Manager is an option for on-premises, Azure, or Cloud Only backup and recovery.

Azure Site Recovery: ASR’s enhanced VMware to Azure scenario is now Generally Available. This GA release, among other enhancements, is designed to help customers benefit from the following key functionality:

Elimination of IaaS-based replication and orchestration components/appliance

MSI-based unified setup of on-premises components, which significantly reduces the time and complexity to onboard to the scenario

Non-disruptive disaster recovery testing with Test Failover

ASR-integrated failback experience without vContinuum, with support for alternate location recovery, and original location recovery

Disk-based replication from source machines, and driver installation without needing a source reboot

Multi-VM Application and Crash-Consistent Replication for Windows and Linux

Migration of protected machines from the in-market – Legacy – VMware to Azure scenario to the Enhanced VMware to Azure scenario

Enterprise-grade enhancements such as support for FQDNs, custom ports, and installation paths

Support for CentOS & RHEL 6.7, vCenter Server 6.0

For Microsoft Government Community Cloud (GCC) Services as defined in Microsoft's service terms and conditions, customer content is stored at rest in the United States. For the GCC versions of Exchange Online, SharePoint Online, Skype for Business, and Dynamics CRM Online, customer content is stored in encrypted format. In the GCC version of Azure Core Services, customers can encrypt nonpublic customer content. For the non-GCC (public) versions of the equivalent services, as well as for Microsoft Intune Online Services, certain types of customer content are stored at rest in the United States, if users in the United States set them up. The terms and conditions governing where customer data is stored can be found in the Microsoft Online Services Terms. For the non-GCC version of Azure Core Services, customers are given the choice of Microsoft worldwide data centers in which they can store and process data. For purposes of GCC Services, customer content means the subset of customer data created by users. For Office 365 Services, customer content at least includes Exchange Online mailbox content (email body, calendar entries, and the content of email attachments), SharePoint Online site content and the files stored at that site, and Skype for Business Online archived conversations. For Microsoft Dynamics CRM Online Services, customer content consists of the entities of customer data managed by Microsoft Dynamics CRM Online Services.

Google Public and Community Cloud for SaaS: Google's data storage methodology ensures that data is replicated across multiple servers and data centers. If a single server begins to exhibit signs of poor performance the customer data will be seamlessly moved to a different health server in a manner that is seamless to the end user. There is no likely instance of an extended downtime.

The likelihood of an unrecoverable loss of data would be more likely due to an end user omission or bad act. Customers have the ability, via Google Apps Vault, to set retention policies on email data and have Data Recovery tools in the Admin Panel that will allow for end user deleted data to be restored up to 25 days after it was deleted.

Systems fail all the time but Google designed its infrastructure to shield customers from that with the methods that will move customer data away from hardware that having issues.

A system outage would not necessarily lead to data loss. Google's infrastructure would be able to serve up the data from any of the redundant locations that was used as part of the real time data replication.

Google has both objectives set to zero based on the inherent design of the infrastructure.

The following methodologies are used for backup and restore services:

Online tape backups

Standard server backup techniques

Encrypted tapes stored on site in locked vaults. Some tapes may be stored with an offsite data storage company.

Google leverages all the data centers all the time to ensure rapid and real time data replication.

8.9 (E) DATA PROTECTION

8.9.1 *Specify standard encryption technologies and options to protect sensitive data, depending on the particular service model that you intend to provide under this Master Agreement, while in transit or at rest.*

Unisys Response:

All of the offerings in this proposal that are provided by our partners use SSL/TLS encryption for browser-based transactions. For non-browser based access, Unisys provides Secure Shell, Secure FTP based access. Additionally, various security APIs with documentation are available for customers to transact with our environments securely.

8.9.2 *Describe whether or not it is willing to sign relevant and applicable Business Associate Agreement or any other agreement that may be necessary to protect data with a Purchasing Entity.*

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS

After an opportunity for review, Unisys is willing to sign a relevant and applicable Business Associate Agreement or any other agreement that may be necessary to protect data with a Purchasing Entity.

8.9.3 *Offeror must describe how it will only use data for purposes defined in the Master Agreement, participating addendum, or related service level agreement. Offeror shall not use the government data or government related data for any other purpose including but not limited to data mining. Offeror or its subcontractors shall not resell nor otherwise redistribute information gained from its access to the data received as a result of this RFP.*

Unisys Response: Unisys is contractually obliged to safeguard a client's data by signing a confidentiality agreement; we cascade this clause to our partners as well. When joining Unisys, associates must sign a Non-Disclosure Agreement. We reiterate the importance of maintaining the confidentiality of client data in annual security awareness training. Unisys understands that breaches of confidential agreements are liable to legal action.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Customers maintain ownership of their customer content and select which AWS services process, store and host their customer content. AWS does not access or use customer content for any purpose other than as legally required and for maintaining the AWS services and providing them to our customers and their end users. We never use customer content or derive information from it for marketing or advertising.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Customer data will be used only to provide a Purchasing Entity with Online Services, including purposes compatible with providing them. Microsoft does not use customer data or derive information from it for advertising or similar commercial purposes. As between the parties, the Purchasing Entity retains the right, title, and interest in customer data. Microsoft acquires no rights in customer data other than the rights that the customer grants to Microsoft to provide the customer with Online Services. This paragraph does not affect Microsoft's rights in software or Microsoft Online Services licenses to the Purchasing Entity. Microsoft uses data mining only to provide the cloud services, subject to the above-mentioned restrictions. Microsoft does not use data mining for unrelated commercial purposes, advertising or advertising-related purposes, or for purposes other than security or service delivery analysis that are not explicitly authorized.

Google Public and Community Cloud for SaaS: Google Apps customers own their data. Google does not sell customer data to third parties. Google offers a detailed Data Processing Amendment that describes our commitment to protecting customer data. Google will not process your data for any purpose other than to fulfill our contractual obligations. Further, Google commits to deleting data from our systems within 180 days of the customer deleting it in our services. Finally, we provide tools to make it easy for the customer to take their data if they choose to stop using our services altogether, without penalty or additional cost imposed by Google.

8.10 (E) SERVICE LEVEL AGREEMENTS

8.10.1 *Offeror must describe whether your sample Service Level Agreement is negotiable. If not describe how it benefits purchasing entity's not to negotiate your Service Level Agreement.*

8.10.2 *Offeror, as part of its proposal, must provide a sample of its Service Level Agreement, which should define the performance and other operating parameters in which the infrastructure must operate to meet IT System and Purchasing Entity's requirements.*

The sample SLA's cover a broad range of services. They are not negotiable. We are constantly evaluating and updating our SLA's to meet the needs of our customers. We are confident that our SLA's meet the overall requirements of the participating entity. The sample SLA's are available in the appendix sections itemized below.

AWS SLA is available in the Appendix 8.

Microsoft SLA is available in the Appendix 9.

Google Apps SLA is available in the Appendix 10.

8.11 (E) DATA DISPOSAL

Specify your data disposal procedures and policies and destruction confirmation process.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys data disposal and retention policies guide the requirement to maintain information for a predetermined time. Additionally, we develop procedures for archiving the information and guidelines for destroying it. We customize storage and disposal according to a client's requirements or Unisys defined policies, with a high priority on the client's requirements. The Unisys standard for media handling lists the types of disposal mechanisms to be followed for various media used for data storage. Disposal mechanisms include simple erasure, secure erasure, physical destruction, and the use of commercial data destruction services.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: It is important for customers to understand the following important basics of data ownership and management in the cloud shared responsibility model:

- Customers continue to own their data.
- Customers choose the geographic locations in which to store their data—it does not move unless the customer decides to move it.
- Customers can download or delete their data whenever they like.
- Customers should consider the sensitivity of their data, and decide if and how to encrypt the data while it is in transit and at rest.

AWS provides customers with the ability to delete their data. However, AWS customers retain control and ownership of their data; customers must manage their data. When a storage device reaches the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from exposure to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M, National Industrial Security Program Operating Manual or NIST 800-88, Guidelines for Media Sanitization to destroy data as part of the decommissioning process. Decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft uses best practice procedures and a wiping solution that is NIST 800-88 compliant. For hard drives that cannot be wiped, Microsoft uses a destruction process that destroys it (shredding) and renders the recovery of information impossible (e.g., disintegrating, shredding, pulverizing, or incinerating). The asset type determines the appropriate means of disposal. Microsoft retains records of the destruction. Microsoft Azure services use approved media storage and disposal management services. Paper documents are destroyed by approved means at the predetermined end-of-life cycle.

Data destruction techniques vary, depending on the type of data object being destroyed, whether it be subscriptions, storage, virtual machines, or databases. In Azure's multitenant environment, Microsoft pays careful attention to prevent one customer's data from leaking into another customer's data. When a customer deletes data, no other customer (usually including the customer that once owned the data) can gain access to that deleted data.

Azure follows NIST 800-88, Guidelines on Media Sanitization, which addresses the principal concern of preventing from unintentional release. These guidelines encompass electronic and physical sanitization.

Microsoft restricts access to administer Microsoft Azure Active Directory portal and services in accordance with assigned privileges and the associated subscription of the customer account. When approved for deletion, Microsoft follows the procedures outlined in DSI-07.1 to remove customer data.

Google Public and Community Cloud for SaaS: When retired from Google's systems, hard disks containing customer information are subjected to a data destruction process before leaving Google's premises. First, disks are logically wiped by authorized individuals using a process approved by the Google Security Team. Then, another authorized individual performs a second inspection to confirm that the disk has been successfully wiped. These erase results are logged by the drive's serial number for

tracking. Finally, the erased drive is released to inventory for reuse and redeployment. If the drive cannot be erased due to hardware failure, it is securely stored until it can be physically destroyed. Each facility is audited on a weekly basis to monitor compliance with the disk erase policy.

8.12 (E) PERFORMANCE MEASURES AND REPORTING

8.12.1 Describe your ability to guarantee reliability and uptime greater than 99.5%. Additional points will be awarded for 99.9% or greater availability.

Unisys Response:

Unisys and our partners in this proposal guarantee availability of 99.5 percent on some service and 99.9 percent availability on other services. The uptime availabilities are contained in the various SLA's provided with this proposal. We are able to provide this availability from our fully redundant, fault-tolerant, geographically dispersed data centers that are connected with fully redundant, fault-tolerant, high-speed connections. Additionally, our talented and skilled 24x365 staff and our best practices and processes can provide this reliability.

8.12.2 Provide your standard uptime service and related Service Level Agreement (SLA) criteria.

Unisys Response:

The SLAs in **Appendix 8, 9, and 10** of this proposal provide the SLA uptimes for our offerings.

8.12.3 Specify and provide the process to be used for the participating entity to call/contact you for support, who will be providing the support, and describe the basis of availability.

Unisys Response:

Unisys enables our clients to directly contact our provider's subsystems. Issues raised directly with the provider will be handled by the SLAs provided by the provider. Unisys also identified a contract manager. Our contract manager will act as the single point of contact (SPOC) for each Participating Entity or specific departments. Each client can directly reach the Unisys SPOC for responses to queries.

8.12.4 Describe the consequences/SLA remedies if the Respondent fails to meet incident response time and incident fix time.

Unisys Response:

Details are provided in the SLAs in **Appendix 8, 9, and 10** of this proposal.

8.12.5 Describe the firm's procedures and schedules for any planned downtime.

Unisys Response:

Unisys and our partners will communicate planned downtime to the user community in advance in email notifications and email reminders.

8.12.6 Describe the consequences/SLA remedies if disaster recovery metrics are not met.

Unisys Response:

Details are covered in the SLAs provided in **Appendix 8, 9 and 10** of this proposal.

8.12.7 Provide a sample of performance reports and specify if they are available over the Web and if they are real-time statistics or batch statistics.

Unisys Response:

For customized cloud solutions surrounding private and hybrid implementations we can customize the frequency and granularity of the performance reports as required. For the public and community cloud deployments, we have provided the sample reports in the exhibits below. Please note the exhibits below are only a sample. Our platform interface is robust, extensible and customizable. Furthermore, the platform offers several documented opportunities to integrate our platform with third party tools for data ingestion, manipulation and decision making.

Exhibit 23 AWS Sample Performance Report




Exhibit 24 Microsoft Sample Performance Report




Exhibit 25 Google Sample Performance Report

orangescape Tasks Reports Admin Masters Upgrade dinesh.varadharajan@orangescape.com									
Reports -> Transport Admin									
Dashboard	Process Step performance summary								Last 12 Months
KPI REPORTS	Process Step	No of Requests	Allocated Time	Avg Time	Max Time	Min Time	SLA Breached	Requests Queried	Requests Queried
Workflow Summary	Manager Approval	64	2 days	1 day 16 hours	26 days 4 hours	44 seconds	23	37	13
Workflow Performance	Finance Approval	30	2 days	1 day 6 hours	23 days 4 hours	8 minutes	2	4	1
Process Step Performance	Head of Finance Approval	74	2 days	1 day 19 hours	42 days 6 hours	45 seconds	14	4	2
Participants Performance	Group Manager Approval	76	2 days	3 days 10 hours	26 days 1 hour	2 minutes	29	5	3
ADHOC REPORTS	Head of Admin	66	2 days	20 hours	14 days 17 hours	2 minutes	11	6	3
You haven't created any adhoc reports. Create your first report by adding a shortcut from within the dashboard.	Payment Processing	75	2 days	14 hours	8 days 21 hours	28 seconds	8	1	1
	Payment Approval	70	2 days	3 days 6 hours	22 days 2 hours	14 minutes	39	1	1
	Payment Completion	19	2 days	7 hours	5 days 9 hours	1 minute	3	1	0
	Payment Acknowledgement	66	2 days	6 hours	4 days	23 seconds	3	0	0
	Process Audit and Comments	52	2 days	2 days	17 days 23 hours	2 minutes	13	19	10

8.12.8 Ability to print historical, statistical, and usage reports locally.**Unisys Response:**

Unisys provides each of the Participating Entities access to a Billing and Consumption portal. This portal will have the ability to provide a tabular view of a client's current as well as past billing and usage details of cloud resources on a real-time basis. It can be accessed by browser and provides clients the ability to print historical, statistical, and usage reports locally onto a client's machine in an excel or pdf format. Clients will have the ability to filter information of billing and usage and obtain granular information at a service, instance or a pre-defined group level.

8.12.9 Offeror must describe whether or not its on-demand deployment is supported 24x365.**Unisys Response:**

The offerings in this proposal provide on-demand deployment 24x365.

8.12.10 Offeror must describe its scale-up and scale-down, and whether it is available 24x365.**Unisys Response:**

The offerings in this proposal provide scaling 24x365.

8.13 (E) CLOUD SECURITY ALLIANCE

Describe your level of disclosure with CSA Star Registry for each Solution offered.

- Completion of a CSA STAR Self-Assessment, as described in Section 5.5.3.*
- Completion of Exhibits 1 and 2 to Attachment B.*
- Completion of a CSA STAR Attestation, Certification, or Assessment.*
- Completion CSA STAR Continuous Monitoring.*

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Appendix 3

See Appendix 3

According to the CSA definitions, AWS aligns with Level 2 in the following determinations in its third-party audits for SOC and ISO:

- Level 2 Attestation is based on SOC 2, which can be requested under the NDA. The SOC 2 report audit attests that a third-party auditor validated AWS to confirm that AWS control objectives are appropriately designed and operating effectively.
- Level 2 Certification is based on ISO 27001:2005. The AWS ISO 27001:2005.

The AWS self-assessed assertions in the CSA STAR Registry Self-Assessment are backed by independent third-party audits across multiple compliance programs. AWS continues to assert that it raises the bar on CSA's attestation and certification program.

Per the CSA website, CSA Level 3 Continuous Monitoring is still under development. AWS has implemented and documented a Continuous Monitoring Plan which defines AWS' approach to conducting continuous monitoring with its authorizing officials within the FedRAMP Security Assessment Framework. It is based on the continuous monitoring process described in NIST SP 800-137, Information Security Continuous Monitoring for Federal Information Systems and Organization, and has been reviewed and validated by a third-party assessor as part of our annual FedRAMP Assessment. It is made available to customers within the AWS FedRAMP Package which can be obtained under NDA through <https://aws.amazon.com/compliance/contact/>

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Microsoft is on the Board of Directors of the CSA. All the questionnaires are available in the appendix section as outlined below.

Appendix 4 – Microsoft O365 CCM

See Appendix 4

Appendix 5 – Microsoft Azure CCM

See Appendix 5

Appendix 6 – Microsoft Dynamics CCM

See Appendix 6

Google Public and Community Cloud for SaaS:

Appendix 7 – Google CAIQ

Google has completed the CAIQ and it is available in the Appendix

See Appendix 7

8.14 (E) SERVICE PROVISIONING

8.14.1 Describe in detail how your firm processes emergency or rush services implementation requests by a Purchasing Entity.

Unisys Response: The standard process and the lead times for onboarding a Participating Entity are as follows:

- Assuming the Participating Entity has a Participating Addendum (PA) in place - Upon the qualification and the understanding of the requirements from the Participating Entity Unisys will deliver a statement of work (SOW) within 5 business days.
- Once the Participating Entity sign's and returns the SOW, Unisys will provision the necessary account(s) in 10 business days and provide the required credentials to the Participating Entity so that they can start using the service.

As an incumbent, Unisys has the proven past performance in expediting "rush" service requests. We have delivered the above set of onboarding in less than half the time and we can do even better. It depends on the nature of the request and cooperation and participation from the participating entity.

Service provisioning once the customer has been on-boarded, is delivered on a self-service basis. Such provisioning services are available 24X365. Participating Entities can provision all services available via the Unisys portal or the AWS EC2 portal or the Microsoft Azure platform or the Google Cloud Platform within 5-10 minutes or less.

8.14.2 Describe in detail the standard lead-time for provisioning your Solutions.

Unisys Response: The standard process and the lead times for onboarding a Participating Entity are as follows:

- Assuming the Participating Entity has a Participating Addendum (PA) in place - Upon the qualification and the understanding of the requirements from the Participating Entity Unisys will deliver a statement of work (SOW) within 5 business days.
- Once the Participating Entity sign's and returns the SOW, Unisys will provision the necessary account(s) in 10 business days and provide the required credentials to the Participating Entity so that they can start using the service.

Once a participating entity has been on boarded, provisioning of services requires no lead time. It is based on a self-provisioning system. As an example all services available via the Unisys portal or the AWS EC2 portal or the Microsoft Azure platform or the Google platform can be self-provisioned within 5-10 minutes or less.

For Private and Hybrid cloud solutions that are customized Unisys will provide a detailed project plan based on the scope of services requested.

8.15 (E) BACK UP AND DISASTER PLAN

8.15.1 Ability to apply legal retention periods and disposition by agency per purchasing entity policy and/or legal requirements.

Unisys Response:

Customers can customize retention periods in accordance with purchasing entity policy, legal requirements, or both across all providers in this RFP.

8.15.2 Describe any known inherent disaster recovery risks and provide potential mitigation strategies.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: We feel very confident about our physically and logically redundant, geographically distributed data center architecture that is constantly replicated. Our provided SLA's speak to our performance and capability in providing the needed protection against such risks.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Amazon's Disaster Recovery Approach is explained in detail in section 8.8 (Service or Data Recovery)

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft's Disaster Recovery Approach is explained in detail in section 8.8 (Service or Data Recovery)

Google Public and Community Cloud for SaaS:

Google's Disaster Recovery Approach is explained in detail in section 8.8 (Service or Data Recovery)

8.15.3 Describe the infrastructure that supports multiple data centers within the United States, each of which supports redundancy, failover capability, and the ability to run large scale applications independently in case one data center is lost.

Unisys Private and Hybrid Cloud for IaaS, PaaS, and SaaS: Unisys has set up primary and secondary data centers. If required by a contract, applications and devices are configured in high availability mode, monitored, and tested regularly. Unisys establishes data mirroring controls and performs backups

regularly. We define recovery parameters such as RTOs and RPOs. We perform risk assessments and business impact analyses regularly on our data centers. In case of a data center outage, the disaster recovery trigger would automatically kick in and will enable failover to other data centers.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Amazon's infrastructure has a high level of availability and provides customers with the features to deploy a resilient IT architecture. AWS designed its systems to tolerate system or hardware failures with minimal impact on customers. The Amazon Infrastructure Group directs Business Continuity Management for data centers at AWS.

Availability

Data centers are built in clusters in various global regions. Data centers are on line and serving customers; no data center is "cold." During a failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that during a data center failure, there is sufficient capacity to enable traffic to be load balanced to the remaining sites.

AWS provides customers with the flexibility to move instances and store data in multiple geographic regions as well as across multiple availability zones in each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated in a typical metropolitan region and located on lower risk flood plains (specific flood zone categorization varies by region). In addition to discreet uninterruptible power supply (UPS) and onsite backup generation facilities, they are each fed by different grids from independent utilities to further reduce single points of failure. Availability zones are redundantly connected to multiple Tier 1 transit providers.

Customers should architect their AWS use to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in most failure modes, including natural disasters or system failures.

Fault-tolerant Design

Amazon's infrastructure has a high level of availability and provides customers with the capability to deploy a resilient IT architecture. AWS designed its systems to tolerate system or hardware failures with minimal impact on customers.

Clients should architect their AWS use to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures. However, clients should be aware of location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between regions unless proactively done by the customer, thereby enabling customers with these types of data placement and privacy requirements to establish compliant environments. Note that communication between regions is across public Internet infrastructure; therefore, appropriate encryption methods should be used to protect sensitive data.

As of this proposal, there are 12 regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), EU (Ireland), EU (Frankfurt), Asia-Pacific (Singapore), Asia-Pacific (Tokyo), Asia-Pacific (Sydney), Asia-Pacific (Seoul), South America (São Paulo), and China (Beijing).

AWS GovCloud (US) is an isolated AWS region designed to allow Federal agencies and customers to move workloads to the cloud by helping them meet certain regulatory and compliance requirements. The AWS GovCloud (US) framework allows Federal agencies and their contractors to comply with U.S. International Traffic in Arms Regulations (ITAR) regulations and Federal Risk and Authorization Management Program (FedRAMP) requirements. AWS GovCloud (US) received an Agency Authorization to Operate (ATO) from the U.S. Department of Health and Human Services (HHS) to use a FedRAMP-accredited Third Party Assessment Organization (3PAO) for several AWS services.

The AWS GovCloud (US) Region provides the same fault-tolerant design as other regions, with two Availability Zones. Additionally, the AWS GovCloud (US) region is a mandatory AWS Virtual Private Cloud (VPC) service by default to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft built a fully physically and logically redundant computing platform. Microsoft also provides the flexibility to easily back up critical applications to georedundant cloud data centers and client on-premises or partner data centers. This adds yet another layer of disaster recovery capability. Microsoft delivers more than 200 cloud services that are hosted in more than 100 globally distributed data centers, edge computing nodes, and service operations centers. This infrastructure is supported by one of the world's largest built terabyte global networks, with an extensive dark fiber footprint, that connects them all.


Google Public and Community Cloud for SaaS: The table below lists the Google datacenters the in US.

Google US Data Centers

Product	North America
Compute Engine	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Preemptible VMs	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Container Engine	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Container Registry	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Cloud Dataflow	• Council Bluffs, IA - us-central1 • Berkeley County, SC - us-east1
Cloud Bigtable	Council Bluffs, IA - us-central1
Cloud Dataproc	Council Bluffs, IA - us-central1

The following diagram shows how these Google Cloud Platform components work together to build a scalable, resilient web application.

Exhibit 26 Google Cloud Platform



Each component in the example application architecture plays a role in ensuring the application is both scalable and resilient. This section briefly describes each of these services.

The HTTP load balancer exposes a single public IP address that customers use to access the application. This IP address can be associated with a DNS A record (e.g., example.com) or CNAME (e.g., www.example.com). Incoming requests are distributed across the instance groups in each zone according to each group's capacity. Within the zone, requests are spread evenly over the instances within the group. Although the HTTP load balancer can balance traffic across multiple regions, we're using it in a single region with multiple zones, as described in the next section.

A zone is an isolated location within a region. Zones have high-bandwidth, low-latency network connections to other zones in the same region. Google recommends deploying applications across multiple zones in a region.

An instance is a virtual machine hosted on Google's infrastructure. You can install and configure these instances just like physical servers. In this document, you can use startup scripts and Chef to configure instances with the application server and code for the web application.

An instance group is a collection of homogeneous instances that can be targeted by an HTTP load balancer. Instances are added to and removed from a group by an instance group manager. An instance group and corresponding manager are required for each zone you want to run in.

The Compute Engine Autoscaler adds or removes Google Compute Engine instances to an instance group by interfacing with the group's manager in response to traffic, CPU utilization, or other signals. In the example solution, the Autoscaler responds to the Request Per Second (RPS) metric of the HTTP load balancer. An Autoscaler is required for each instance group that you want scaled automatically.

Google Cloud SQL is a fully managed MySQL database. Replication, encryption, patches, and backups are managed by Google. A Cloud SQL instance is deployed to a single zone, and data is replicated to other zones automatically. The Redmine application used in this example is compatible with MySQL and works seamlessly with Cloud SQL.

Cloud Storage allows objects (usually files) to be stored and retrieved with a simple and scalable interface. In this solution, a Cloud Storage Bucket is used to distribute private SSL keys to the scalable Google Compute Engine instances, and is also used to store all files uploaded to the Redmine application, meaning no stateful information is stored on any instance's disks.

A dynamic decision is made throughout a user's log in session on which servers and which data centers will be used to store primary and backup copies of all data. If one server begins to fail, data is moved to another healthy server. If one data center is having broad localized issues, all data will be served up and stored at the other healthy data centers. No action is required by the customer to invoke this service and it is transparent to your end users.

8.16 (E) SOLUTION ADMINISTRATION

8.16.1 Ability of the Purchasing Entity to fully manage identity and user accounts.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System enables several authentication methods, including username/password, username/password + 2nd factor, identity federation (SAML 2.0), OpenID, and a proprietary Single Sign-On (SSO) API. User authentication credentials can be stored locally in the Unisys Cloud Provisioning System and authenticated there, or they can be leveraged from a customer-hosted LDAP directory (including AD). Note that the LDAP/AD model works only for the on-premises deployment of the Unisys Cloud Provisioning System and is not supported through our SaaS offering.

When username/password is the option for the local Unisys Cloud Provisioning System user store, passwords are stored securely using a versioned one-way hashing algorithm. This algorithm can be replaced if necessary.

The Unisys Cloud Provisioning System also enables users to have one or more API keys for talking to the Unisys Cloud Provisioning System API. These API keys enforce a user's access rights when the user is interacting with the Unisys Cloud Provisioning System API. A user can revoke an API key without affecting the others. These API keys are stored securely using AES-256 encryption based on an administrator-managed encryption key.

Each user can belong to zero or more groups. Users and groups can be maintained in the Unisys Cloud Provisioning System, or synchronized with an external LDAP or AD directory (regardless of whether authentication goes through LDAP or AD). In an account (such as an R&D Amazon Web Services account instead of a production account), each group is assigned a role. A group can therefore have a different role for the R&D account as opposed to the production account. In fact, no role can be assigned that results in no access rights for that role. Each user set up in the Unisys Cloud Provisioning System can share a common access point with Amazon Web Services—freeing system administrators from the requirement to create individual Amazon Web Services users.

Access rights are determined by the role. System administrators can create their own roles and map them to groups to suit business needs. System administrators can also define permissions for each role at a very granular level. An access right is defined in a RESOURCE, ACTION, QUALIFIER triad. The default for a new role is no access rights at all. System administrators can create an administrative role by setting the triad to ANY/ANY/ANY, or they can limit the ability to terminate virtual machines for a specific role by setting the right VM/Terminate/MINE. System administrators can define of permissions for roles.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Each participating agency can have its own AWS account. AWS provides several ways for customers to identify themselves and securely access their AWS account.

AWS Identity and Access Management (IAM) enables system administrators to securely control access to AWS services and resources for users. IAM enables system administrators to create and manage users in

AWS, and it enables them to grant access to AWS resources for users managed outside AWS in a corporate directory.

IAM also enables identity federation between corporate directory and AWS services. This enables system administrators to use existing corporate identities to grant secure and direct access to AWS resources, such as S3 buckets, without creating a new AWS identity for those users.

Exhibit 15 shows the AWS IAM tab on the AWS Management Console.




Exhibit 15. AWS IAM Tab.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Managing identity is just as important in the public cloud as it is on premises. To help with managing identity, Azure supports several different cloud identity technologies, which include the following.

Administrators can run Windows Server Active Directory (commonly called just AD) in the cloud using virtual machines created with Azure Virtual machines. This approach makes sense when administrators using Azure to extend an on-premises data center to the cloud.

Administrators can use Azure Active Directory to give users a single sign-on to Software as a Service (SaaS) applications. Microsoft Office 365 uses this technology, for example, and applications running on Azure or other cloud platforms can also use it.

Applications running in the cloud or on premises can use Azure Active Directory Access Control to let users log in using identities from Facebook, Google, Microsoft, and other identity providers.

There are several options for connecting the domain controllers in the cloud with those running on premises:

- Make them part of a single Active Directory domain.
- Create separate AD domains on-premises and in the cloud that are part of the same forest.
- Create separate AD forests in the cloud and on-premises, then connect the forests using cross-forest trusts or Windows Server Active Directory Federation Services (AD FS), which can also run in virtual machines on Azure.

As SaaS applications become more common, they raise an obvious question: What kind of directory service should these cloud-based applications use? Microsoft's answer to that question is Azure Active Directory.

There are two main options for using this directory service in the cloud:


- Individuals and organizations that use only SaaS applications can rely on Azure Active Directory as their sole directory service.
- Organizations that run Windows Server Active Directory can connect their on-premises directory to Azure Active Directory, then use it to give their users single sign-on (SSO) to SaaS applications.

Administrators can federate on-premises AD or other directory stores with Azure AD. Microsoft recommends that organizations synchronize their on-premises directory information with SSO to enable users removed from on-premises directories to be also removed from the Azure AD to maintain sufficient access controls. Deployment of a highly available Security Token Service (STS) (such as Active Directory Federation Services 2.0 or AD FS) on premises is required because this STS will become part of the authentication flow for every user login. Once federation is configured, Azure AD users whose identities are based on the federated domain can use their existing corporate login to authenticate to Azure AD services. Federation enables secure, token-based authentication and SSO across Azure applications.

Microsoft provides Multi-Factor Authentication for Azure administrators with a phone as the second factor; it also supports integration with third-party authentication solutions via on-premises STS integration.

Google Public and Community Cloud for SaaS: Each participating agency can have its own Google account. Google provides several ways for customers to identify themselves and securely access their AWS account. Google Identity and Access Management (IAM) enables system administrators to securely control access to Google services and resources for users. IAM enables system administrators to create and manage users in the Google Cloud. IAM also enables them to grant access to Google resources.

The Graphic below Google IAM Overview illustrates Google IAM model.



8.16.2 Ability to provide anti-virus protection, for data stores.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys Data centers and assets are protected by Symantec End Point Protection. Unisys infrastructure has policies defined that allows real-time background scanning in which a memory resident scanner runs continuously in background mode and automatically scans files and media upon access. All Unisys assets including their Data centers undergoes periodic full system scans that would detect for malicious and

unwanted programs, files and media. It is also the responsibility of the client to make sure to follow best practices while using Unisys infrastructure to refrain from accessing potential malicious sites and programs.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Moving IT infrastructure to AWS creates a shared responsibility model between the customer and AWS. This shared model can relieve customer operational burden because AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. Customers assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, and the configuration of the AWS-provided security group firewall. Customers should carefully consider the services they choose because their responsibilities vary, depending on the services used, the integration of those services into their IT environment, and applicable laws and regulations. To enhance security or meet more stringent compliance requirements (or do both), customers can use technology such as host-based firewalls, host-based intrusion detection and prevention, antivirus software, encryption, and key management. The nature of this shared responsibility also provides the flexibility and customer control that allows deployment of solutions that meet industry-specific certification requirements.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Microsoft Antimalware offers customers the ability to install an antimalware agent for PaaS roles and virtual machines. Based on System Center Endpoint Protection, this new feature brings proven on-premises security technology to the cloud.

We offer deep integration for Trend's Deep Security and SecureCloud products on the Azure platform. DeepSecurity is an antivirus solution; SecureCloud is an encryption solution. DeepSecurity will be deployed inside virtual machines using an extension model that Azure already announced. Using the portal UI and PowerShell, customers can choose to use DeepSecurity inside new virtual machines that are being spun up, or existing virtual machines that are already deployed.

Symantec End Point Protection (SEP) is also supported on Azure. Through portal integration, customers can specify that they intend to use SEP in a virtual machine. SEP can be installed on a brand new virtual machine at the Azure Portal or on an existing virtual machine using PowerShell.

Google Public and Community Cloud for SaaS:

Google's malware strategy begins with infection prevention by using manual and automated scanners to scour Google's search index for websites that may be vehicles for malware or phishing. Approximately one billion people use Google's Safe Browsing on a regular basis. Google's Safe Browsing technology examines billions of URLs per day looking for unsafe websites. Every day, we discover thousands of new unsafe sites, many of which are legitimate websites that have been compromised. When we detect unsafe sites, we show warnings on Google Search and in web browsers. In addition to our Safe Browsing solution, Google operates VirusTotal, a free online service that analyzes files and URLs enabling the identification of viruses, worms, trojans and other kinds of malicious content detected by antivirus engines and website scanners. VirusTotal's mission is to help in improving the antivirus and security industry and make the Internet a safer place through the development of free tools and services.

8.16.3 Ability to migrate all Purchasing Entity data, metadata, and usage data to a successor Cloud Hosting solution provider.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

Unisys provides client the ability to download or migrate their data and workloads as and when required by the client. Unisys provides Cloud Application Migration Services that starts with a complete and analytical view of your application portfolio and cloud/data center capabilities. Using a set of criteria refined during customer engagements, Unisys can help you narrow the application pool of various cloud

portable candidates and then provide the detailed assessment of costs and other factors that go into a cloud migration decision. Once the decision has been made, Unisys can help you with the planning and implementation of the application migration process including on-boarding of applications to AWS and other clouds.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Customers are free to export data. For large amounts of data, they can use the AWS Import/Export service to copy data to customer-supplied storage and deliver it to a customer-specified location. AWS Import/Export Disk accelerates moving large amounts of data into and out of the AWS cloud by using portable storage devices for transport. Disk transfers your data directly onto and off of your storage devices using the Amazon high-speed internal network and bypassing the Internet. For gigabyte-scale data sets that would transfer at less than 10 Mbps, Disk is often faster than Internet transfer and more cost-effective than upgrading your connectivity.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

A customer's use of an Import/Export Service is conditioned on its compliance with the instructions provided by Microsoft for the preparation, treatment, and shipment of physical media containing its data (storage media). The customer is solely responsible for verifying the storage media and data are provided in compliance with laws and regulations. Microsoft has no duty toward the storage media and no liability for lost, damaged, or destroyed storage media. Storage media shipped to Microsoft must be shipped DAP Microsoft DCS Data Center (INCOTERMS 2010). Storage media shipped to the customer is shipped DAP Customer Dock (INCOTERMS 2010).

Google Public and Community Cloud for SaaS: Customers are free to export data. For large amounts of data, they can use our documented API's to Import/Export /copy data to customer-supplied storage.

8.16.4 Ability to administer the solution in a distributed manner to different participating entities

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS: The Unisys Cloud Provisioning System can also be managed in a distributed way. Access controls enable managers to have a specific level of control in the system. For example, a person in the QA department in one city can manage the QA systems and budget controls, and a person in the Development department in another city can manage a different set of infrastructure and the corresponding budget.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS: Each Participating Entity can be provisioned with one or more separate accounts. The AWS environment is a virtualized, multitenant environment. AWS implemented security management processes, PCI controls, and other security controls designed to isolate each customer account from other customer accounts. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture was validated by an independent PCI Qualified Security Assessor (QSA) and found to comply with the requirements of PCI DSS version 2.0 published in October 2010.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS: Azure Role-Based Access Control (RBAC) enables fine-grained access management for Azure. Using RBAC, administrators can segregate duties on a DevOps team and grant users only the amount of access that they need to perform their jobs.

Each Azure subscription is associated with one Azure Active Directory. Only users, groups, and applications from that directory can be granted access to manage resources in the Azure subscription, using the Azure portal, Azure Command-Line tools and Azure Management APIs.

To grant access, administrators assign the appropriate RBAC role to users, groups, and applications at the right scope. To grant access to the entire subscription, administrators assign a role at the subscription scope. To grant access to a specific resource group in a subscription, administrators assign a role at the

resource group scope. They also can assign roles for specific resources, such as websites, virtual machines and subnets, to grant access only to a resource.

Google Public and Community Cloud for SaaS: Each Participating Entity can be provisioned with one or more separate accounts. The Google environment is a virtualized, multitenant environment. Google provides RBAC to the underlying infrastructure. Google systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software.

8.16.5 Ability to apply a participating entity's defined administration policies in managing a solution.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System supports major cloud providers with a unified management solution that includes access controls, financial tracking, and automation. The Unisys Cloud Provisioning System also has a REST-based API to allow for integration with internal and third-party systems. Each Participating Entity can set up its own policies for who has access, how budgets are handled, and more.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

Each Participating Entity can be provisioned with one or more separate accounts. The AWS environment is a virtualized, multitenant environment. AWS implemented security management processes, PCI controls, and other security controls designed to isolate each customer account from other customer accounts. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture was validated by an independent PCI Qualified Security Assessor (QSA) and found to comply with the requirements of PCI DSS version 2.0 published in October 2010.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

Each subscription in Azure belongs to only one directory, each resource group belongs to only one subscription, and each resource belongs to only one resource group. Access that administrators grant at parent scopes is inherited at child scopes. If administrators grant a reader role to an Azure AD group at the subscription scope, that group's members will be able to view every resource group and every resource in the subscription. If administrators grant the contributor role to an application at the resource group scope, the application will be able to manage resources in that resource group, but not other resource groups in the subscription.

The finer grained authorization model (Azure RBAC) is supported only for management operations of the Azure resources at the Azure portal and in Azure Resource Manager APIs. Not all data-level operations for Azure resources can be authorized by Azure RBAC. For example, the creation, reading, updating, and deletion of storage accounts can be controlled by RBAC, but the creation, reading, updating, and deletion of blobs or tables in the storage account cannot yet be controlled by RBAC. Similarly, the creation, reading, updating, and deletion of a SQL database can be controlled by RBAC, but the creation, reading, updating, and deletion of SQL tables in the database cannot yet be controlled by RBAC.

In Azure API Management, policies are a powerful capability of the system that allow a publisher to change the API's behavior through configuration. Policies are a collection of statements that are executed sequentially on the request or response of an API. Popular statements include format conversion from XML to JSON and call-rate limiting to restrict the amount of incoming calls from a developer. Many more policies are available out of the box.

Google Public and Community Cloud for SaaS:

Each Participating Entity can be provisioned with one or more separate accounts. The Google environment is a virtualized, multitenant environment. Google provides RBAC to the underlying infrastructure. Google systems are designed to prevent customers from accessing physical hosts or

instances not assigned to them by filtering through the virtualization software.

8.17 (E) HOSTING AND PROVISIONING

8.17.1 Documented cloud hosting provisioning processes, and your defined/standard cloud provisioning stack.

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System includes not only basic provisioning (e.g., starting a server, creating a machine image, creating a new volume, attaching the volume to a server, creating a snapshot of a volume, creating a volume from a snapshot, reserving an IP address, creating a standard network or port, or creating or modifying a firewall), but also robust automation features such as creating and deploying multitier applications and keeping the application 100 percent available through autoscaling, autorecovery, and automatic backups. This enhancement is more significant than that in the base Amazon provisioning capability. Additionally, using the Unisys Cloud Provisioning System enables this capability across our supported cloud hosts and private cloud solutions.

AWS Public and Private Cloud for IaaS, PaaS, and SaaS:

There are several methods by which customers can provision a server or application services on AWS. Usually, customers use the web-based EC2 portal, which guides them through the provisioning process. The EC2 portal offers a wide array of preconfigured, fully automated deployable workloads to customize every step of a workload instantiation. Beyond EC2, customers can use external tools to deploy servers or application services to the EC2 portal. These external tools can be vendor products, Open Source solutions, and homegrown scripts and technologies. Exhibit # illustrates the various methods and tools of provisioning available from AWS.




Exhibit 26. AWS Provisioning Process, Methods, and Tools.**Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:**


Microsoft Azure Provisioning is an automated process that allows interaction from the Azure Portal or completely remotely using Azure APIs. For example, a virtual machine can be allowed to interact from the Azure portal in just a few minutes. The normal steps taken are planning for the deployment, selecting the virtual machine image, creating the virtual machine, and then managing the virtual machine. The virtual machine image can come from the images available in the gallery, including Windows and Linux. The processes support customized images from a customer's on-premises standard image template library. These processes can be repeated one at a time or automated with PowerShell to create hundreds of virtual machines at once.

This example is for automating a virtual machine, but the properties in Microsoft Azure can be automated using the same methods, at the portal or in PowerShell. Other examples for automating include cloud storage, SQL Servers, and virtual networks.

Microsoft offers the Azure Portal as the access gateway to create, modify, and decommission workloads on the Microsoft cloud platform. The Azure portal can be accessed over the Internet. **Appendix 19 – Getting Started with VMs on Windows Azure** outlines the provisioning process and the provisioning stack.


Google Public and Community Cloud for SaaS: The Cloud Platform provides options for computing and hosting. Customers can choose to work with a managed application platform, leverage container technologies to gain lots of flexibility, or build their own cloud-based infrastructure to have the most control and flexibility.

The exhibit below Google Hosting and Provisioning Services illustrates the spectrum.



The exhibit below illustrates how the Google Compute Engine works.

Google Compute Engine in pictures



8.17.2 Provide tool sets at minimum for:

1. **Deploying new servers (determining configuration for both stand alone or part of an existing server farm, etc.)**
2. **Creating and storing server images for future multiple deployments**
3. **Securing additional storage space**
4. **Monitoring tools for use by each jurisdiction's authorized personnel – and this should ideally cover components of a public (respondent hosted) or hybrid cloud (including Participating entity resources)**

Unisys Public and Community Cloud for IaaS, PaaS, and SaaS:

The Unisys Cloud Provisioning System enables the deployment of cloud servers in new and existing cloud server environments in the context of advanced governance and security to protect users from overspending, which violates ownership and governance rules.

The Unisys Cloud Provisioning System gives system administrators complete control over the creation, configuration, and deployment of not only server images, but also complete multitier application topologies that can be delivered across AWS regions or across cloud providers complete with firewall rules, load balancers, databases, and more.

The Unisys Cloud Provisioning System supports the functionality provided by other vendors such as Amazon and Microsoft. It also supports storage provisioning across many public and private clouds.

The Unisys Cloud Provisioning System supports the direct tracking of CPU usage, memory, network statistics, and disk I/O from the guest operating through our agent. The Unisys Cloud Provisioning System also tracks metrics provided by the cloud provider (e.g., AWS CloudWatch). These can be used for Unisys Cloud Provisioning System autoscaling. The Unisys Cloud Provisioning System graphs some of these metrics, but the metrics can be accessed at the API. The scripts for tracking loads in the guest using the Unisys agent can be extended to use a customer's own scripts or third-party monitoring tools.

Virtual machine load for autoscaling is based on these metrics or on custom metrics. System administrators can write their own scripts that execute on the guest and "vote" on autoscaling. The Unisys Cloud Provisioning System regularly tabulates votes and determines a course of action for scaling.

Amazon Public and Community Cloud for IaaS, PaaS, and SaaS:

The AWS Management Console is a single destination for managing AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. System administrators use the AWS Management Console to perform tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console also enables customers to manage the aspects of their AWS account, including accessing monthly spending by service, managing security credentials, or setting up new AWS Identity and Access Management (AWS IAM) users. The AWS Management Console supports AWS regions and allows customers to provision resources across multiple regions.

Command Line Interface

The AWS Command Line Interface (CLI) is a unified tool used to manage AWS cloud services. With just one tool to download and configure, customers can control multiple AWS resources from the command line and automate them through scripts. The AWS CLI introduces a new set of simple file commands for efficient transfer of files to and from Amazon Simple Storage Service (Amazon S3).

Use of Existing Management Tools

Many of the tools that organizations use to manage on-premises environments also can be integrated with AWS. Integrating an AWS environment can provide a simpler and quicker path for cloud adoption because a customer's operations team does not have to learn new tools or develop completely new processes. For example, AWS provides the following tools.

The AWS Management Portal for vCenter enables customers to manage their AWS resources using VMware vCenter. The portal installs as a vCenter plugin in the existing vCenter environment. Once installed, it enables customers to migrate VMware virtual machines to Amazon EC2 and manage AWS resources from within vCenter. The AWS resources that customers create at the portal can be located in their AWS account, even though those resources were created using vCenter. For experienced VMware administrators, AWS Management Portal for vCenter provides a familiar look and feel that can make it easy to start using AWS. AWS Management Portal for vCenter is available at no additional charge.

The Amazon EC2 VM Import Connector extends the capabilities of VMware vCenter to provide a familiar graphical user interface that customers can use to import their preexisting virtual machines to Amazon EC2. Using the connector, importing a virtual machine is as simple as selecting a virtual machine from the vSphere infrastructure and specifying the AWS region, Availability Zone, operating system, instance size, security group, and Amazon Virtual Private Cloud (Amazon VPC) details (if desired) to which the virtual machine should be imported. Once the virtual machine is imported, system administrators can launch it as an instance from the AWS Management Console and immediately take advantage of the features of Amazon EC2.

AWS Management Pack for Microsoft System Center enables system administrators to view and monitor their AWS resources directly in the Operations Manager console. This way, they can use a single, familiar console to monitor their resources, whether they are on the premises or in the AWS cloud. System administrators get a consolidated view of AWS resources across regions and Availability Zones. It also has built-in integration with Amazon CloudWatch so that the metrics and alarms defined in Amazon CloudWatch surface as performance counters and alerts at the Operations Manager console.

Microsoft Public and Community Cloud for IaaS, PaaS, and SaaS:

The Azure Portal provides the flexibility to create one or thousands of virtual machines (VMs). A single VM can be created, and two or more VMs can be created and moved to an "availability set" so that the uptime SLA is applied.

The Azure Portal and the underlying Azure APIs are used to extend or scale systems easily. Sites (PaaS) options in Azure are scaled by simply configuring the Azure Portal to extend the web farm from 1 to 10 or more in accordance with customized needs. For example, the number of client connections, CPU

utilization, and other metrics could be configured once and the scaled-out systems would deploy automatically, without interaction from administrators.

The Azure Portal supports VMs to be deployed to a previously created VM farm. The VMs are created using automated processes and then configured either through the VM operating system processes or automated by using scripting such as PowerShell.

Additionally, custom virtual machines can be created using the same on-premises processes. They are then uploaded to the Azure Portal to include in the customer's VM library. Once the image is added to the library, the automation processes for single VM deployment or joining to a VM farm are exactly the same.

With a cloud service previously created, the addition of cloud storage can be automated using the Azure Portal or PowerShell. Azure Blob storage is a service that stores file data in the cloud. Blob storage can store text or binary data, such as a document, a media file, or an application installer. Blob storage is sometimes referred to as object storage.

Azure Blob storage can be completely automated using .NET, Node.js, Java, C++, PHP, Ruby, Python, IOS, and Xamarin.

Blob storage includes the following common uses:

- Serving images or documents directly to a browser
- Storing files for distributed access
- Streaming video and audio
- Performing secure backup and disaster recovery
- Storing data for analysis by an on-premises or Azure-hosted service.

Monitoring cloud services are included at the Azure portal and available from APIs at the Azure portal and on the individual VM available in the operating system. Administrators can monitor key performance metrics for their cloud services at the Azure classic portal. They can set the level of monitoring to minimal and verbose for each service role and can customize the monitoring displays. Verbose monitoring data is stored in a storage account, which administrators can access outside the portal.

Monitoring displays at the Azure classic portal are highly configurable. Administrators can choose the metrics they want to monitor on the metrics list on the Monitor page. They can choose which metrics to plot on metrics charts on the Monitor page and the dashboard.

By default, Azure provides minimal for a new cloud service using performance counters gathered from the host operating system for the roles' instances (virtual machines). The minimal metrics are limited to CPU Percentage, Data In, Data Out, Disk Read Throughput, and Disk Write Throughput. By configuring verbose monitoring, administrators can receive additional metrics based on performance data in the virtual machines (role instances). The verbose metrics enable closer analysis of issues that occur during application operations.

By default, performance counter data from role instances is sampled and transferred from the role instance at 3-minute intervals. When administrators enable verbose monitoring, the raw performance counter data is aggregated for each role instance and across role instances for each role at intervals of 5 minutes, 1 hour, and 12 hours. The aggregated data is purged after 10 days.

Google Public and Community Cloud for SaaS: The Google Cloud Platform consists of a set of physical assets, such as computers and hard disk drives, and virtual resources, such as virtual machines (VMs), that are contained in Google's data centers around the globe. Each data center location is in a global region. Regions include Central US, Western Europe, and East Asia. Each region is a collection of zones, which are isolated from each other within the region. Each zone is identified by a name that combines a letter identifier with the name of the region.

The Google Cloud Platform Console provides a web-based, graphical user interface that customers can use to manage their Cloud Platform projects and resources. Using the Cloud Platform Console, customers can create a new project, or choose an existing project, and use the resources that were create in the context of that project.

Google Cloud SDK provides the gcloud command-line tool, which gives customers another method of access. The gcloud tool can be used to manage both development workflow and the Cloud Platform resources.

Cloud Platform provides tools for logging and monitoring so you can keep track keep track of the performance and availability of your resources and applications.

Cloud Logging collects and stores logs from applications and services running on Cloud Platform. Customers can use Cloud Logging with App Engine or Compute Engine. The Logs Viewer in the Cloud Platform Console lets one see their logs. Logs can be exported to Cloud Storage, BigQuery, and Cloud Pub/Sub so cunstomers can process them more easily. The Cloud Logging Agent enables you to integrate third-party logs.

Cloud Monitoring provides dashboards and alerts for applications that run on Cloud Platform. Customers configure Cloud Monitoring using the Cloud Monitoring Console. Review performance metrics for cloud services, Compute Engine instances, and common open source servers such as MongoDB, Apache, Nginx, and Elasticsearch. Customers can use the Cloud Monitoring API to retrieve monitoring data and create custom metrics.

8.18 (E) TRIAL AND TESTING PERIODS (PRE- AND POST- PURCHASE)

8.18.1 Describe your testing and training periods that your offer for your service offerings.

Unisys Response:

We offer generous free trials in some cases for up to a full year. The product set offered during this free period is identical to the paid version of the product. Potential clients, existing clients, or both can use this benefit to take a test drive to learn the environment

8.18.2 Describe how you intend to provide a test and/or proof of concept environment for evaluation that verifies your ability to meet mandatory requirements.

Unisys Response:

In the public and the community spaces for IaaS, PaaS, and SaaS, our partners offer free trials that are generous and come with identical features and functionality of the paid version of the product. Potential clients, existing clients, or both can use this benefit to take a test drive to learn the environment. We also provide a significant amount of audiovisual material and other knowledge transfer media such as white papers and how-to documents. Potential clients, existing clients, or both can use this material for training. For the private and the hybrid spaces for IaaS, PaaS, and SaaS, Unisys customizes most of the solution offerings to the client's specifications. We provide access to user interfaces, onsite proofs of concept, and site visits to our facility, where clients can interact with our staff and test drive environments for feasibility and learning.

8.18.3 Offeror must describe what training and support it provides at no additional cost.

Unisys Response:

Unisys will provide Participating Entities with the following value-add across our cloud models and service models at no additional cost:

- **Provisioning:** Unisys will provision the access login ID for the Participating Entity.
- **Billing Portal:** Unisys will provide a customized, online web-based portal for the Participating Entities to monitor and report on their consumption, utilization, current balances, etc.

- **Billing:** Unisys will bill Participating Entities each month.
- **Consultation:** Unisys will guide Participating Entities on how to publicly available learning tools and content from our partners.

Additionally, the public-facing websites of Unisys and of our partners in this proposal are rich in training content, support material, and tools that our prospects, clients, or both can access at no additional cost.

8.19 (E) INTEGRATION AND CUSTOMIZATION

8.19.1 Describe how the Solutions you provide can be integrated to other complementary applications, and if you offer standard-based interface to enable additional integrations.

8.19.2 Describe the ways to customize and personalize the Solutions you provide to meet the needs of specific Purchasing Entities.

Unisys and its partners platforms on this RFP provide a completely integrated and built on the concept of pluggable, reusable components. The interfaces conform to a standard framework, protocols, and design tools. We provide formal framework documentation; how-to documents; toolsets; and prebuilt, reusable, and customizable components that clients can leverage for additional integration or customization. A significant amount of related content, training material, Software Development Kits (SDK) is freely available on our website. Our sales architects are available to provide free consultation on an as needed basis. Our professional services staff is available for contracted engagements. Typically, our customers will leverage the information available on our portals to integrate and customize the environment internally or customer can contract with us for this work.

8.20 (E) MARKETING PLAN

Describe your how you intend to market your Solutions to NASPO ValuePoint and Participating Entities.

Unisys Response:

Unisys designated sales and business development and our NASPO contract manager will perform this function. Additionally, we plan to take an active role in the NASPO community. We see ourselves bringing value to the NASPO community by sponsoring and attending NASPO events. We will raise awareness of NASPO in our own client base and work very closely with the State of Utah and the Participating Entities of this contract vehicle. Unisys has been a participant of NASCIO; we will continue to demonstrate our support through participation and sponsorships.

8.21 (E) RELATED VALUE-ADDED SERVICES TO CLOUD SOLUTIONS

Describe the valued-added services that you can provide as part of an awarded contract, e.g. consulting services pre- and post- implementation. Offerors may detail professional services in the RFP limited to assisting offering activities with initial setup, training and access to the services.

Value-add services at No Cost to the Participating Entities

If Unisys is awarded the contract, besides the onboarding consultation services, contract management, contract support organization, billing, and fee disbursements that are required by this contract, we will offer the following additional services at no cost to the Participating Entities:

- **Provisioning:** Unisys will provision the access login ID for the Participating Entity.
- **Billing Portal:** Unisys will provide a customized, online web-based portal for the Participating Entities to monitor and report on their consumption, utilization, current balances, etc.
- **Billing:** Unisys will bill Participating Entities each month.
- **Consultation:** Unisys will guide Participating Entities on how to locate publicly available learning tools and content from our partners.

Fee-based Public Sector Cloud based Solutions and Value-Add Professional Services:

Following is a summary of the public sector solutions and professional services offered by Unisys to the Participating Entities.

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud: Outlined below are the cloud based solution offerings from Unisys.

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
Unisys Justice Law Enforcement and Border Security Solutions for the Cloud		
Secure Image Management Solution	Unisys developed a highly secure solution to meet the forensic integrity standards and policy requirements of several of the largest police departments in the world. The secure storage of images and other multimedia content occurs in a tamperproof information archive – in effect, a digital evidence vault.	All
U-LEAF	U-LEAF provides a POLE-type data model (Person, Object, Location and Event) for the storage and recording of incidents and entities. The POLE model allows entities to be recorded in the system once. The recorded entities, however, can be linked to other entities and events as many times as necessary, to build the picture of an incident, or a network of associations.	All
Unisys Social Services Solutions for the Cloud		
Unisys 311	The Unisys “311 solution” is a Multi- Channel Citizen Service Delivery and Engagement solution that can be used by a Local Government or any Governmental Agency to deliver and manage non-emergency services, respond to inquiries and engage with their constituents. This offers multi-channel interaction including social media and mobile applications, supported by a knowledgebase and GIS systems and also interfaces with existing systems to manage work orders and service requests. The solution can be configured to meet specific needs and extend its functionality and flexibility through other apps, and integration to mobile apps, social media and other systems.	All
State/Local Government Enterprise Regulatory System (AMANDA)	The Unisys AMANDA Platform (from CSDC) is designed to provide a collection of back office functions such as Licensing, Permitting, Inspections, Land Use, Planning and a number of other functions in various form factors (desktop, mobile and tablet) and can also be deployed either on premise, hosted, or delivered via the Cloud. The solution includes the design, delivery, deployment and ongoing support and maintenance of the solution. The AMANDA platform can be configured specific to client needs in order to extend its functionality and interface with other systems as required.	All
Unisys Enterprise Content Management for the Cloud		
Infomage	Unisys Infomage is the Enterprise Content Management (ECM), Business Process Management (BPM) and Record Management (RM) solution for organizations looking to significantly improve business processes that depend heavily high-volume paper documents, documents generated from Internet transactions, Office documents, and other electronic documents that need to be accessed for automated and manual processes. Infomage can easily capture, manage, store, and access the content required for cases, inquiries, and process-centric work, regardless of	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	data structure or document origination, with a single intuitive user interface. Unisys InfoImage brings together ECM, imaging, workflow, document management, record management and integration technologies to form an integrated end-to-end solution.	
Unisys Horizontal Solutions for the Cloud		
Stealth	The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. We built a better way to deal with advanced threats for our clients by applying novel approaches to new threats. By substituting traditional hardware topology for software-based cryptography, our Stealth Microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users.	Private and Hybrid Cloud Only. For Public and Community Cloud, see AWS offering above. IaaS, PaaS and SaaS
VantagePoint	VantagePoint can also extend beyond the boundaries of IT with uses cases to support relevant business scenarios in, • Security: Strengthens the security posture of organizations through greater visibility into the Stealth enabled infrastructure. • Advanced Data Analytics: Cuts across data silos and sources relevant data to optimize the performance of Advanced Analytics solution. • Cloud and Infrastructure Services: Enables hybrid IT by seamlessly integrating with a variety of data sources and services and presenting them via a digital control panel, irrespective of technology underpinnings. • End User Services: Supports millennial users via personalized and secure access to relevant data and services, across endpoints. • Service Management: Cuts across ITSM platforms to improve service delivery through service automation, orchestration, and aggregation capabilities. • Facilities/Crisis Management: Helps organizations monitor their facilities and keep their Facilities Managers abreast of any security threats and risks. Also acts as a standard communication vehicle for users at times of crisis. (178) With VantagePoint, our clients gain a personalized, secure, and intuitive data and service aggregation platform that cuts across strategic and operational dimensions of business and accelerates digital transformation.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys Cloud Hosting	Unisys Hosted Private Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can provide, integrate and scale the cloud infrastructure to meet client's needs. Our management platform enables a single pane of glass to effectively manage client's workloads. Unisys Hybrid Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can source, integrate and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our management	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	platform enables a single pane of glass to effectively manage a variety of clouds within an organization, whether they be public or private.	
ServiceNow	ServiceNow offers enterprise service management software for human resources, law, facilities management, finance, marketing, and field operations in the cloud. ServiceNow has its STAR Self-Assessment available on the CSA's website. ServiceNow specializes in ITSM applications and provides forms-based workflow application development. ServiceNow has open integration options to variety platforms such as: Salesforce, SharePoint, and BMC Remedy Action Request System.	All
Unisys ClearPath Forward!	Unisys ClearPath Forward is an Intel based fabric computing platform designed to run mission critical applications that require predictable performance and low transport latency. The ClearPath Forward platform provides hardware partitioning technology similar to the Mainframes, but is designed to run Windows and Linux operating environments on commodity server hardware. The fabric interconnect included with the ClearPath Forward platform provides transport speeds at 56Gbps or higher. The ClearPath Forward platform fabric can be used in a hybrid mode with VMware.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys White Label Offerings	Unisys has strategic, long lasting relationships in the industry and is pleased to offer our partner products surrounding the cloud market space. This is an evolving list. Our current partner list includes: SalesForce, Oracle, Verizon, SHI, Decision Lens, Aptio, NetApp, VMware, EMC, NetApp, Birst, Okta, Box, WatchDox. We have created webpage on our public site that is updated with our list of current partners. The page is available here: http://www.unisys.com/ms/wsca-cloud-hosting/	All

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud.

Exhibit Y. Unisys Services Offerings: Outlined below are our professional services surrounding the cloud.

Unisys Professional Services	Description
Cloud Advisory Services	Unisys Cloud Advisory Services provide strategic and financial guidance on aligning IT with business objectives. This starts with a roadmap that outlines the vision of a hybrid IT based on a combination of existing data center, internal cloud, and external cloud resources to provide agility, flexibility, and control.
Data Center Planning, Design, and Implementation Services	Unisys Data Center Planning, Design, and Implementation Services offer a complete range of services that delivers cohesive, end-to-end optimization of data centers. With a wide range of services for discovery, analysis, optimization, virtualization, consolidation, and migration of data centers that can complement client efforts and fill gaps in skills and capacities and a combination of world-class people, processes, and technology with program and project management expertise, Unisys transforms clients' existing data centers to a business engine that provides agility at a lower cost.
CloudBuild Services	Unisys CloudBuild Services enable organizations to successfully build a cloud that is integrated with the overall business process, transforming their existing infrastructure to an agile IT-as-a-service model. The Unisys "8-Tracks" model, a 360-degree cloud view approach, covers eight critical data center domains. Together with ConOps, which includes industry best practices, Unisys enables a client's cloud infrastructure to meet the security, regulatory, and compliance requirements that enable us to deliver the most secure and reliable cloud in the marketplace.
Hybrid Cloud Strategy	Unisys Hybrid Cloud strategy helps organizations overcome key challenges when

Unisys Professional Services	Description
	planning to implement a cloud environment. Unisys provides governments and businesses with a comprehensive cloud architecture that gives clients cloud on their terms. We can source, integrate, and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our homegrown VantagePoint management platform enables a single pane of glass to effectively manage a variety of clouds in an organization, whether they be public or private.
Unisys Platform Services	
Platform Assessment Services	Unisys Platform Assessment Services enable governments and businesses to foster their platform adoption initiatives. Unisys analyzes organizations' business goals and objectives along with their technology landscape of their enterprise application portfolio and recommend a best-fit Platform Suitability Analysis to help them make the right decision in their move to PaaS.
Architecture Design Services	Unisys Architecture Design Services assist governments and businesses with designing an effective architecture suited for cloud. We provide several concepts and best practices that are essential to build highly scalable applications in the cloud – be it on-premises, public, or hybrid.
PaaS Enablement Services	Unisys PaaS Enablement Services analyze an organization's current set of applications and provide a strategy to use the best practices and engagement models of various platform providers such as Microsoft Azure, IBM Bluemix, SAP HCP, and SFDC AppCloud.
Application Migration Services	Unisys Application Migration Services assist organizations in their cloud migration activities by using tools and frameworks adhering to Unisys best practices that promise minimal downtime without affecting the day-to-day business processes.
Application Portability Services	Unisys Application Portability Services help organizations to develop cross-platform applications that can be scaled across multiple cloud platforms.
Application Development Services	Unisys Application Development Services help organizations to build rich, interactive applications focused on business logic and workflows using visual tools as well as cloud-based tools, architectures, and services that make their applications cloud ready.
IoT Development Services	Unisys IoT Development Services enable organizations to use IoT Services across various cloud platforms that enable apps to communicate and consume data collected by the connected devices, sensors, and gateways.
API Management Services	Unisys API Management Services provide a solution that addresses the aspects of the application programming interface (API) life cycle for on-premises and cloud environments and offer capabilities to create, run, manage, secure, and monetize APIs and microservices that deliver an integrated user experience and enable rapid deployment and simplified administration of APIs.
Horizontal Technology Integration Services	
Horizontal Technology Integration Services	Unisys is a world-recognized leader in integrating technology infrastructures across platforms and vendors. We particularly specialize in integrating solutions surrounding service desks, cloud provisioning, data center management, and the ITIL framework.
Unisys Service Management Services	
Maturity and Platform Technology Assessments	Unisys Service Management Consulting Services leverage subject matter experts with more than 15 years of experience in working and delivering services for ITIL and service management. We leverage that knowledge to provide best practices, lessons learned, and roadmap development for Service Management disciplines and benchmark them with industry and operational best practices.
Service Management Platform Implementation	Unisys brings experience in moving existing ITSM tool information, processes, and requirements to new cloud enterprise Service Management platforms. By leveraging industry best practices and structured methodologies, Unisys can establish transition and transformation plans that minimize risk to clients' ongoing operations and deliver speed to value with the new cloud-based Service Management platform.
Enterprise Cloud Service Management Platform Support	Unisys can provide ongoing day-to-day support in the ongoing management of a client's Service Management platform. This service allows the organization to focus on the business aspects instead of the routine activities necessary to support and maintain a platform. Unisys provides cost-effective solutions that leverage staff in many locations, driving the cost of servicing the environment to effective and efficient levels.

Exhibit Y. Unisys Services Offerings:

8.22 (E) SUPPORTING INFRASTRUCTURE

8.22.1 *Describe what infrastructure is required by the Purchasing Entity to support your Solutions or deployment models.*

8.22.2 *If required, who will be responsible for installation of new infrastructure and who will incur those costs?*

Unisys Response:

Our proposed cloud models and service delivery options across all providers and solution offerings require only a computer connected to the Internet. The Participating Entity will install this infrastructure and incur those costs.

8.23 (E) ALIGNMENT OF CLOUD COMPUTING REFERENCE ARCHITECTURE

Clarify how your architecture compares to the NIST Cloud Computing Reference Architecture, in particular, to describe how they align with the three domains e.g. Infrastructure as a Service (IaaS), Software as a Service (SaaS), and Platform as a Service (PaaS).

Unisys Response:

For information on the public and the community IaaS, PaaS, and SaaS domains, refer to the responses to RFP requirements 8.6.1, 8.6.2, and 8.6.6. We have provided our comprehensive list of compliance certifications. In the cases relevant to this proposal, the architecture provided by Amazon, Microsoft, and Google exceed the NIST Cloud Computing Reference Architecture requirements. For the private and hybrid cloud environments, Unisys builds and integrates the environment from the ground up, based on client's requirements, which typically exceed NIST requirements.

7.0 CONFIDENTIAL, PROTECTED OR PROPRIETARY INFORMATION

Section Title: Confidential, Protected or Proprietary Information. All confidential, protected or proprietary Information must be included in this section of proposal response. Do not incorporate protected information throughout the Proposal. Rather, provide a reference in the proposal response directing Lead State to the specific area of this protected Information section.

- *If there is no protected information, write "None" in this section.*

Failure to comply with this Section and Section 3.13 of the RFP releases the Lead State, NASPO ValuePoint, and Participating Entities from any obligation or liability arising from the inadvertent release of Offeror information.

Unisys Response:

None.

PRODUCTS & SERVICES

- Amazon EC2 >
- Product Details >
- Instances >
- Pricing >
- Purchasing Options >
- Developer Resources >
- FAQs >
- Getting Started >
- Amazon EC2 Run Command >

RELATED LINKS

- Amazon EC2 Dedicated Hosts
- Amazon EC2 Spot Instances
- Amazon EC2 Reserved Instances
- Amazon EC2 Dedicated Instances
- Windows Instances
- VM Import/Export
- AWS Management Portal for vCenter Management Console

Amazon EC2 Service Level Agreement

Last Updated June 1, 2013

This Amazon EC2 Service Level Agreement (“SLA”) is a policy governing the use of Amazon Elastic Compute Cloud (“Amazon EC2”) and Amazon Elastic Block Store (“Amazon EBS”) under the terms of the Amazon Web Services Customer Agreement (the “AWS Agreement”) between Amazon Web Services, Inc. and its affiliates (“AWS”, “us” or “we”) and users of AWS’ services (“you”). This SLA applies separately to each account using Amazon EC2 or Amazon EBS. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Service Commitment

AWS will use commercially reasonable efforts to make Amazon EC2 and Amazon EBS each available with a Monthly Uptime Percentage (defined below) of at least 99.95%, in each case during any monthly billing cycle (the “Service Commitment”). In the event Amazon EC2 or Amazon EBS does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- “Monthly Uptime Percentage” is calculated by subtracting from 100% the percentage of minutes during the month in which Amazon EC2 or Amazon EBS, as applicable, was in the state of “Region Unavailable.” Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon EC2 SLA Exclusion (defined below).
- “Region Unavailable” and “Region Unavailability” mean that more than one Availability Zone in which you are running an instance, within the same Region, is “Unavailable” to you.
- “Unavailable” and “Unavailability” mean:

 Menu

English ▼

My Account ▼

[Create an AWS Account](#)[Discussion Forum](#)[Amazon EC2 SLA](#)[Run Command](#)

- For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue.

- A “Service Credit” is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Commitments and Service Credits

Service Credits are calculated as a percentage of the total charges paid by you (excluding one-time payments such as upfront payments made for Reserved Instances) for either Amazon EC2 or Amazon EBS (whichever was Unavailable, or both if both were Unavailable) in the Region affected for the monthly billing cycle in which the Region Unavailability occurred in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0%	30%

We will apply any Service Credits only against future Amazon EC2 or Amazon EBS payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the Unavailability occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar (\$1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability, non-performance, or other failure by us to provide Amazon EC2 or Amazon EBS is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you must submit a claim by [opening a case in the AWS Support Center](#). To be eligible, the credit request must be received by us

1. the words “SLA Credit Request” in the subject line;

2. the dates and times of each Unavailability incident that you are claiming;
3. the affected EC2 instance IDs or the affected EBS volume IDs; and
4. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the Monthly Uptime Percentage of such request is confirmed by us and is less than the Service Commitment, then we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon EC2 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon EC2 or Amazon EBS, or any other Amazon EC2 or Amazon EBS performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon EC2 or Amazon EBS; (iii) that result from any actions or inactions of you or any third party, including failure to acknowledge a recovery volume; (iv) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); (v) that result from failures of individual instances or volumes not attributable to Region Unavailability; (vi) that result from any maintenance as provided for pursuant to the AWS Agreement; or (vii) arising from our suspension and termination of your right to use Amazon EC2 or Amazon EBS in accordance with the AWS Agreement (collectively, the “Amazon EC2 SLA Exclusions”). If availability is impacted by factors other than those used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

[Prior Version\(s\)](#)

**Menu****English** ▼**My Account** ▼**Create an AWS Account****Create a Free Account** [AWS on Twitter](#) [AWS on Facebook](#) [AWS on Google+](#) [AWS Blog](#) [What's New? RSS](#)**AWS & Cloud Computing**[What is Cloud Computing?](#)[Products & Services](#)[Customer Success](#)[Economics Center](#)[Architecture Center](#)[Security Center](#)[What's New](#)[Whitepapers](#)[AWS Blog](#)[Events](#)[Sustainable Energy](#)[Press Releases](#)[Analyst Reports](#)[Legal](#)**Solutions**[Websites & Website Hosting](#)[Business Applications](#)[Backup & Recovery](#)[Disaster Recovery](#)[Data Archive](#)[Big Data](#)[High Performance Computing](#)[Mobile Services](#)[Digital Marketing](#)[Game Development](#)[Digital Media](#)[Government & Education](#)**Resources & Training**[Developers](#)[Java on AWS](#)[JavaScript on AWS](#)[Mobile on AWS](#)[PHP on AWS](#)

 Menu

English ▼ My Account ▼

Create an AWS Account

[SDKs & Tools](#)
[AWS Marketplace](#)
[User Groups](#)
[Support Plans](#)
[Service Health Dashboard](#)
[Discussion Forums](#)
[FAQs](#)
[Documentation](#)
[Articles & Tutorials](#)
[Test Drives](#)

Manage Your Account
[Management Console](#)
[Billing & Cost Management](#)
[Personal Information](#)
[Payment Method](#)
[AWS Identity & Access Management](#)
[Security Credentials](#)
[Request Service Limit Increases](#)
[Contact Us](#)

Amazon Web Services is Hiring.

Amazon Web Services (AWS) is a dynamic, growing business unit within Amazon.com. We are currently hiring Software Development Engineers, Product Managers, Account Managers, Solutions Architects, Support Engineers, System Engineers, Designers and more. Visit our [Careers](#) page or our [Developer-specific Careers](#) page to learn more.

Amazon Web Services is an Equal Opportunity Employer.

Language [Deutsch](#) | [English](#) | [Español](#) | [Français](#) | [Italiano](#) | [Português](#) | [Русский](#) | [日本語](#) | [한국어](#) | [中文 \(简体\)](#)
[中文 \(繁體\)](#)

[Site Terms](#) | [Privacy](#)

© 2016, Amazon Web Services, Inc. or its affiliates. All rights reserved.

SLA summary for Azure services

Last updated: September 2015

Azure Active Directory

We guarantee at least 99.9% availability of the Azure Active Directory Basic and Premium services. The services are considered available in the following scenarios:

- Users are able to login to the service, login to the Access Panel, access applications on the Access Panel and reset passwords.
- IT administrators are able to create, read, write and delete entries in the directory or provision or de-provision users to applications in the directory.

No SLA is provided for the Free tier of Azure Active Directory.

Download Active Directory SLA (<http://www.microsoftvolumelicensing.com/DocumentSearch.aspx?Mode=3&DocumentTypeId=37>)

View full details (</en-us/support/legal/sla/active-directory/>)

API Management

- We guarantee that API Management Service instances running in the Standard tier will respond to requests to perform operations at least 99.9% of the time.
- We guarantee that API Management Service instances running in the Premium tier deployed across two or more regions will respond to requests to perform operations at least 99.95% of the time.

No SLA is provided for the Developer tier of the API Management Service.

View full details (</en-us/support/legal/sla/api-management/>)

App Service

We guarantee that Web Apps running in a customer subscription will be available 99.95% of the time. No SLA is provided for Mobile Apps, Logic Apps, or API Apps while such services are still in Preview or for Apps under either the Free or Shared tiers.

View full details (</en-us/support/legal/sla/app-service/>)

Application Gateway

We guarantee that each Application Gateway Cloud Service having two or more medium or larger instances will be available at least 99.9% of the time.

[View full details \(/en-us/support/legal/sla/application-gateway/\)](/en-us/support/legal/sla/application-gateway/)

Automation

We guarantee that at least 99.9% of runbook jobs will start within 30 minutes of their planned start times.

We guarantee at least 99.9% availability of the Azure Automation DSC agent service.

No SLA is provided for the Free tier of Azure Automation.

[View full details \(/en-us/support/legal/sla/automation/\)](/en-us/support/legal/sla/automation/)

Backup

We guarantee at least 99.9% availability of the backup and restore functionality of the Azure Backup service.

[View full details \(/en-us/support/legal/sla/backup/\)](/en-us/support/legal/sla/backup/)

BizTalk Services

We guarantee that at least 99.9% of the time customers will have connectivity between their BizTalk Service Environments in the Basic, Standard and Premium tiers and our Internet gateway. We do not offer an SLA for the BizTalk Services Developer tier.

[View full details \(/en-us/support/legal/sla/biztalk-services/\)](/en-us/support/legal/sla/biztalk-services/)

Cache

We guarantee at least 99.9% of the time that customers will have connectivity between the Cache endpoints and our Internet gateway.

[View full details \(/en-us/support/legal/sla/cache/\)](/en-us/support/legal/sla/cache/)

CDN

We guarantee that at least 99.9% of the time CDN will respond to client requests and deliver the requested content without error. We will review and accept data from any commercially reasonable independent measurement system that you choose to monitor your content. You must select a set of agents from the measurement system's list of standard agents that are generally available and represent at least five geographically diverse locations in major worldwide metropolitan areas

(excluding PR of China).

[View full details \(/en-us/support/legal/sla/cdn/\)](/en-us/support/legal/sla/cdn/)

Cloud Services and Virtual Machines

- For Cloud Services, we guarantee that when you deploy two or more role instances in different fault and upgrade domains, your Internet facing roles will have external connectivity at least 99.95% of the time.
- For all Internet facing Virtual Machines that have two or more instances deployed in the same Availability Set, we guarantee you will have external connectivity at least 99.95% of the time.

[View full details \(/en-us/support/legal/sla/virtual-machines/\)](/en-us/support/legal/sla/virtual-machines/)

DocumentDB

We guarantee at least 99.95% of the time we will successfully process requests to perform operations against DocumentDB Resources.

[View full details \(/en-us/support/legal/sla/documentdb/\)](/en-us/support/legal/sla/documentdb/)

ExpressRoute

We guarantee a minimum of 99.9% ExpressRoute dedicated circuit availability.

[View full details \(/en-us/support/legal/sla/expressroute/\)](/en-us/support/legal/sla/expressroute/)

HDInsight

For HDInsight, we guarantee that any HDInsight Cluster that you deploy will have external connectivity at least 99.9% of the time over a monthly billing cycle.

[View full details \(/en-us/support/legal/sla/hdinsight/\)](/en-us/support/legal/sla/hdinsight/)

IoT Hub

- For IoT Hub, we promise that at least 99.9% of the time deployed IoT hubs will be able to send messages to and receive messages from registered devices and the Service will be able to perform create, read, update, and delete operations on IoT hubs.
- No SLA is provided for the Free Tier of IoT Hub.

[View full details \(/en-us/support/legal/sla/iot-hub/\)](/en-us/support/legal/sla/iot-hub/)

Key Vault

We guarantee that we will process Key Vault transactions within 5 seconds at least 99.9% of the time.

[View full details \(/en-us/support/legal/sla/key-vault/\)](/en-us/support/legal/sla/key-vault/)

Machine Learning

- For the Request Response Service (RRS), we guarantee 99.95% availability of API transactions.
- For the Batch Execution Service (BES) and management APIs, we guarantee 99.9% availability of API transactions.

No SLA is provided for the Free tier of Machine Learning.

[View full details \(/en-us/support/legal/sla/machine-learning/\)](/en-us/support/legal/sla/machine-learning/)

Media Services

- For Media Services Encoding, we guarantee 99.9% availability of REST API transactions.
- For Streaming, we will successfully service requests with a 99.9% availability guarantee for existing media content when at least one Streaming Unit is purchased.
- For Live Channels, we guarantee that running Channels will have external connectivity at least 99.9% of the time.
- For Content Protection, we guarantee that we will successfully fulfill key requests at least 99.9% of the time.
- For Indexer, we will successfully service Indexer Task requests processed with an Encoding Reserved Unit 99.9% of the time.

[View full details \(/en-us/support/legal/sla/media-services/\)](/en-us/support/legal/sla/media-services/)

Mobile Engagement

We guarantee at least 99.9% availability of REST API calls to the Azure Mobile Engagement Service.

No SLA is provided for the Free tier.

[View full details \(/en-us/support/legal/sla/mobile-engagement/\)](/en-us/support/legal/sla/mobile-engagement/)

Mobile Services

We guarantee 99.9% availability of REST API calls to all provisioned Azure Mobile Services running in Standard and Premium tiers in a customer subscription. No SLA is provided for the Free tier of Mobile Services.

[View full details \(/en-us/support/legal/sla/mobile-services/\)](/en-us/support/legal/sla/mobile-services/)

Multi-Factor Authentication

We guarantee 99.9% availability of Azure Multi-Factor Authentication. The service is considered unavailable when it is unable to receive or process authentication requests for the Multi-Factor authentication provider deployed in a customer subscription.

[View full details \(/en-us/support/legal/sla/multi-factor-authentication/\)](/en-us/support/legal/sla/multi-factor-authentication/)

Operational Insights

We guarantee that at least 99.9% of the time, log data will be indexed within six hours of the data being queued for indexing by the Operational Insights Service.

No SLA is provided for the Free tier of Azure Operational Insights.

[View full details \(/en-us/support/legal/sla/operational-insights/\)](/en-us/support/legal/sla/operational-insights/)

RemoteApp

We guarantee at least 99.9% of the time users will have connectivity to their applications through the RemoteApp service. No SLA is provided for the Free tier of RemoteApp.

[View full details \(/en-us/support/legal/sla/remoteapp/\)](/en-us/support/legal/sla/remoteapp/)

Scheduler

We guarantee that at least 99.9% of the time all scheduled jobs will initiate within 30 minutes of their planned execution times.

[View full details \(/en-us/support/legal/sla/scheduler/\)](/en-us/support/legal/sla/scheduler/)

Search

We guarantee at least 99.9% availability for index query requests when an Azure Search Service Instance is configured with two or more replicas, and index update requests when an Azure Search Service Instance is configured with three or more replicas. No SLA is provided for the Free tier.

[View full details \(/en-us/support/legal/sla/search/\)](/en-us/support/legal/sla/search/)

Service Bus

- For Service Bus Relays, we guarantee that at least 99.9% of the time, properly configured applications will be able to establish a connection to a deployed Relay.
- For Service Bus Queues and Topics, we guarantee that at least 99.9% of the time, properly configured applications will be able to send or receive messages or perform other operations on a deployed Queue or Topic.

- For Service Bus Basic and Standard Notification Hub tiers, we guarantee that at least 99.9% of the time, properly configured applications will be able to send notifications or perform registration management operations with respect to a Notification Hub.
- For Event Hubs Basic and Standard tiers, we guarantee that at least 99.9% of the time, properly configured applications will be able to send or receive messages or perform other operations on the Event Hub.

[View full details \(/en-us/support/legal/sla/service-bus/\)](/en-us/support/legal/sla/service-bus/)

Site Recovery

- For each Protected Instance configured for On-Premises-to-On-Premises Failover, we guarantee at least 99.9% availability of the Site Recovery service.
- For each Protected Instance configured for On-Premises-to-Azure planned and unplanned Failover, we guarantee a four-hour Recovery Time Objective for unencrypted Protected Instances, and a six-hour Recovery Time Objective for encrypted Protected Instance, depending on the size of the Protected Instance.

[View full details \(/en-us/support/legal/sla/site-recovery/\)](/en-us/support/legal/sla/site-recovery/)

SQL Database

Web and Business Tiers

We guarantee at least 99.9% of the time customers will have connectivity between their Web or Business Microsoft Azure SQL Database and our Internet gateway.

Basic, Standard, and Premium Tiers

We guarantee at least 99.99% of the time customers will have connectivity between their Basic, Standard, or Premium Microsoft Azure SQL Database and our Internet gateway.

[View full details \(/en-us/support/legal/sla/sql-database/\)](/en-us/support/legal/sla/sql-database/)

Storage

- We guarantee that at least 99.99% of the time, we will successfully process requests to read data from Read Access-Geo Redundant Storage (RA-GRS) Accounts, provided that failed attempts to read data from the primary region are retried on the secondary region.
- We guarantee that at least 99.9% of the time, we will successfully process requests to read data from Locally Redundant Storage (LRS), Zone Redundant Storage (ZRS), and Geo Redundant Storage (GRS) Accounts.
- We guarantee that at least 99.9% of the time, we will successfully process requests to write data to

Locally Redundant Storage (LRS), Zone Redundant Storage (ZRS), and Geo Redundant Storage (GRS) Accounts and Read Access-Geo Redundant Storage (RA-GRS) Accounts.

[View full details \(/en-us/support/legal/sla/storage/\)](/en-us/support/legal/sla/storage/)

StorSimple

We guarantee at least 99.9% availability of the backup, cloud tiering, and restore functionality of the Azure StorSimple service.

[View full details \(/en-us/support/legal/sla/storsimple/\)](/en-us/support/legal/sla/storsimple/)

Stream Analytics

- We guarantee at least 99.9% availability of the Stream Analytics API.
- We guarantee that 99.9% of the time, deployed Stream Analytics jobs will be either processing data or available to process data.

[View full details \(/en-us/support/legal/sla/stream-analytics/\)](/en-us/support/legal/sla/stream-analytics/)

Traffic Manager

We guarantee that DNS queries will receive a valid response from at least one of our Azure Traffic Manager name server clusters at least 99.99% of the time.

[View full details \(/en-us/support/legal/sla/traffic-manager/\)](/en-us/support/legal/sla/traffic-manager/)

Visual Studio Team Services

- We guarantee at least 99.9% availability of Visual Studio Team Services for paid Visual Studio Team Services users to access the associated Visual Studio Team Services account.
- We guarantee at least 99.9% availability to execute build operations using the paid Visual Studio Team Services Build Service.
- We guarantee at least 99.9% availability to execute load testing operations using the paid Visual Studio Team Services Load Testing Service.
- We guarantee at least 99.9% availability to execute build and deployment operations using the paid Visual Studio Team Services Build & Deployment Service.

[View full details \(/en-us/support/legal/sla/visual-studio-team-services/\)](/en-us/support/legal/sla/visual-studio-team-services/)

VPN Gateway

We guarantee 99.9% availability for each VPN Gateway.

[View full details \(/en-us/support/legal/sla/vpn-gateway/\)](/en-us/support/legal/sla/vpn-gateway/)

Microsoft will provide at least 90 days' notice for adverse material changes to any of the SLAs listed above.

Availability for all Azure services is calculated over a monthly billing cycle. Click here (<http://go.microsoft.com/fwlink/?linkid=517017&clid=0x409>) to download SLA for most Microsoft Azure Services. The SLA for Active Directory can be found here (<http://www.microsoftvolumelicensing.com/DocumentSearch.aspx?Mode=3&DocumentTypeId=37>).



Amazon Web Services: Risk and Compliance

January 2016

(Consult <http://aws.amazon.com/compliance/aws-whitepapers/>

for the latest version of this paper)

This document is intended to provide information to assist AWS customers with integrating AWS into their existing control framework supporting their IT environment. This document includes a basic approach to evaluating AWS controls and provides information to assist customers with integrating control environments. This document also addresses AWS-specific information around general cloud computing compliance questions.

Table of Contents

Risk and Compliance Overview	3
<i>Shared Responsibility Environment</i>	<i>3</i>
<i>Strong Compliance Governance</i>	<i>4</i>
Evaluating and Integrating AWS Controls	4
<i>AWS IT Control Information</i>	<i>5</i>
<i>AWS Global Regions</i>	<i>5</i>
AWS Risk and Compliance Program	6
<i>Risk Management</i>	<i>6</i>
<i>Control Environment</i>	<i>6</i>
<i>Information Security</i>	<i>7</i>
AWS Certifications, Programs, Reports, and Third-Party Attestations	7
<i>CJIS</i>	<i>7</i>
<i>CSA</i>	<i>7</i>
<i>Cyber Essentials Plus</i>	<i>8</i>
<i>DoD SRG Levels 2 and 4</i>	<i>8</i>
<i>FedRAMP SM</i>	<i>8</i>
<i>FERPA</i>	<i>9</i>
<i>FIPS 140-2</i>	<i>9</i>
<i>FISMA and DIACAP</i>	<i>9</i>
<i>GxP</i>	<i>10</i>
<i>HIPAA</i>	<i>10</i>
<i>IRAP</i>	<i>11</i>
<i>ISO 9001</i>	<i>11</i>
<i>ISO 27001</i>	<i>12</i>
<i>ISO 27017</i>	<i>14</i>
<i>ISO 27018</i>	<i>14</i>
<i>ITAR</i>	<i>15</i>
<i>MPAA</i>	<i>16</i>
<i>MTCS Tier 3 Certification</i>	<i>16</i>

<i>NIST</i>	16
<i>PCI DSS Level 1</i>	17
<i>SOC 1/ISAE 3402</i>	17
<i>SOC 2</i>	19
<i>SOC 3</i>	19
<i>Key Compliance Questions and AWS</i>	20
AWS Contact	24
Appendix A: CSA Consensus Assessments Initiative Questionnaire v3.0.1	25
Appendix B: AWS alignment with the Australian Signals Directorate (ASD) Cloud Computing Security Considerations	62
Appendix C: Glossary of Terms	82

Risk and Compliance Overview

AWS and its customers share control over the IT environment, both parties have responsibility for managing the IT environment. AWS' part in this shared responsibility includes providing its services on a highly secure and controlled platform and providing a wide array of security features customers can use. The customers' responsibility includes configuring their IT environments in a secure and controlled manner for their purposes. While customers don't communicate their use and configurations to AWS, AWS does communicate its security and control environment relevant to customers. AWS does this by doing the following:

- Obtaining industry certifications and independent third-party attestations described in this document
- Publishing information about the AWS security and control practices in whitepapers and web site content
- Providing certificates, reports, and other documentation directly to AWS customers under NDA (as required)

For a more detailed description of AWS security please see:

[AWS Security Center](https://aws.amazon.com/security/): <https://aws.amazon.com/security/>

For a more detailed description of AWS Compliance please see

[AWS Compliance page](https://aws.amazon.com/compliance/): <https://aws.amazon.com/compliance/>

Additionally, The [AWS Overview of Security Processes Whitepaper](#) covers AWS' general security controls and service-specific security.

Shared Responsibility Environment

Moving IT infrastructure to AWS services creates a model of shared responsibility between the customer and AWS. This shared model can help relieve customer's operational burden as AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall. Customers should carefully consider the services they choose as their responsibilities vary depending on the services used, the integration of those



services into their IT environment, and applicable laws and regulations. It is possible for customers to enhance security and/or meet their more stringent compliance requirements by leveraging technology such as host based firewalls, host based intrusion detection/prevention, encryption and key management. The nature of this shared responsibility also provides the flexibility and customer control that permits the deployment of solutions that meet industry-specific certification requirements.

This customer/AWS shared responsibility model also extends to IT controls. Just as the responsibility to operate the IT environment is shared between AWS and its customers, so is the management, operation and verification of IT controls shared. AWS can help relieve customer burden of operating controls by managing those controls associated with the physical infrastructure deployed in the AWS environment that may previously have been managed by the customer. As every customer is deployed differently in AWS, customers can take advantage of shifting management of certain IT controls to AWS which results in a (new) distributed control environment. Customers can then use the AWS control and compliance documentation available to them (described in the [AWS Certifications and Third-party Attestations](#) section of this document) to perform their control evaluation and verification procedures as required.

The next section provides an approach on how AWS customers can evaluate and validate their distributed control environment effectively.

Strong Compliance Governance

As always, AWS customers are required to continue to maintain adequate governance over the entire IT control environment regardless of how IT is deployed. Leading practices include an understanding of required compliance objectives and requirements (from relevant sources), establishment of a control environment that meets those objectives and requirements, an understanding of the validation required based on the organization's risk tolerance, and verification of the operating effectiveness of their control environment. Deployment in the AWS cloud gives enterprises different options to apply various types of controls and various verification methods.

Strong customer compliance and governance might include the following basic approach:

1. Review information available from AWS together with other information to understand as much of the entire IT environment as possible, and then document all compliance requirements.
2. Design and implement control objectives to meet the enterprise compliance requirements.
3. Identify and document controls owned by outside parties.
4. Verify that all control objectives are met and all key controls are designed and operating effectively.

Approaching compliance governance in this manner will help companies gain a better understanding of their control environment and will help clearly delineate the verification activities to be performed.

Evaluating and Integrating AWS Controls

AWS provides a wide range of information regarding its IT control environment to customers through white papers, reports, certifications, and other third-party attestations. This documentation assists customers in understanding the controls in place relevant to the AWS services they use and how those controls have been validated. This information also assists customers in their efforts to account for and to validate that controls in their extended IT environment are operating effectively.

Traditionally, the design and operating effectiveness of control objectives and controls are validated by internal and/or external auditors via process walkthroughs and evidence evaluation. Direct observation/verification, by the customer or customer's external auditor, is generally performed to validate controls. In the case where



service providers, such as AWS, are used, companies request and evaluate third-party attestations and certifications in order to gain reasonable assurance of the design and operating effectiveness of control objective and controls. As a result, although customer's key controls may be managed by AWS, the control environment can still be a unified framework where all controls are accounted for and are verified as operating effectively. Third-party attestations and certifications of AWS can not only provide a higher level of validation of the control environment, but may relieve customers of the requirement to perform certain validation work themselves for their IT environment in the AWS cloud.

AWS IT Control Information

AWS provides IT control information to customers in the following two ways:

1. **Specific control definition.** AWS customers are able to identify key controls managed by AWS. Key controls are critical to the customer's control environment and require an external attestation of the operating effectiveness of these key controls in order to comply with compliance requirements—such as the annual financial audit. For this purpose, AWS publishes a wide range of specific IT controls in its Service Organization Controls 1 (SOC 1) Type II report. The SOC 1 report, formerly the Statement on Auditing Standards (SAS) No. 70, Service Organizations report, is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). The SOC 1 audit is an in-depth audit of both the design and operating effectiveness of AWS' defined control objectives and control activities (which include control objectives and control activities over the part of the infrastructure AWS manages). "Type II" refers to the fact that each of the controls described in the report are not only evaluated for adequacy of design, but are also tested for operating effectiveness by the external auditor. Because of the independence and competence of AWS' external auditor, controls identified in the report should provide customers with a high level of confidence in AWS' control environment. AWS' controls can be considered designed and operating effectively for many compliance purposes, including Sarbanes-Oxley (SOX) Section 404 financial statement audits. Leveraging SOC 1 Type II reports is also generally permitted by other external certifying bodies (e.g., ISO 27001 auditors may request a SOC 1 Type II report in order to complete their evaluations for customers).

Other specific control activities relate to AWS' Payment Card Industry (PCI) and Federal Information Security Management Act (FISMA) compliance. As discussed below, AWS is compliant with FISMA Moderate standards and with the PCI Data Security Standard. These PCI and FISMA standards are very prescriptive and require independent validation that AWS adheres to the published standard.

2. **General control standard compliance.** If an AWS customer requires a broad set of control objectives to be met, evaluation of AWS' industry certifications may be performed. With the AWS ISO 27001 certification, AWS complies with a broad, comprehensive security standard and follows best practices in maintaining a secure environment. With the PCI Data Security Standard (PCI DSS), AWS complies with a set of controls important to companies that handle credit card information. With AWS' compliance with the FISMA standards, AWS complies with a wide range of specific controls required by US government agencies. Compliance with these general standards provides customers with in-depth information on the comprehensive nature of the controls and security processes in place and can be considered when managing compliance.

AWS Global Regions

Data centers are built in clusters in various global regions. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Frankfurt), EU (Ireland), Asia Pacific (Seoul) Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing), and South America (Sao Paulo).



AWS Risk and Compliance Program

AWS provides information about its risk and compliance program to enable customers to incorporate AWS controls into their governance framework. This information can assist customers in documenting a complete control and governance framework with AWS included as an important part of that framework.

Risk Management

AWS management has developed a strategic business plan which includes risk identification and the implementation of controls to mitigate or manage risks. AWS management re-evaluates the strategic business plan at least biannually. This process requires management to identify risks within its areas of responsibility and to implement appropriate measures designed to address those risks.

In addition, the AWS control environment is subject to various internal and external risk assessments. AWS' Compliance and Security teams have established an information security framework and policies based on the Control Objectives for Information and related Technology (COBIT) framework and have effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.1, and the National Institute of Standards and Technology (NIST) Publication 800-53 Rev 3 (Recommended Security Controls for Federal Information Systems). AWS maintains the security policy, provides security training to employees, and performs application security reviews. These reviews assess the confidentiality, integrity, and availability of data, as well as conformance to the information security policy.

AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. These scans are done in a manner for the health and viability of the underlying AWS infrastructure and are not meant to replace the customer's own vulnerability scans required to meet their specific compliance requirements. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these types of scans can be initiated by submitting a request via the [AWS Vulnerability / Penetration Testing Request Form](#).

Control Environment

AWS manages a comprehensive control environment that includes policies, processes and control activities that leverage various aspects of Amazon's overall control environment. This control environment is in place for the secure delivery of AWS' service offerings. The collective control environment encompasses the people, processes, and technology necessary to establish and maintain an environment that supports the operating effectiveness of AWS' control framework. AWS has integrated applicable cloud-specific controls identified by leading cloud computing industry bodies into the AWS control framework. AWS continues to monitor these industry groups for ideas on which leading practices can be implemented to better assist customers with managing their control environment.

The control environment at Amazon begins at the highest level of the Company. Executive and senior leadership play important roles in establishing the Company's tone and core values. Every employee is provided with the Company's Code of Business Conduct and Ethics and completes periodic training. Compliance audits are performed so that employees understand and follow the established policies.



The AWS organizational structure provides a framework for planning, executing and controlling business operations. The organizational structure assigns roles and responsibilities to provide for adequate staffing, efficiency of operations, and the segregation of duties. Management has also established authority and appropriate lines of reporting for key personnel. Included as part of the Company's hiring verification processes are education, previous employment, and, in some cases, background checks as permitted by law and regulation for employees commensurate with the employee's position and level of access to AWS facilities. The Company follows a structured on-boarding process to familiarize new employees with Amazon tools, processes, systems, policies and procedures.

Information Security

AWS has implemented a formal information security program designed to protect the confidentiality, integrity, and availability of customers' systems and data. AWS publishes a security whitepaper that is available on the public website that addresses how AWS can help customers secure their data.

AWS Certifications, Programs, Reports, and Third-Party Attestations

AWS engages with external certifying bodies and independent auditors to provide customers with considerable information regarding the policies, processes, and controls established and operated by AWS.

CJIS

AWS complies with the FBI's Criminal Justice Information Services (CJIS) standard. We sign CJIS security agreements with our customers, including allowing or performing any required employee background checks according to the [CJIS Security Policy](#).

Law enforcement customers (and partners who manage CJI) are taking advantage of AWS services to improve the security and protection of CJI data, using the advanced security services and features of AWS, such as activity logging ([AWS CloudTrail](#)), encryption of data in motion and at rest (S3's Server-Side Encryption with the option to bring your own key), comprehensive key management and protection ([AWS Key Management Service](#) and [CloudHSM](#)), and integrated permission management (IAM federated identity management, multi-factor authentication).

AWS has created a Criminal Justice Information Services (CJIS) [Workbook](#) in a security plan template format aligned to the CJIS Policy Areas. Additionally, a CJIS Whitepaper has been developed to help guide customers in their journey to cloud adoption.

Visit the CJIS Hub Page: <https://aws.amazon.com/compliance/cjis/>

CSA

In 2011, the Cloud Security Alliance (CSA) launched [STAR](#), an initiative to encourage transparency of security practices within cloud providers. The [CSA Security, Trust & Assurance Registry](#) (STAR) is a free, publicly accessible registry that documents the security controls provided by various cloud computing offerings, thereby helping users assess the security of cloud providers they currently use or are considering contracting with. [AWS is a CSA STAR registrant](#) and has completed the Cloud Security Alliance (CSA) Consensus Assessments Initiative Questionnaire (CAIQ). This CAIQ published by the CSA provides a way to reference and



document what security controls exist in AWS' Infrastructure as a Service offerings. The CAIQ provides 298 questions a cloud consumer and cloud auditor may wish to ask of a cloud provider.

See: [Appendix A: CSA Consensus Assessments Initiative Questionnaire v3.0.1](#)

Cyber Essentials Plus

[Cyber Essentials Plus](#) is a UK Government-backed, industry-supported certification scheme introduced in the UK to help organizations demonstrate operational security against common cyber-attacks.

It demonstrates the baseline controls AWS implements to mitigate the risk from common Internet-based threats, within the context of the UK Government's "[10 Steps to Cyber Security](#)". It is backed by industry, including the Federation of Small Businesses, the Confederation of British Industry and a number of insurance organizations that offer incentives for businesses holding this certification.

Cyber Essentials sets out the necessary technical controls; the related assurance framework shows how the independent assurance process works for Cyber Essentials Plus certification through an annual external assessment conducted by an accredited assessor. Due to the regional nature of the certification, the certification scope is limited to EU (Ireland) region.

DoD SRG Levels 2 and 4

[The Department of Defense \(DoD\) Cloud Security Model \(SRG\)](#) provides a formalized assessment and authorization process for cloud service providers (CSPs) to gain a DoD Provisional Authorization, which can subsequently be leveraged by DoD customers. A Provisional Authorization under the SRG provides a reusable certification that attests to our compliance with DoD standards, reducing the time necessary for a DoD mission owner to assess and authorize one of their systems for operation on AWS. AWS currently holds provisional authorizations at Levels 2 and 4 of the SRG.

Additional information of the security control baselines defined for [Levels 2, 4, 5, and 6 can be found at: \[http://iase.disa.mil/cloud_security/Pages/index.aspx\]\(http://iase.disa.mil/cloud_security/Pages/index.aspx\)](#).

Visit the DoD Hub Page: <https://aws.amazon.com/compliance/dod/>

FedRAMP SM

AWS is a Federal Risk and Authorization Management Program (FedRAMPSM) Compliant Cloud Service Provider. AWS has completed the testing performed by a FedRAMPSM accredited Third-Party Assessment Organization (3PAO) and has been granted two Agency Authority to Operate (ATOs) by the US Department of Health and Human Services (HHS) after demonstrating compliance with FedRAMPSM requirements at the Moderate impact level. All U.S. government agencies can leverage the AWS Agency ATO packages stored in the FedRAMPSM repository to evaluate AWS for their applications and workloads, provide authorizations to use AWS, and transition workloads into the AWS environment. The two FedRAMPSM Agency ATOs encompass all U.S. regions (the AWS GovCloud (US) region and the AWS US East/West regions).

The following services are in the accreditation boundary for the regions stated above:



- [Amazon Redshift](#). Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools.
- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#). Amazon EC2 provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.
- [Amazon Simple Storage Service \(S3\)](#). Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.
- [Amazon Virtual Private Cloud \(VPC\)](#). Amazon VPC provides the ability for you to provision a logically isolated section of AWS where you can launch AWS resources in a virtual network that you define.
- [Amazon Elastic Block Store \(EBS\)](#). Amazon EBS provides highly available, highly reliable, predictable storage volumes that can be attached to a running Amazon EC2 instance and exposed as a device within the instance.
- [AWS Identity and Access Management \(IAM\)](#). IAM enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources.

For more information on AWS FedRAMPsm compliance please see the [AWS FedRAMPsm FAQs](#) at: <https://aws.amazon.com/compliance/fedramp/>

FERPA

[The Family Educational Rights and Privacy Act](#) (FERPA) (20 U.S.C. § 1232g; 34 CFR Part 99) is a Federal law that protects the privacy of student education records. The law applies to all schools that receive funds under an applicable program of the U.S. Department of Education. FERPA gives parents certain rights with respect to their children's education records. These rights transfer to the student when he or she reaches the age of 18, or attends a school beyond the high school level. Students to whom the rights have transferred are "eligible students."

AWS enables covered entities and their business associates subject to FERPA to leverage the secure AWS environment to process, maintain, and store protected education information.

AWS also offers a [FERPA-focused whitepaper](#) for customers interested in learning more about how they can leverage AWS for the processing and storage of educational data.

The "[FERPA Compliance on AWS Whitepaper](#)" outlines how companies can use AWS to process systems that facilitate FERPA compliance:
https://do.awsstatic.com/whitepapers/compliance/AWS_FERPA_Whitepaper.pdf

FIPS 140-2

[The Federal Information Processing Standard \(FIPS\) Publication 140-2](#) is a US government security standard that specifies the security requirements for cryptographic modules protecting sensitive information. To support customers with FIPS 140-2 requirements, SSL terminations in [AWS GovCloud \(US\)](#) operate using FIPS 140-2 validated hardware. AWS works with AWS GovCloud (US) customers to provide the information they need to help manage compliance when using the [AWS GovCloud \(US\) environment](#).

FISMA and DIACAP



AWS enables US government agencies to achieve and sustain compliance with the Federal Information Security Management Act ([FISMA](#)). The AWS infrastructure has been evaluated by independent assessors for a variety of government systems as part of their system owners' approval process. Numerous Federal Civilian and Department of Defense (DoD) organizations have successfully achieved security authorizations for systems hosted on AWS in accordance with the Risk Management Framework (RMF) process defined in NIST 800-37 and DoD Information Assurance Certification and Accreditation Process ([DIACAP](#)).

GxP

GxP is an acronym that refers to the regulations and guidelines applicable to life sciences organizations that make food and medical products such as drugs, medical devices, and medical software applications. The overall intent of GxP requirements is to ensure that food and medical products are safe for consumers and to ensure the integrity of data used to make product-related safety decisions.

AWS offers a [GxP whitepaper](#) which details a comprehensive approach for using AWS for GxP systems. This whitepaper provides guidance for using [AWS Products in the context of GxP](#) and the content has been developed in conjunction with AWS pharmaceutical and medical device customers, as well as software partners, who are currently using AWS Products in their validated GxP systems.

For more information on the GxP on AWS [please contact AWS Sales and Business Development](#).

For additional information please see our GxP Compliance FAQs:

<https://aws.amazon.com/compliance/gxp-part-11-annex-11/>

HIPAA

AWS enables covered entities and their business associates subject to the U.S. Health Insurance Portability and Accountability Act (HIPAA) to leverage the secure AWS environment to process, maintain, and store protected health information and AWS will be signing business associate agreements with such customers. AWS also offers a HIPAA-focused whitepaper for customers interested in learning more about how they can leverage AWS for the processing and storage of health information. The [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) whitepaper outlines how companies can use AWS to process systems that facilitate HIPAA and Health Information Technology for Economic and Clinical Health (HITECH) compliance.

Customers may use any AWS service in an account designated as a HIPAA account, but they should only process, store and transmit PHI in the HIPAA-eligible services defined in the BAA. There are nine HIPAA-eligible services today, including:

- [Amazon DynamoDB](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Cloud Compute \(EC2\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon Glacier](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(Amazon RDS\)](#) using only MySQL and Oracle engines
- [Amazon Simple Storage Service \(S3\)](#)



AWS follows a standards-based risk management program to ensure that the HIPAA-eligible services specifically support the security, control, and administrative processes required under HIPAA. Using these services to store and process PHI allows our customers and AWS to address the HIPAA requirements applicable to our utility-based operating model. AWS prioritizes and adds new eligible services based on customer demand.

For additional information please see our HIPAA Compliance FAQs:

<https://aws.amazon.com/compliance/hipaa-compliance/>

Architecting for HIPAA Security and Compliance on Amazon Web Services:

https://do.awsstatic.com/whitepapers/compliance/AWS_HIPAA_Compliance_Whitepaper.pdf

IRAP

The Information Security Registered Assessors Program (IRAP) enables Australian government customers to validate that appropriate controls are in place and determine the appropriate responsibility model for addressing the needs of the Australian Signals Directorate (ASD) Information Security Manual (ISM).

Amazon Web Services **has completed an independent assessment** that has determined all applicable ISM controls are in place relating to the processing, storage and transmission of Unclassified (DLM) for the AWS Sydney Region.

IRAP Compliance FAQs:

<https://aws.amazon.com/compliance/irap/>

For more information see: **Appendix B: AWS alignment with the Australian Signals Directorate (ASD) Cloud Computing Security Considerations**

ISO 9001

AWS has achieved ISO 9001 certification, AWS' ISO 9001 certification directly supports customers who develop, migrate and operate their quality-controlled IT systems in the AWS cloud. Customers can leverage AWS' compliance reports as evidence for their own ISO 9001 programs and industry-specific quality programs, such as GxP in life sciences, ISO 13485 in medical devices, AS9100 in aerospace, and ISO/TS 16949 in automotive. AWS customers who don't have quality system requirements will still benefit from the additional assurance and transparency that an ISO 9001 certification provides.

The ISO 9001 certification covers the quality management system over a specified scope of AWS services and Regions of operations (below) and services including:

- [AWS CloudFormation](#)
- [AWS Cloud Hardware Security Model \(HSM\)](#)
- [Amazon CloudFront](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [AWS Directory Service](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 VM Import/Export](#)
- [AWS Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Cloud Compute \(EC2\)](#)



- [Amazon EC2 Container Service \(ECS\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [AWS Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [AWS Storage Gateway](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [AWS WAF - Web Application Firewall](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- The underlying physical infrastructure and the AWS Management Environment

AWS' ISO 9001 accreditation covers AWS Regions including US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), South America (Sao Paulo), EU (Ireland), EU (Frankfurt) and Asia Pacific (Singapore), Asia Pacific (Sydney), and Asia Pacific (Tokyo).

ISO 9001:2008 is a global standard for managing the quality of products and services. The 9001 standard outlines a quality management system based on eight principles defined by the International Organization for Standardization (ISO) Technical Committee for Quality Management and Quality Assurance. They include:

- Customer focus
- Leadership
- Involvement of people
- Process approach
- System approach to management
- Continual Improvement
- Factual approach to decision-making
- Mutually beneficial supplier relationships

The AWS ISO 9001 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_9001_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 9001 certification at:

<https://aws.amazon.com/compliance/iso-9001-faqs/>

ISO 27001

AWS has achieved ISO 27001 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services including:



- [AWS CloudFormation](#)
- [Amazon CloudFront](#)
- [AWS Cloudtrail](#)
- [AWS Directory Service](#)
- [Amazon DynamoDB](#)
- [AWS Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Cloud Compute \(EC2\)](#)
- [Amazon EC2 Container Service \(ECS\)](#)
- [AWS Direct Connect](#)
- [Amazon EC2 VM Import/Export](#)
- [AWS Cloud Hardware Security Model \(HSM\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [AWS Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [AWS Storage Gateway](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [AWS WAF - Web Application Firewall](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- The underlying physical infrastructure (including GovCloud) and the AWS Management Environment

ISO 27001/27002 is a widely-adopted global security standard that sets out requirements and best practices for a systematic approach to managing company and customer information that's based on periodic risk assessments appropriate to ever-changing threat scenarios. In order to achieve the certification, a company must show it has a systematic and ongoing approach to managing information security risks that affect the confidentiality, integrity, and availability of company and customer information. This certification reinforces Amazon's commitment to providing significant information regarding our security controls and practices.

AWS' ISO 27001 accreditation covers AWS Regions including US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), South America (Sao Paulo), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Sydney), and Asia Pacific (Tokyo).

The AWS ISO 27001 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_27001_global_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 27001 certification at:

<https://aws.amazon.com/compliance/iso-27001-faqs/>



ISO 27017

ISO 27017 is the newest code of practice released by the International Organization for Standardization (ISO). It provides implementation guidance on information security controls that specifically relate to cloud services.

AWS has achieved ISO 27017 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services including:

- [Amazon CloudFront](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 Container Service \(ECS\)](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [Amazon SimpleDB](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [AWS Directory Service](#)
- [AWS Elastic Beanstalk](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Storage Gateway](#)
- [AWS WAF \(Web Application Firewall\)](#)
- [Elastic Load Balancing](#)
- [VM Import/Export](#)

The AWS ISO 27017 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_27017_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 27017 certification at:

<https://aws.amazon.com/compliance/iso-27017-faqs/>

ISO 27018



ISO 27018 is the first International code of practice that focuses on protection of personal data in the cloud. It is based on ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls applicable to public cloud Personally Identifiable Information (PII). It also provides a set of additional controls and associated guidance intended to address public cloud PII protection requirements not addressed by the existing ISO 27002 control set.

AWS has achieved ISO 27018 certification of our Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services including:

- [Amazon CloudFront](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 Container Service \(ECS\)](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon ElastiCache](#)
- [Amazon Glacier](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow Service \(SWF\)](#)
- [Amazon SimpleDB](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [Amazon WorkDocs](#)
- [Amazon WorkMail](#)
- [Amazon WorkSpaces](#)
- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [AWS Directory Service](#)
- [AWS Elastic Beanstalk](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Storage Gateway](#)
- [AWS WAF \(Web Application Firewall\)](#)
- [Elastic Load Balancing](#)
- [VM Import/Export](#)

The AWS ISO 27018 certification can be downloaded at:

https://do.awsstatic.com/certifications/iso_27018_certification.pdf

AWS provides additional information and frequently asked questions about its ISO 27018 certification at:

<https://aws.amazon.com/compliance/iso-27018-faqs/>

ITAR

The [AWS GovCloud \(US\)](#) region supports US International Traffic in Arms Regulations ([ITAR](#)) compliance. As a part of managing a comprehensive ITAR compliance program, companies subject to ITAR export regulations must control unintended exports by restricting access to protected data to US Persons and restricting physical location of that data to the US. AWS GovCloud (US) provides an environment physically located in the US and where access by AWS Personnel is limited to US Persons, thereby allowing qualified companies to transmit, process, and store protected articles and data subject to ITAR restrictions. The AWS GovCloud (US) environment has been audited by an independent third-party to validate the proper controls are in place to support customer export compliance programs for this requirement.

MPAA

The Motion Picture Association of America (MPAA) has established a set of best practices for securely storing, processing and delivering protected media and content (<http://www.fightfilmtheft.org/facility-security-program.html>). Media companies use these best practices as a way to assess risk and security of their content and infrastructure. AWS has demonstrated alignment with the MPAA best practices and the AWS infrastructure is compliant with all applicable MPAA infrastructure controls. While the MPAA does not offer a “certification,” media industry customers can use the AWS MPAA documentation to augment their risk assessment and evaluation of MPAA-type content on AWS.

See the [AWS Compliance MPAA hub page](#) for additional details:
<https://aws.amazon.com/compliance/mpaa/>

MTCS Tier 3 Certification

The [Multi-Tier Cloud Security \(MTCS\)](#) is an operational Singapore security management Standard (SPRING SS 584:2013), based on ISO 27001/02 Information Security Management System (ISMS) standards. The certification assessment requires us to:

- Systematically evaluate our information security risks, taking into account the impact of company threats and vulnerabilities
- Design and implement a comprehensive suite of information security controls and other forms of risk management to address company and architecture security risks
- Adopt an overarching management process to ensure that the information security controls meet the our information security needs on an ongoing basis

View the MTCS Hub Page at:
<https://aws.amazon.com/compliance/aws-multitiered-cloud-security-standard-certification/>

NIST

In June 2015 The National Institute of Standards and Technology (NIST) released guidelines [800-171](#), "Final Guidelines for Protecting Sensitive Government Information Held by Contractors". This guidance is applicable to the protection of Controlled Unclassified Information (CUI) on nonfederal systems.

AWS is already compliant with these guidelines, and customers can effectively comply with NIST 800-171 immediately. NIST 800-171 outlines a subset of the NIST 800-53 requirements, a guideline under which AWS has already been audited under the FedRAMP program. The FedRAMP Moderate security control baseline is more rigorous than the recommended requirements established in Chapter 3 of 800-171, and includes a significant number of security controls above and beyond those required of FISMA Moderate systems that



protect CUI data. A detailed mapping is available in the [NIST Special Publication 800-171](#), starting on page D2 (which is page 37 in the PDF).

PCI DSS Level 1

AWS is Level 1 compliant under the Payment Card Industry (PCI) Data Security Standard (DSS). Customers can run applications on our PCI-compliant technology infrastructure for storing, processing, and transmitting credit card information in the cloud. In February 2013, the PCI Security Standards Council released PCI DSS Cloud Computing Guidelines. These guidelines provide customers who are managing a cardholder data environment with considerations for maintaining PCI DSS controls in the cloud. AWS has incorporated the PCI DSS Cloud Computing Guidelines into the AWS PCI Compliance Package for customers. The AWS PCI Compliance Package includes the AWS PCI Attestation of Compliance (AoC), which shows that AWS has been successfully validated against standards applicable to a Level 1 service provider under PCI DSS Version 3.1, and the AWS PCI Responsibility Summary, which explains how compliance responsibilities are shared between AWS and our customers in the cloud.

The following services are in scope for PCI DSS Level 1:

- [Auto Scaling](#)
- [AWS CloudFormation](#)
- [Amazon CloudFront](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)
- [AWS Direct Connect](#)
- [Amazon DynamoDB](#)
- [AWS Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon Glacier](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Workflow Service SWF](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- The underlying physical infrastructure (including GovCloud) and the AWS Management Environment

The latest scope of services and regions for the AWS PCI DSS Level 1 certification can be found at: <https://aws.amazon.com/compliance/pci-dss-level-1-faqs/>

SOC 1/ISAE 3402

Amazon Web Services publishes a Service Organization Controls 1 (SOC 1), Type II report. The audit for this report is conducted in accordance with American Institute of Certified Public Accountants (AICPA): AT 801



(formerly SSAE 16) and the International Standards for Assurance Engagements No. 3402 (ISAE 3402). This dual-standard report is intended to meet a broad range of financial auditing requirements for U.S. and international auditing bodies. The SOC 1 report audit attests that AWS' control objectives are appropriately designed and that the individual controls defined to safeguard customer data are operating effectively. This report is the replacement of the Statement on Auditing Standards No. 70 (SAS 70) Type II Audit report.

The AWS SOC 1 control objectives are provided here. The report itself identifies the control activities that support each of these objectives and the independent auditor's results of their testing procedures of each control.

Objective Area	Objective Description
Security Organization	Controls provide reasonable assurance that information security policies have been implemented and communicated throughout the organization.
Employee User Access	Controls provide reasonable assurance that procedures have been established so that Amazon employee user accounts are added, modified and deleted in a timely manner and reviewed on a periodic basis.
Logical Security	Controls provide reasonable assurance that policies and mechanisms are in place to appropriately restrict unauthorized internal and external access to data and customer data is appropriately segregated from other customers.
Secure Data Handling	Controls provide reasonable assurance that data handling between the customer's point of initiation to an AWS storage location is secured and mapped accurately.
Physical Security and Environmental Protection	Controls provide reasonable assurance that physical access to data centers is restricted to authorized personnel and that mechanisms are in place to minimize the effect of a malfunction or physical disaster to data center facilities.
Change Management	Controls provide reasonable assurance that changes (including emergency / non-routine and configuration) to existing IT resources are logged, authorized, tested, approved and documented.
Data Integrity, Availability and Redundancy	Controls provide reasonable assurance that data integrity is maintained through all phases including transmission, storage and processing.
Incident Handling	Controls provide reasonable assurance that system incidents are recorded, analyzed, and resolved.

The SOC 1 reports are designed to focus on controls at a service organization that are likely to be relevant to an audit of a user entity's financial statements. As AWS' customer base is broad, and the use of AWS services is equally as broad, the applicability of controls to customer financial statements varies by customer. Therefore, the AWS SOC 1 report is designed to cover specific key controls likely to be required during a financial audit, as well as covering a broad range of IT general controls to accommodate a wide range of usage and audit scenarios. This allows customers to leverage the AWS infrastructure to store and process critical data, including that which is integral to the financial reporting process. AWS periodically reassesses the selection of these controls to consider customer feedback and usage of this important audit report.

AWS' commitment to the SOC 1 report is ongoing, and AWS will continue the process of periodic audits. The SOC 1 report scope covers:

- [AWS CloudFormation](#)
- [AWS CloudHSM](#)
- [AWS CloudTrail](#)



- [AWS Direct Connect](#)
- [Amazon DynamoDB](#)
- [Amazon EC2 VM Import/Export](#)
- [Amazon Elastic Beanstalk](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon ElastiCache](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Elastic Load Balancing \(ELB\)](#)
- [Amazon Elastic MapReduce \(EMR\)](#)
- [Amazon Glacier](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Route 53](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Email Service \(SES\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Simple Workflow \(SWF\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [AWS Storage Gateway](#)
- [Amazon Virtual Private Cloud \(VPC\)](#)
- [Amazon Workspaces](#)

SOC 2

In addition to the SOC 1 report, AWS publishes a Service Organization Controls 2 (SOC 2), Type II report. Similar to the SOC 1 in the evaluation of controls, the SOC 2 report is an attestation report that expands the evaluation of controls to the criteria set forth by the American Institute of Certified Public Accountants (AICPA) Trust Services Principles. These principles define leading practice controls relevant to security, availability, processing integrity, confidentiality, and privacy applicable to service organizations such as AWS. The AWS SOC 2 is an evaluation of the design and operating effectiveness of controls that meet the criteria for the security and availability principles set forth in the AICPA's Trust Services Principles criteria. This report provides additional transparency into AWS security and availability based on a pre-defined industry standard of leading practices and further demonstrates AWS' commitment to protecting customer data. The SOC 2 report scope covers the same services covered in the SOC 1 report. See the SOC 1 description above for the in-scope services.

SOC 3

AWS publishes a Service Organization Controls 3 (SOC 3) report. The SOC 3 report is a publically-available summary of the AWS SOC 2 report. The report includes the external auditor's opinion of the operation of controls (based on the [AICPA's Security Trust Principles](#) included in the SOC 2 report), the assertion from AWS management regarding the effectiveness of controls, and an overview of AWS Infrastructure and Services. The AWS SOC 3 report includes all AWS data centers worldwide that support in-scope services. This is a great resource for customers to validate that AWS has obtained external auditor assurance without going through the process to request a SOC 2 report. The SOC 3 report scope covers the same services covered in the SOC 1 report. See the SOC 1 description above for the in-scope services. [View the AWS SOC 3 report here.](#)



Key Compliance Questions and AWS

This section addresses generic cloud computing compliance questions specifically for AWS. These common compliance questions listed may be of interest when evaluating and operating in a cloud computing environment and may assist in AWS customers' control management efforts.

Ref	Cloud Computing Question	AWS Information
1	Control ownership. Who owns which controls for cloud-deployed infrastructure?	For the portion deployed into AWS, AWS controls the physical components of that technology. The customer owns and controls everything else, including control over connection points and transmissions. To help customers better understand what controls we have in place and how effectively they are operating, we publish a SOC 1 Type II report with controls defined around EC2, S3 and VPC, as well as detailed physical security and environmental controls. These controls are defined at a high level of specificity that should meet most customer needs. AWS customers that have signed a non-disclosure agreement with AWS may request a copy of the SOC 1 Type II report.
2	Auditing IT. How can auditing of the cloud provider be accomplished?	Auditing for most layers and controls above the physical controls remains the responsibility of the customer. The definition of AWS-defined logical and physical controls is documented in the SOC 1 Type II report, and the report is available for review by audit and compliance teams. AWS ISO 27001 and other certifications are also available for auditors to review.
3	Sarbanes-Oxley compliance. How is SOX compliance achieved if in-scope systems are deployed in the cloud provider environment?	If a customer processes financial information in the AWS cloud, the customer's auditors may determine that some AWS systems come into scope for Sarbanes-Oxley (SOX) requirements. The customer's auditors must make their own determination regarding SOX applicability. Because most of the logical access controls are managed by customer, the customer is best positioned to determine if its control activities meet relevant standards. If the SOX auditors request specifics regarding AWS' physical controls, they can reference the AWS SOC 1 Type II report which details the controls that AWS provides.
4	HIPAA compliance. Is it possible to meet HIPAA compliance requirements while deployed in the cloud provider environment?	HIPAA requirements apply to and are controlled by the AWS customer. The AWS platform allows for the deployment of solutions that meet industry-specific certification requirements such as HIPAA. Customers can use AWS services to maintain a security level that is equivalent or greater than those required to protect electronic health records. Customers have built healthcare applications compliant with HIPAA's Security and Privacy Rules on AWS. AWS provides additional information about HIPAA compliance on its web site, including a whitepaper on this topic .
5	GLBA compliance. Is it possible to meet GLBA certification requirements while deployed in the cloud provider environment?	Most GLBA requirements are controlled by the AWS customer. AWS provides means for customers to protect data, manage permissions, and build GLBA-compliant applications on AWS infrastructure. If the customer requires specific assurance that physical security controls are operating effectively, they can reference the AWS SOC 1 Type II report as relevant.

Ref	Cloud Computing Question	AWS Information
6	Federal regulation compliance. Is it possible for a US Government agency to be compliant with security and privacy regulations while deployed in the cloud provider environment?	US Federal agencies can be compliant under a number of compliance standards, including the Federal Information Security Management Act (FISMA) of 2002, Federal Risk and Authorization Management Program (FedRAMP), the Federal Information Processing Standard (FIPS) Publication 140-2, and the International Traffic in Arms Regulations (ITAR). Compliance with other laws and statutes may also be accommodated depending on the requirements set forth in the applicable legislation.
7	Data location. Where does customer data reside?	AWS customers designate in which physical region their data and their servers will be located. Data replication for S3 data objects is done within the regional cluster in which the data is stored and is not replicated to other data center clusters in other regions. AWS customers designate in which physical region their data and their servers will be located. AWS will not move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Frankfurt), EU (Ireland), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing), and South America (Sao Paulo).
8	E-Discovery. Does the cloud provider meet the customer's needs to meet electronic discovery procedures and requirements?	AWS provides infrastructure, and customers manage everything else, including the operating system, the network configuration, and the installed applications. Customers are responsible for responding appropriately to legal procedures involving the identification, collection, processing, analysis, and production of electronic documents they store or process using AWS. Upon request, AWS may work with customers who require AWS' assistance in legal proceedings.
9	Data center tours. Are data center tours by customers allowed by the cloud provider?	No. Due to the fact that our data centers host multiple customers, AWS does not allow data center tours by customers, as this exposes a wide range of customers to physical access of a third party. To meet this customer need, an independent and competent auditor validates the presence and operation of controls as part of our SOC 1 Type II report. This broadly accepted third-party validation provides customers with the independent perspective of the effectiveness of controls in place. AWS customers that have signed a non-disclosure agreement with AWS may request a copy of the SOC 1 Type II report. Independent reviews of data center physical security is also a part of the ISO 27001 audit, the PCI assessment, ITAR audit, and the FedRAMP sm testing programs.
10	Third-party access. Are third parties allowed access to the cloud provider data centers?	AWS strictly controls access to data centers, even for internal employees. Third parties are not provided access to AWS data centers except when explicitly approved by the appropriate AWS data center manager per the AWS access policy. See the SOC 1 Type II report for specific controls related to physical access, data center access authorization, and other related controls.

Ref	Cloud Computing Question	AWS Information
11	Privileged actions. Are privileged actions monitored and controlled?	Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored. In addition, customer data is and server instances are logically isolated from other customers by default. Privileged user access control is reviewed by an independent auditor during the AWS SOC 1, ISO 27001, PCI, ITAR, and FedRAMP sm audits.
12	Insider access. Does the cloud provider address the threat of inappropriate insider access to customer data and applications?	AWS provides specific SOC 1 controls to address the threat of inappropriate insider access, and the public certification and compliance initiatives covered in this document address insider access. All certifications and third-party attestations evaluate logical access preventative and detective controls. In addition, periodic risk assessments focus on how insider access is controlled and monitored.
13	Multi-tenancy. Is customer segregation implemented securely?	The AWS environment is a virtualized, multi-tenant environment. AWS has implemented security management processes, PCI controls, and other security controls designed to isolate each customer from other customers. AWS systems are designed to prevent customers from accessing physical hosts or instances not assigned to them by filtering through the virtualization software. This architecture has been validated by an independent PCI Qualified Security Assessor (QSA) and was found to be in compliance with all requirements of PCI DSS version 3.1 published in April 2015. Note that AWS also has single-tenancy options. Dedicated Instances are Amazon EC2 instances launched within your Amazon Virtual Private Cloud (Amazon VPC) that run hardware dedicated to a single customer. Dedicated Instances let you take full advantage of the benefits of Amazon VPC and the AWS cloud while isolating your Amazon EC2 compute instances at the hardware level.
14	Hypervisor vulnerabilities. Has the cloud provider addressed known hypervisor vulnerabilities?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. See the AWS security whitepaper for more information on the Xen hypervisor and instance isolation.
15	Vulnerability management. Are systems patched appropriately?	AWS is responsible for patching systems supporting the delivery of service to customers, such as the hypervisor and networking services. This is done as required per AWS policy and in accordance with ISO 27001, NIST, and PCI requirements. Customers control their own guest operating systems, software and applications and are therefore responsible for patching their own systems.
16	Encryption. Do the provided services support encryption?	Yes. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB, and EC2. IPsec tunnels to VPC are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Customers may also use third-party encryption technologies. Refer to the AWS Security white paper for more information.

Ref	Cloud Computing Question	AWS Information
17	Data ownership. What are the cloud provider's rights over customer data?	AWS customers retain control and ownership of their data. AWS errs on the side of protecting customer privacy and is vigilant in determining which law enforcement requests we must comply with. AWS does not hesitate to challenge orders from law enforcement if we think the orders lack a solid basis.
18	Data isolation. Does the cloud provider adequately isolate customer data?	All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Amazon S3 provides advanced data access controls. Please see the AWS security whitepaper for more information about specific data services' security.
19	Composite services. Does the cloud provider layer its service with other providers' cloud services?	AWS does not leverage any third-party cloud providers to deliver AWS services to customers.
20	Physical and environmental controls. Are these controls operated by the cloud provider specified?	Yes. These are specifically outlined in the SOC 1 Type II report. In addition, other certifications AWS supports such as ISO 27001 and FedRAMP sm require best practice physical and environmental controls.
21	Client-side protection. Does the cloud provider allow customers to secure and manage access from clients, such as PC and mobile devices?	Yes. AWS allows customers to manage client and mobile applications to their own requirements.
22	Server security. Does the cloud provider allow customers to secure their virtual servers?	Yes. AWS allows customers to implement their own security architecture. See the AWS security whitepaper for more details on server and network security.
23	Identity and Access Management. Does the service include IAM capabilities?	AWS has a suite of identity and access management offerings, allowing customers to manage user identities, assign security credentials, organize users in groups, and manage user permissions in a centralized way. Please see the AWS web site for more information.
24	Scheduled maintenance outages. Does the provider specify when systems will be brought down for maintenance?	AWS does not require systems to be brought offline to perform regular maintenance and system patching. AWS' own maintenance and system patching generally do not impact customers. Maintenance of instances themselves is controlled by the customer.
25	Capability to scale. Does the provider allow customers to scale beyond the original agreement?	The AWS cloud is distributed, highly secure and resilient, giving customers massive scale potential. Customers may scale up or down, paying for only what they use.
26	Service availability. Does the provider commit to a high level of availability?	AWS does commit to high levels of availability in its service level agreements (SLA). For example, Amazon EC2 commits to annual uptime percentage of at least 99.95% during the service year. Amazon S3 commits to monthly uptime percentage of at least 99.9%. Service credits are provided in the case these availability metrics are not met.

Ref	Cloud Computing Question	AWS Information
27	Distributed Denial Of Service (DDoS) attacks. How does the provider protect their service against DDoS attacks?	The AWS network provides significant protection against traditional network security issues and the customer can implement further protection. See the AWS Security Whitepaper for more information on this topic, including a discussion of DDoS attacks.
28	Data portability. Can the data stored with a service provider be exported by customer request?	AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport.
29	Service provider business continuity. Does the service provider operate a business continuity program?	AWS does operate a business continuity program. Detailed information is provided in the AWS Security Whitepaper.
30	Customer business continuity. Does the service provider allow customers to implement a business continuity plan?	AWS provides customers with the capability to implement a robust continuity plan, including the utilization of frequent server instance back-ups, data redundancy replication, and multi-region/availability zone deployment architectures.
31	Data durability. Does the service specify data durability?	Amazon S3 provides a highly durable storage infrastructure. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 Region. Once stored, Amazon S3 maintains the durability of objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. Data stored in S3 is designed to provide 99.99999999% durability and 99.99% availability of objects over a given year.
32	Backups. Does the service provide backups to tapes?	AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 service is designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy. For information on data durability and redundancy, please refer to the AWS web site.
33	Price increases. Will the service provider raise prices unexpectedly?	AWS has a history of frequently reducing prices as the cost to provide these services reduces over time. AWS has reduced prices consistently over the past several years.
34	Sustainability. Does the service provider company have long term sustainability potential?	AWS is a leading cloud provider and is a long-term business strategy of Amazon.com. AWS has very high long term sustainability potential.

AWS Contact

Customers can request the reports and certifications produced by our third-party auditors or can request more information about AWS Compliance by contacting [AWS Sales and Business Development](#). The representative will route customers to the proper team depending on nature of the inquiry. For additional information on AWS Compliance, see the [AWS Compliance site](#) or send questions directly to awscompliance@amazon.com.



Appendix A: CSA Consensus Assessments Initiative Questionnaire v3.0.1

The Cloud Security Alliance (CSA) is a “not-for-profit organization with a mission to promote the use of best practices for providing security assurance within Cloud Computing, and to provide education on the uses of Cloud Computing to help secure all other forms of computing.” [Reference

<https://cloudsecurityalliance.org/about/>] A wide range of industry security practitioners, corporations, and associations participate in this organization to achieve its mission.

The CSA Consensus Assessments Initiative Questionnaire provides a set of questions the CSA anticipates a cloud consumer and/or a cloud auditor would ask of a cloud provider. It provides a series of security, control, and process questions which can then be used for a wide range of uses, including cloud provider selection and security evaluation. AWS has completed this questionnaire with the answers below.

Control Group	CID	Consensus Assessment Questions	AWS Response
Application & Interface Security <i>Application Security</i>	AIS-01.1	Do you use industry standards (Build Security in Maturity Model [BSIMM] benchmarks, Open Group ACS Trusted Technology Provider Framework, NIST, etc.) to build in security for your Systems/Software Development Lifecycle (SDLC)?	The AWS system development lifecycle incorporates industry best practices which include formal design reviews by the AWS Security Team, threat modeling and completion of a risk assessment. Refer to the AWS Overview of Security Processes for further details. AWS has in place procedures to manage new development of resources. Refer to ISO 27001 standard, Annex A, domain 14 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	AIS-01.2	Do you use an automated source code analysis tool to detect security defects in code prior to production?	
	AIS-01.3	Do you use manual source-code analysis to detect security defects in code prior to production?	
	AIS-01.4	Do you verify that all of your software suppliers adhere to industry standards for Systems/Software Development Lifecycle (SDLC) security?	
	AIS-01.5	(SaaS only) Do you review your applications for security vulnerabilities and address any issues prior to deployment to production?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Application & Interface Security <i>Customer Access Requirements</i>	AIS-02.1	Are all identified security, contractual and regulatory requirements for customer access contractually addressed and remediated prior to granting customers access to data, assets and information systems?	AWS Customers retain responsibility to ensure their usage of AWS is in compliance with applicable laws and regulations. AWS communicates its security and control environment to customers through industry certifications and third-party attestations, white papers (available at http://aws.amazon.com/compliance) and providing certifications, reports and other relevant documentation directly to AWS Customers.
	AIS-02.2	Are all requirements and trust levels for customers' access defined and documented?	
Application & Interface Security <i>Data Integrity</i>	AIS-03.1	Are data input and output integrity routines (i.e., reconciliation and edit checks) implemented for application interfaces and databases to prevent manual or systematic processing errors or corruption of data?	AWS data integrity controls as described in AWS SOC reports illustrates the data integrity controls maintained through all phases including transmission, storage and processing. In addition, refer to ISO 27001 standard, Annex A, domain 14 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Application & Interface Security <i>Data Security / Integrity</i>	AIS-04.1	Is your Data Security Architecture designed using an industry standard (e.g., CDSA, MULITSAFE, CSA Trusted Cloud Architectural Standard, FedRAMP, CAESARS)?	AWS Data Security Architecture was designed to incorporate industry leading practices. Refer to AWS Certifications, reports and whitepapers for additional details on the various leading practices that AWS adheres to (available at http://aws.amazon.com/compliance).
Audit Assurance & Compliance <i>Audit Planning</i>	AAC-01.1	Do you produce audit assertions using a structured, industry accepted format (e.g., CloudAudit/A6 URI Ontology, CloudTrust, SCAP/CYBEX, GRC XML, ISACA's Cloud Computing Management Audit/Assurance Program, etc.)?	AWS obtains certain industry certifications and independent third-party attestations and provides certain certifications, reports and other relevant documentation directly to AWS Customers.
Audit Assurance & Compliance <i>Independent Audits</i>	AAC-02.1	Do you allow tenants to view your SOC2/ISO 27001 or similar third-party audit or certification reports?	AWS provides third-party attestations, certifications, Service Organization Controls (SOC) reports and other relevant compliance reports directly to our customers under NDA. The AWS ISO 27001 certification can be downloaded here: http://do.awsstatic.com/certifications/iso_27001_global_certification.pdf. The AWS SOC 3 report can be downloaded here: https://do.awsstatic.com/whitepapers/compliance/soc3_amazon_web_services.pdf. AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat
	AAC-02.2	Do you conduct network penetration tests of your cloud service infrastructure regularly as prescribed by industry best practices and guidance?	
	AAC-02.3	Do you conduct application penetration tests of your cloud infrastructure regularly as prescribed by industry best practices and guidance?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	AAC - 02.4	Do you conduct internal audits regularly as prescribed by industry best practices and guidance?	assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership.
	AAC - 02.5	Do you conduct external audits regularly as prescribed by industry best practices and guidance?	In addition, the AWS control environment is subject to regular internal and external audits and risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment.
	AAC - 02.6	Are the results of the penetration tests available to tenants at their request?	
	AAC - 02.7	Are the results of internal and external audits available to tenants at their request?	
	AAC - 02.8	Do you have an internal audit program that allows for cross-functional audit of assessments?	
Audit Assurance & Compliance <i>Information System Regulatory Mapping</i>	AAC -03.1	Do you have the ability to logically segment or encrypt customer data such that data may be produced for a single tenant only, without inadvertently accessing another tenant's data?	All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers retain control and ownership of their data, thus it is their responsibility to choose to encrypt the data. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security
	AAC - 03.2	Do you have capability to recover data for a specific customer in the case of a failure or data loss?	
	AAC - 03.3	Do you have the capability to restrict the storage of customer data to specific countries or geographic locations?	AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 and Glacier services are designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy. For information on data durability and redundancy, please refer to the AWS website. AWS Customers designate in which physical region their content will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US)(Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region, and South America (Sao Paulo).

Control Group	CID	Consensus Assessment Questions	AWS Response
	AAC - 03.4	Do you have a program in place that includes the ability to monitor changes to the regulatory requirements in relevant jurisdictions, adjust your security program for changes to legal requirements, and ensure compliance with relevant regulatory requirements?	AWS monitors relevant legal and regulatory requirements. Refer to ISO 27001 standard Annex 18 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Business Continuity Management & Operational Resilience <i>Business Continuity Planning</i>	BCR -01.1	Do you provide tenants with geographically resilient hosting options?	Data centers are built in clusters in various global regions. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Refer to AWS Overview of Cloud Security whitepaper for additional details - available at http://aws.amazon.com/security .
	BCR -01.2	Do you provide tenants with infrastructure service failover capability to other providers?	
Business Continuity Management & Operational Resilience <i>Business Continuity Testing</i>	BCR -02.1	Are business continuity plans subject to test at planned intervals or upon significant organizational or environmental changes to ensure continuing effectiveness?	AWS Business Continuity Policies and Plans have been developed and tested in alignment with ISO 27001 standards. Refer to ISO 27001 standard, annex A domain 17 for further details on AWS and business continuity.
Business Continuity Management & Operational Resilience <i>Power / Telecommunications</i>	BCR -03.1	Do you provide tenants with documentation showing the transport route of their data between your systems?	AWS Customers designate in which physical region their data and servers will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. AWS SOC reports provides additional details. Customers can also choose their network path to AWS facilities, including over dedicated, private networks where the customer controls the traffic routing.
	BCR - 03.2	Can tenants define how their data is transported and through which legal jurisdictions?	
Business Continuity Management & Operational Resilience Documentation	BCR -04.1	Are information system documents (e.g., administrator and user guides, architecture diagrams, etc.) made available to authorized personnel to ensure configuration, installation and operation of the information system?	Information System Documentation is made available internally to AWS personnel through the use of Amazon's Intranet site. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security/ . Refer to ISO 27001 Appendix A Domain 12.
Business Continuity Management & Operational Resilience <i>Environmental Risks</i>	BCR -05.1	Is physical protection against damage (e.g., natural causes, natural disasters, deliberate attacks) anticipated and designed with countermeasures applied?	AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 11.

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Equipment Location</i>	BCR -06.1	Are any of your data centers located in places that have a high probability/occurrence of high-impact environmental risks (floods, tornadoes, earthquakes, hurricanes, etc.)?	AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 11.
Business Continuity Management & Operational Resilience <i>Equipment Maintenance</i>	BCR -07.1	If using virtual infrastructure, does your cloud solution include independent hardware restore and recovery capabilities?	EBS Snapshot functionality allows customers to capture and restore virtual machine images at any time. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	BCR -07.2	If using virtual infrastructure, do you provide tenants with a capability to restore a Virtual Machine to a previous state in time?	
	BCR -07.3	If using virtual infrastructure, do you allow virtual machine images to be downloaded and ported to a new cloud provider?	
	BCR -07.4	If using virtual infrastructure, are machine images made available to the customer in a way that would allow the customer to replicate those images in their own off-site storage location?	
	BCR -07.5	Does your cloud solution include software/provider independent restore and recovery capabilities?	
Business Continuity Management & Operational Resilience <i>Equipment Power Failures</i>	BCR -08.1	Are security mechanisms and redundancies implemented to protect equipment from utility service outages (e.g., power failures, network disruptions, etc.)?	AWS equipment is protected from utility service outages in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS SOC reports provides additional details on controls in place to minimize the effect of a malfunction or physical disaster to the computer and data center facilities. In addition, refer to the AWS Cloud Security Whitepaper - available at http://aws.amazon.com/security.

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Impact Analysis</i>	BCR-09.1	Do you provide tenants with ongoing visibility and reporting of your operational Service Level Agreement (SLA) performance?	AWS CloudWatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com .
	BCR-09.2	Do you make standards-based information security metrics (CSA, CAMM, etc.) available to your tenants?	
	BCR-09.3	Do you provide customers with ongoing visibility and reporting of your SLA performance?	
Business Continuity Management & Operational Resilience <i>Policy</i>	BCR-10.1	Are policies and procedures established and made available for all personnel to adequately support services operations' roles?	Policies and Procedures have been established through AWS Security framework based upon NIST 800-53, ISO 27001, ISO 27017, ISO 27018, ISO 9001 standard and the PCI DSS requirements. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/compliance.
Business Continuity Management & Operational Resilience <i>Retention Policy</i>	BCR-11.1	Do you have technical control capabilities to enforce tenant data retention policies?	AWS provide customers with the ability to delete their data. However, AWS Customers retain control and ownership of their data so it is the customer's responsibility to manage data retention to their own requirements. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security. AWS errs on the side of protecting customer privacy and is vigilant in determining which law enforcement requests we must comply with. AWS does not hesitate to challenge orders from law enforcement if we think the orders lack a solid basis. For additional information refer to https://aws.amazon.com/compliance/data-privacy-faq/. AWS backup and redundancy mechanisms have been developed and tested in alignment with ISO 27001 standards. Refer to ISO 27001 standard, annex A domain 12 and the AWS SOC 2 report for additional information on AWS backup and redundancy mechanisms.
	BCR-11.2	Do you have a documented procedure for responding to requests for tenant data from governments or third parties?	
	BCR-11.4	Have you implemented backup or redundancy mechanisms to ensure compliance with regulatory, statutory, contractual or business requirements?	
	BCR-11.5	Do you test your backup or redundancy mechanisms at least annually?	
Change Control & Configuration Management <i>New Development / Acquisition</i>	CCC-01.1	Are policies and procedures established for management authorization for development or acquisition of new applications, systems, databases, infrastructure, services, operations and facilities?	Policies and Procedures have been established through AWS Security framework based upon NIST 800-53, ISO 27001, ISO 27017, ISO 27018, ISO 9001 standard and the PCI DSS requirements. Whether a customer is new to AWS or an advanced user, useful information about the services, ranging from introductions to advanced features, can be found on the AWS Documentation section of our website at https://aws.amazon.com/documentation/.
	CCC-01.2	Is documentation available that describes the installation, configuration and use of products/services/features?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Change Control & Configuration Management <i>Outsourced Development</i>	CCC-02.1	Do you have controls in place to ensure that standards of quality are being met for all software development?	AWS does not generally outsource development of software. AWS incorporates standards of quality as part of the system development lifecycle (SDLC) processes.
	CCC-02.2	Do you have controls in place to detect source code security defects for any outsourced software development activities?	Refer to ISO 27001 standard, Annex A, domain 14 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Quality Testing</i>	CCC-03.1	Do you provide your tenants with documentation that describes your quality assurance process?	AWS maintains an ISO 9001 certification. This is an independent validation of AWS quality system and determined that AWS activities comply with ISO 9001 requirements. AWS Security Bulletins notify customers of security and privacy events. Customers can subscribe to the AWS Security Bulletin RSS feed on our website. Refer to aws.amazon.com/security/security-bulletins/ .
	CCC-03.2	Is documentation describing known issues with certain products/services available?	AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com .
	CCC-03.3	Are there policies and procedures in place to triage and remedy reported bugs and security vulnerabilities for product and service offerings?	The AWS system development lifecycle (SDLC) incorporates industry best practices which include formal design reviews by the AWS Security Team, threat modeling and completion of a risk assessment. Refer to the AWS Overview of Security Processes for further details.
	CCC-03.4	Are mechanisms in place to ensure that all debugging and test code elements are removed from released software versions?	In addition, refer to ISO 27001 standard, Annex A, domain 14 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Unauthorized Software Installations</i>	CCC-04.1	Do you have controls in place to restrict and monitor the installation of unauthorized software onto your systems?	AWS' program, processes and procedures for managing malicious software is in alignment with ISO 27001 standards. Refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Production Changes</i>	CCC-05.1	Do you provide tenants with documentation that describes your production change management procedures and their roles / rights / responsibilities within it?	AWS SOC reports provides an overview of the controls in place to manage change management in the AWS environment. In addition, refer to ISO 27001 standard, Annex A, domain 12 for further details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Data Security & Information Lifecycle Management <i>Classification</i>	DSI-01.1	Do you provide a capability to identify virtual machines via policy tags/metadata (e.g., tags can be used to limit guest operating systems from booting/instantiating/transporting data in the wrong country)?	Virtual Machines are assigned to customers as a part of the EC2 service. Customers retain control over what resources are being used and where resources reside. Refer to the AWS website for additional details - http://aws.amazon.com .
	DSI-01.2	Do you provide a capability to identify hardware via policy tags/metadata/hardware tags (e.g., TXT/TPM, VN-Tag, etc.)?	AWS provides the ability to tag EC2 resources. A form of metadata, EC2 tags can be used to create user-friendly names, enhance searchability, and improve coordination between multiple users. The AWS Management Console has also supports tagging.
	DSI-01.3	Do you have a capability to use system geographic location as an authentication factor?	AWS provides the capability of conditional user access based on IP address. Customers can add conditions to control how users can use AWS, such as time of day, their originating IP address, or whether they are using SSL.
	DSI-01.4	Can you provide the physical location/geography of storage of a tenant's data upon request?	AWS provides customers the flexibility to place instances and store data within multiple geographic Regions. AWS Customers designate in which physical region their data and their servers will be located. AWS will not move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region and South America (Sao Paulo).
	DSI-01.5	Can you provide the physical location/geography of storage of a tenant's data in advance?	
	DSI-01.6	Do you follow a structured data-labeling standard (e.g., ISO 15489, Oasis XML Catalog Specification, CSA data type guidance)?	AWS Customers retain control and ownership of their data and may implement a structured data-labeling standard to meet their requirements.
	DSI-01.7	Do you allow tenants to define acceptable geographical locations for data routing or resource instantiation?	AWS provides customers the flexibility to place instances and store data within multiple geographic regions. AWS Customers designate in which physical region their data and their servers will be located. AWS will not move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region and South America (Sao Paulo).
Data Security & Information Lifecycle Management <i>Data Inventory / Flows</i>	DSI-02.1	Do you inventory, document, and maintain data flows for data that is resident (permanent or temporary) within the services' applications and infrastructure network and systems?	AWS Customers designate in which physical region their content will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US)(Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region, and South America (Sao Paulo).

Control Group	CID	Consensus Assessment Questions	AWS Response
	DSI-02.2	Can you ensure that data does not migrate beyond a defined geographical residency?	
Data Security & Information Lifecycle Management <i>eCommerce Transactions</i>	DSI-03.1	Do you provide open encryption methodologies (3.4ES, AES, etc.) to tenants in order for them to protect their data if it is required to move through public networks (e.g., the Internet)?	All of the AWS APIs are available via SSH-protected endpoints which provide server authentication. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Customers may also use third-party encryption technologies.
	DSI-03.2	Do you utilize open encryption methodologies any time your infrastructure components need to communicate with each other via public networks (e.g., Internet-based replication of data from one environment to another)?	Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Data Security & Information Lifecycle Management <i>Handling / Labeling / Security Policy</i>	DSI-04.1	Are policies and procedures established for labeling, handling and the security of data and objects that contain data?	AWS Customers retain control and ownership of their data and may implement a labeling and handling policy and procedures to meet their requirements.
	DSI-04.2	Are mechanisms for label inheritance implemented for objects that act as aggregate containers for data?	
Data Security & Information Lifecycle Management <i>Nonproduction Data</i>	DSI-05.1	Do you have procedures in place to ensure production data shall not be replicated or used in non-production environments?	AWS Customers retain control and ownership of their own data. AWS provides customers the ability to maintain and develop production and non-production environments. It is the responsibility of the customer to ensure that their production data is not replicated to non-production environments.
Data Security & Information Lifecycle Management <i>Ownership / Stewardship</i>	DSI-06.1	Are the responsibilities regarding data stewardship defined, assigned, documented and communicated?	AWS Customers retain control and ownership of their own data. Refer to the AWS Customer Agreement for additional information.
Data Security & Information Lifecycle Management <i>Secure Disposal</i>	DSI-07.1	Do you support secure deletion (e.g., degaussing/cryptographic wiping) of archived and backed-up data as determined by the tenant?	When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be

Control Group	CID	Consensus Assessment Questions	AWS Response
	DSI-07.2	Can you provide a published procedure for exiting the service arrangement, including assurance to sanitize all computing resources of tenant data once a customer has exited your environment or has vacated a resource?	decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security. Amazon EBS volumes are presented to you as raw unformatted block devices that have been wiped prior to being made available for use. Wiping occurs immediately before reuse so that you can be assured that the wipe process completed. If you have procedures requiring that all data be wiped via a specific method, such as those detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization"), you have the ability to do so on Amazon EBS. You should conduct a specialized wipe procedure prior to deleting the volume for compliance with your established requirements. Encryption of sensitive data is generally a good security practice, and AWS provides the ability to encrypt EBS volumes and their snapshots with AES-256. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. In order to be able to do this efficiently and with low latency, the EBS encryption feature is only available on EC2's more powerful instance types (e.g., M3, C3, R3, G2).
Datacenter Security <i>Asset Management</i>	DCS-01.1	Do you maintain a complete inventory of all of your critical assets that includes ownership of the asset?	In alignment with ISO 27001 standards, AWS Hardware assets are assigned an owner, tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 8 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	DCS-01.2	Do you maintain a complete inventory of all of your critical supplier relationships?	
Datacenter Security <i>Controlled Access Points</i>	DCS-02.1	Are physical security perimeters (e.g., fences, walls, barriers, guards, gates, electronic surveillance, physical authentication mechanisms, reception desks and security patrols) implemented?	Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. The AWS SOC reports provides additional details on the specific control activities executed by AWS. Refer to ISO 27001 standards; Annex A, domain 11 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Datacenter Security <i>Equipment Identification</i>	DCS-03.1	Is automated equipment identification used as a method to validate connection authentication integrity based on known equipment location?	AWS manages equipment identification in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Datacenter Security <i>Offsite Authorization</i>	DCS -04.1	Do you provide tenants with documentation that describes scenarios in which data may be moved from one physical location to another? (e.g., offsite backups, business continuity failovers, replication)	AWS Customers can designate which physical region their data will be located. AWS will not move customers' content from the selected Regions without notifying the customer unless required to comply with the law or requests of governmental entities. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Datacenter Security <i>Offsite equipment</i>	DCS -05.1	Can you provide tenants with evidence documenting your policies and procedures governing asset management and repurposing of equipment?	In alignment with ISO 27001 standards, when a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to ISO 27001 standards; Annex A, domain 8 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Datacenter Security <i>Policy</i>	DCS -06.1	Can you provide evidence that policies, standards and procedures have been established for maintaining a safe and secure working environment in offices, rooms, facilities and secure areas?	AWS engages with external certifying bodies and independent auditors to review and validate our compliance with compliance frameworks. AWS SOC reports provides additional details on the specific physical security control activities executed by AWS. Refer to ISO 27001 standards; Annex A, domain 11 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	DCS -06.2	Can you provide evidence that your personnel and involved third parties have been trained regarding your documented policies, standards and procedures?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security . AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification. In addition AWS SOC 1 and SOC 2 reports provides further information.
Datacenter Security <i>Secure Area Authorization</i>	DCS -07.1	Do you allow tenants to specify which of your geographic locations their data is allowed to move into/out of (to address legal jurisdictional considerations based on where data is stored vs. accessed)?	AWS Customers designate which physical region their data will be located. AWS will not move customers' content from the selected Regions without notifying the customer unless required to comply with the law or requests of governmental entities. As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US)(Oregon), EU (Ireland), EU (Frankfurt), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), China (Beijing) Region, and South America (Sao Paulo).

Control Group	CID	Consensus Assessment Questions	AWS Response
Datacenter Security <i>Unauthorized Persons Entry</i>	DCS -08.1	Are ingress and egress points, such as service areas and other points where unauthorized personnel may enter the premises, monitored, controlled and isolated from data storage and process?	Physical access is strictly controlled both at the perimeter and at building ingress points and includes, but is not limited to, professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. Physical access points to server locations are recorded by closed circuit television camera (CCTV) as defined in the AWS Data Center Physical Security Policy.
Datacenter Security <i>User Access</i>	DCS -09.1	Do you restrict physical access to information assets and functions by users and support personnel?	AWS Physical Security Mechanisms are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Encryption & Key Management <i>Entitlement</i>	EKM -01.1	Do you have key management policies binding keys to identifiable owners?	AWS provides customers the ability to use their own encryption mechanism for nearly all services including S3, EBS and EC2. VPC sessions are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications. AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Encryption & Key Management <i>Key Generation</i>	EKM -02.1	Do you have a capability to allow creation of unique encryption keys per tenant?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/).
	EKM -02.2	Do you have a capability to manage encryption keys on behalf of tenants?	Refer to AWS SOC reports for more details on KMS.
	EKM -02.3	Do you maintain key management procedures?	In addition, refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	EKM -02.4	Do you have documented ownership for each stage of the lifecycle of encryption keys?	Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications.
	EKM -02.5	Do you utilize any third party/open source/proprietary frameworks to manage encryption keys?	AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Encryption & Key	EKM -03.1	Do you encrypt tenant data at rest (on disk/storage) within your environment?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key

Control Group	CID	Consensus Assessment Questions	AWS Response
Management Encryption	EKM - 03.2	Do you leverage encryption to protect data and virtual machine images during transport across and between networks and hypervisor instances?	Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS SOC reports for more details on KMS. In addition, refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
	EKM - 03.3	Do you support tenant-generated encryption keys or permit tenants to encrypt data to an identity without access to a public key certificate (e.g. identity-based encryption)?	
	EKM - 03.4	Do you have documentation establishing and defining your encryption management policies, procedures and guidelines?	
Encryption & Key Management Storage and Access	EKM - 04.1	Do you have platform and data appropriate encryption that uses open/validated formats and standard algorithms?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS SOC reports for more details on KMS.
	EKM - 04.2	Are your encryption keys maintained by the cloud consumer or a trusted key management provider?	AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications.
	EKM - 04.3	Do you store encryption keys in the cloud?	
	EKM - 04.4	Do you have separate key management and key usage duties?	AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Governance and Risk Management Baseline Requirements	GR M- 01.1	Do you have documented information security baselines for every component of your infrastructure (e.g., hypervisors, operating systems, routers, DNS servers, etc.)?	In alignment with ISO 27001 standards, AWS maintains system baselines for critical components. Refer to ISO 27001 standards, Annex A, domain 14 and 18 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. Customers can provide their own virtual machine image. VM Import enables customers to easily import virtual machine images from your existing environment to Amazon EC2 instances.
	GR M- 01.2	Do you have a capability to continuously monitor and report the compliance of your infrastructure against your information security baselines?	
	GR M- 01.3	Do you allow your clients to provide their own trusted virtual machine image to ensure conformance to their own internal standards?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Governance and Risk Management <i>Risk Assessments</i>	GR M-02.1	Do you provide security control health data in order to allow tenants to implement industry standard Continuous Monitoring (which allows continual tenant validation of your physical and logical control status)?	AWS does publish independent auditor reports and certifications to provide customers with considerable information regarding the policies, processes, and controls established and operated by AWS. The relevant certifications and reports can be provided to AWS Customers. Continuous Monitoring of logical controls can be executed by customers on their own systems.
	GR M-02.2	Do you conduct risk assessments associated with data governance requirements at least once a year?	In alignment with ISO 27001 standard, AWS maintains a Risk Management program to mitigate and manage risk. In addition AWS maintains an AWS ISO 27018 certification. Alignment with ISO 27018 demonstrates to customers that AWS has a system of controls in place that specifically address the privacy protection of their content. For more information refer to the AWS Compliance ISO 27018 FAQ http://aws.amazon.com/compliance/iso-27018-faqs/ .
Governance and Risk Management <i>Management Oversight</i>	GR M-03.1	Are your technical, business, and executive managers responsible for maintaining awareness of and compliance with security policies, procedures, and standards for both themselves and their employees as they pertain to the manager and employees' area of responsibility?	The Control environment at Amazon begins at the highest level of the Company. Executive and senior leadership play important roles in establishing the Company's tone and core values. Every employee is provided with the Company's Code of Business Conduct and Ethics and completes periodic training. Compliance audits are performed so that employees understand and follow the established policies. Refer to AWS Risk & Compliance whitepaper for additional details - available at http://aws.amazon.com/compliance .
Governance and Risk Management <i>Management Program</i>	GR M-04.1	Do you provide tenants with documentation describing your Information Security Management Program (ISMP)?	AWS provides our customers with our ISO 27001 certification. The ISO 27001 certification is specifically focused on the AWS ISMS and measures how AWS internal processes follow the ISO standard. Certification means a third party accredited independent auditor has performed an assessment of our processes and controls and confirms they are operating in alignment with the ISO 27001 certification standard. For additional information refer to the AWS Compliance ISO 27001 FAQ website: http://aws.amazon.com/compliance/iso-27001-faqs/ .
	GR M-04.2	Do you review your Information Security Management Program (ISMP) least once a year?	
Governance and Risk Management <i>Management Support / Involvement</i>	GR M-05.1	Do you ensure your providers adhere to your information security and privacy policies?	AWS has established information security framework and policies which have integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, PCI DSS v3.1 and National Institute of Standards and Technology (NIST) Publication 800-53 (Recommended Security Controls for Federal Information Systems).
Governance and Risk Management <i>Policy</i>	GR M-06.1	Do your information security and privacy policies align with industry standards (ISO-27001, ISO-22307, CoBIT, etc.)?	AWS manages third-party relationships in alignment with ISO 27001 standards. AWS Third Party requirements are reviewed by independent external

Control Group	CID	Consensus Assessment Questions	AWS Response
	GR M-06.2	Do you have agreements to ensure your providers adhere to your information security and privacy policies?	auditors during audits for our PCI DSS, ISO 27001 and FedRAMP compliance. Information about the AWS Compliance programs is published publicly on our website at http://aws.amazon.com/compliance/.
	GR M-06.3	Can you provide evidence of due diligence mapping of your controls, architecture and processes to regulations and/or standards?	
	GR M-06.4	Do you disclose which controls, standards, certifications and/or regulations you comply with?	
Governance and Risk Management <i>Policy Enforcement</i>	GR M-07.1	Is a formal disciplinary or sanction policy established for employees who have violated security policies and procedures?	AWS provides security policies and security training to employees to educate them as to their role and responsibilities concerning information security. Employees who violate Amazon standards or protocols are investigated and appropriate disciplinary action (e.g. warning, performance plan, suspension, and/or termination) is followed.
	GR M-07.2	Are employees made aware of what actions could be taken in the event of a violation via their policies and procedures?	Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security . Refer to ISO 27001 Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Governance and Risk Management <i>Business / Policy Change Impacts</i>	GR M-08.1	Do risk assessment results include updates to security policies, procedures, standards and controls to ensure they remain relevant and effective?	Updates to AWS security policies, procedures, standards and controls occur on an annual basis in alignment with the ISO 27001 standard. Refer to ISO 27001 for additional information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification.
Governance and Risk Management <i>Policy Reviews</i>	GR M-09.1	Do you notify your tenants when you make material changes to your information security and/or privacy policies?	Our AWS Cloud Security Whitepaper and Risk and Compliance whitepapers, available at http://aws.amazon.com/security and http://aws.amazon.com/compliance , are updated on a regular basis to reflect updates to the AWS policies.
	GR M-09.2	Do you perform, at minimum, annual reviews to your privacy and security policies?	The AWS SOC reports provide details related to privacy and security policy review.
Governance and Risk Management <i>Assessments</i>	GR M-10.1	Are formal risk assessments aligned with the enterprise-wide framework and performed at least annually, or at planned intervals, determining the likelihood and impact of all identified risks, using qualitative and quantitative methods?	In alignment with ISO 27001 AWS has developed a Risk Management program to mitigate and manage risk. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification. Refer to AWS Risk and Compliance Whitepaper (available at aws.amazon.com/security) for additional details on AWS Risk Management Framework.

Control Group	CID	Consensus Assessment Questions	AWS Response
	GR M-10.2	Is the likelihood and impact associated with inherent and residual risk determined independently, considering all risk categories (e.g., audit results, threat and vulnerability analysis, and regulatory compliance)?	
Governance and Risk Management Program	GR M-11.1	Do you have a documented, organization-wide program in place to manage risk?	In alignment with ISO 27001, AWS maintains a Risk Management program to mitigate and manage risk.
	GR M-11.2	Do you make available documentation of your organization-wide risk management program?	AWS management has a strategic business plan which includes risk identification and the implementation of controls to mitigate or manage risks. AWS management re-evaluates the strategic business plan at least biannually. This process requires management to identify risks within its areas of responsibility and to implement appropriate measures designed to address those risks. AWS Risk Management program is reviewed by independent external auditors during audits for our PCI DSS, ISO 27001 and FedRAMP compliance.
Human Resources Asset Returns	HRS -01.1	Are systems in place to monitor for privacy breaches and notify tenants expeditiously if a privacy event may have impacted their data?	AWS Customers retain the responsibility to monitor their own environment for privacy breaches. The AWS SOC reports provides an overview of the controls in place to monitor AWS managed environment.
	HRS -01.2	Is your Privacy Policy aligned with industry standards?	
Human Resources Background Screening	HRS -02.1	Pursuant to local laws, regulations, ethics and contractual constraints, are all employment candidates, contractors and involved third parties subject to background verification?	AWS conducts criminal background checks, as permitted by applicable law, as part of pre-employment screening practices for employees commensurate with the employee's position and level of access to AWS facilities. The AWS SOC reports provides additional details regarding the controls in place for background verification.
Human Resources Employment Agreements	HRS -03.1	Do you specifically train your employees regarding their specific role and the information security controls they must fulfill?	In alignment with ISO 27001 standard, all AWS employees complete periodic role based training that includes AWS Security training and requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. Refer to SOC reports for additional details.
	HRS -03.2	Do you document employee acknowledgment of training they have completed?	All personnel supporting AWS systems and devices must sign a non-disclosure agreement prior to being granted access. Additionally, upon hire, personnel are required to read and accept the Acceptable Use Policy and the Amazon Code of Business Conduct and Ethics (Code of Conduct) Policy.
	HRS -03.3	Are all personnel required to sign NDA or Confidentiality Agreements as a condition of employment to protect customer/tenant information?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	HRS - 03.4	Is successful and timed completion of the training program considered a prerequisite for acquiring and maintaining access to sensitive systems?	
	HRS - 03.5	Are personnel trained and provided with awareness programs at least once a year?	
Human Resources <i>Employment Termination</i>	HRS -04.1	Are documented policies, procedures and guidelines in place to govern change in employment and/or termination?	AWS Human Resources team defines internal management responsibilities to be followed for termination and role change of employees and vendors. AWS SOC reports provide additional details.
	HRS - 04.2	Do the above procedures and guidelines account for timely revocation of access and return of assets?	Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked. AWS SOC reports provide further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Human Resources <i>Portable / Mobile Devices</i>	HRS -05.1	Are policies and procedures established and measures implemented to strictly limit access to your sensitive data and tenant data from portable and mobile devices (e.g. laptops, cell phones and personal digital assistants (PDAs)), which are generally higher-risk than non-portable devices (e.g., desktop computers at the provider organization's facilities)?	Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
Human Resources <i>Nondisclosure Agreements</i>	HRS -06.1	Are requirements for non-disclosure or confidentiality agreements reflecting the organization's needs for the protection of data and operational details identified, documented and reviewed at planned intervals?	Amazon Legal Counsel manages and periodically revises the Amazon NDA to reflect AWS business needs.
Human Resources <i>Roles / Responsibilities</i>	HRS -07.1	Do you provide tenants with a role definition document clarifying your administrative responsibilities versus those of the tenant?	The AWS Cloud Security Whitepaper and the AWS Risk and Compliance Whitepaper provide details on the roles and responsibilities of AWS and those of our Customers. The whitepapers are available at: http://aws.amazon.com/security and http://aws.amazon.com/compliance .

Control Group	CID	Consensus Assessment Questions	AWS Response
Human Resources <i>Acceptable Use</i>	HRS -08.1	Do you provide documentation regarding how you may or access tenant data and metadata?	AWS has a formal access control policy that is reviewed and updated on an annual basis (or when any major change to the system occurs that impacts the policy). The policy addresses purpose, scope, roles, responsibilities and management commitment. AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content. Refer to the ISO 27001 standard and 27018 code of practice for additional information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 and ISO 27018.
	HRS -08.2	Do you collect or create metadata about tenant data usage through inspection technologies (search engines, etc.)?	
	HRS -08.3	Do you allow tenants to opt out of having their data/metadata accessed via inspection technologies?	
Human Resources <i>Training / Awareness</i>	HRS -09.1	Do you provide a formal, role-based, security awareness training program for cloud-related access and data management issues (e.g., multi-tenancy, nationality, cloud delivery model segregation of duties implications and conflicts of interest) for all persons with access to tenant data?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. AWS roles and responsibilities are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
	HRS -09.2	Are administrators and data stewards properly educated on their legal responsibilities with regard to security and data integrity?	
Human Resources <i>User Responsibility</i>	HRS -10.1	Are users made aware of their responsibilities for maintaining awareness and compliance with published security policies, procedures, standards and applicable regulatory requirements?	AWS has implemented various methods of internal communication at a global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employee as well as electronic mail messages and the posting of information via the Amazon intranet. Refer to ISO 27001 standard, Annex A, domain 7 and 8. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. In addition the AWS Cloud Security Whitepaper provides further details - available at http://aws.amazon.com/security.
	HRS -10.2	Are users made aware of their responsibilities for maintaining a safe and secure working environment?	
	HRS -10.3	Are users made aware of their responsibilities for leaving unattended equipment in a secure manner?	
Human Resources <i>Workspace</i>	HRS -11.1	Do your data management policies and procedures address tenant and service level conflicts of interests?	AWS data management policies are in alignment with ISO 27001 standard. Refer to ISO 27001 standard, Annex A, domain 8 and 9. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS SOC reports

Control Group	CID	Consensus Assessment Questions	AWS Response
	HRS-11.2	Do your data management policies and procedures include a tamper audit or software integrity function for unauthorized access to tenant data?	provides additional details on the specific control activities executed by AWS to prevent unauthorized access to AWS resources.
	HRS-11.3	Does the virtual machine management infrastructure include a tamper audit or software integrity function to detect changes to the build/configuration of the virtual machine?	AWS has identified auditable event categories across systems and devices within the AWS system. Service teams configure the auditing features to record continuously the security-related events in accordance with requirements. Audit records contain a set of data elements in order to support necessary analysis requirements. In addition, audit records are available for AWS Security team or other appropriate teams to perform inspection or analysis on demand, and in response to security-related or business-impacting events.
Identity & Access Management <i>Audit Tools Access</i>	IAM-01.1	Do you restrict, log and monitor access to your information security management systems? (E.g., hypervisors, firewalls, vulnerability scanners, network sniffers, APIs, etc.)	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	IAM-01.2	Do you monitor and log privileged access (administrator level) to information security management systems?	AWS has identified auditable event categories across systems and devices within the AWS system. Service teams configure the auditing features to record continuously the security-related events in accordance with requirements. The log storage system is designed to provide a highly scalable, highly available service that automatically increases capacity as the ensuing need for log storage grows. Audit records contain a set of data elements in order to support necessary analysis requirements. In addition, audit records are available for AWS Security team or other appropriate teams to perform inspection or analysis on demand, and in response to security-related or business-impacting events. Designated personnel on AWS teams receive automated alerts in the event of an audit processing failure. Audit processing failures include, for example, software/hardware errors. When alerted, on-call personnel issue a trouble ticket and track the event until it is resolved. AWS logging and monitoring processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Identity & Access Management <i>User Access Policy</i>	IAM-02.1	Do you have controls in place ensuring timely removal of systems access that is no longer required for business purposes?	AWS SOC reports provides further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 9 for additional details. AWS has

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM - 02.2	Do you provide metrics to track the speed with which you are able to remove systems access that is no longer required for business purposes?	been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Identity & Access Management <i>Diagnostic / Configuration Ports Access</i>	IAM -03.1	Do you use dedicated secure networks to provide management access to your cloud service infrastructure?	Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored per the AWS access policy. In addition, customer data is and server instances are logically isolated from other customers by default. Privileged user access controls are reviewed by an independent auditor during the AWS SOC, ISO 27001, PCI, ITAR, and FedRAMP audits.
Identity & Access Management <i>Policies and Procedures</i>	IAM -04.1	Do you manage and store the identity of all personnel who have access to the IT infrastructure, including their level of access?	
	IAM - 04.2	Do you manage and store the user identity of all personnel who have network access, including their level of access?	
Identity & Access Management <i>Segregation of Duties</i>	IAM -05.1	Do you provide tenants with documentation on how you maintain segregation of duties within your cloud service offering?	Customers retain the ability to manage segregations of duties of their AWS resources. Internally, AWS aligns with ISO 27001 standards for managing segregation of duties. Refer to ISO 27001 standard, Annex A, domain 6 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Identity & Access Management <i>Source Code Access Restriction</i>	IAM -06.1	Are controls in place to prevent unauthorized access to your application, program or object source code, and assure it is restricted to authorized personnel only?	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IAM - 06.2	Are controls in place to prevent unauthorized access to tenant application, program or object source code, and assure it is restricted to authorized personnel only?	
Identity & Access Management <i>Third Party Access</i>	IAM -07.1	Do you provide multi-failure disaster recovery capability?	AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS SOC reports provides further details. ISO 27001 standard Annex A, domain 15 provides additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification.
	IAM - 07.2	Do you monitor service continuity with upstream providers in the event of provider failure?	
	IAM - 07.3	Do you have more than one provider for each service you depend on?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM - 07.4	Do you provide access to operational redundancy and continuity summaries, including the services you depend on?	
	IAM -07.5	Do you provide the tenant the ability to declare a disaster?	
	IAM - 07.6	Do you provided a tenant-triggered failover option?	
	IAM -07.7	Do you share your business continuity and redundancy plans with your tenants?	
Identity & Access Management <i>User Access Restriction / Authorization</i>	IAM -08.1	Do you document how you grant and approve access to tenant data?	AWS Customers retain control and ownership of their data. Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored. In addition, customer data and server instances are logically isolated from other customers by default. Privileged user access controls are reviewed by an independent auditor during the AWS SOC, ISO 27001, PCI, ITAR, and FedRAMP audits.
	IAM - 08.2	Do you have a method of aligning provider and tenant data classification methodologies for access control purposes?	
Identity & Access Management <i>User Access Authorization</i>	IAM -09.1	Does your management provision the authorization and restrictions for user access (e.g. employees, contractors, customers (tenants), business partners and/or suppliers) prior to their access to data and any owned or managed (physical and virtual) applications, infrastructure systems and network components?	Unique user identifiers are created as part of the onboarding workflow process in the AWS human resources management system. The device provisioning process helps ensure unique identifiers for devices. Both processes include manager approval to establish the user account or device. Initial authenticators are delivered to user's in-person and to devices as part of the provisioning process. Internal users can associate SSH public keys with their account. System account authenticators are provided to the requestor as part of the account creation process after the identity of the requestor is verified.
	IAM - 09.2	Do you provide upon request user access (e.g. employees, contractors, customers (tenants), business partners and/or suppliers) to data and any owned or managed (physical and virtual) applications, infrastructure systems and network components?	AWS has established controls to address the threat of inappropriate insider access. All certifications and third-party attestations evaluate logical access preventative and detective controls. In addition, periodic risk assessments focus on how insider access is controlled and monitored.
Identity & Access Management <i>User Access Reviews</i>	IAM -10.1	Do you require at least annual certification of entitlements for all system users and administrators (exclusive of users maintained by your tenants)?	In alignment with ISO 27001 standard, all access grants are reviewed on a periodic basis; explicit re-approval is required or access to the resource is automatically revoked. Controls specific to User Access reviews are outlined in the SOC reports. Exceptions in the User entitlement controls are documented in the SOC reports. Refer to ISO 27001 standards, Annex A, domain 9 for additional details.

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM-10.2	If users are found to have inappropriate entitlements, are all remediation and certification actions recorded?	AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IAM-10.3	Will you share user entitlement remediation and certification reports with your tenants, if inappropriate access may have been allowed to tenant data?	
Identity & Access Management <i>User Access Revocation</i>	IAM-11.1	Is timely deprovisioning, revocation or modification of user access to the organizations systems, information assets and data implemented upon any change in status of employees, contractors, customers, business partners or involved third parties?	Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked. AWS SOC reports provides further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 9 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IAM-11.2	Is any change in user access status intended to include termination of employment, contract or agreement, change of employment or transfer within the organization?	
Identity & Access Management <i>User ID Credentials</i>	IAM-12.1	Do you support use of, or integration with, existing customer-based Single Sign On (SSO) solutions to your service?	The AWS Identity and Access Management (IAM) service provides identity federation to the AWS Management Console. Multi-factor authentication is an optional feature that a customer can utilize. Refer to the AWS website for additional details - http://aws.amazon.com/mfa . AWS Identity and Access Management (IAM) supports identity federation for delegated access to the AWS Management Console or AWS APIs. With identity federation, external identities (federated users) are granted secure access to resources in your AWS account without having to create IAM users. These external identities can come from your corporate identity provider (such as Microsoft Active Directory or from the AWS Directory Service) or from a web identity provider, such as Amazon Cognito, Login with Amazon, Facebook, Google or any OpenID Connect (OIDC) compatible provider.
	IAM-12.2	Do you use open standards to delegate authentication capabilities to your tenants?	
	IAM-12.3	Do you support identity federation standards (SAML, SPML, WS-Federation, etc.) as a means of authenticating/authorizing users?	
	IAM-12.4	Do you have a Policy Enforcement Point capability (e.g., XACML) to enforce regional legal and policy constraints on user access?	
	IAM-12.5	Do you have an identity management system (enabling classification of data for a tenant) in place to enable both role-based and context-based entitlement to data?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IAM -12.6	Do you provide tenants with strong (multifactor) authentication options (digital certs, tokens, biometrics, etc.) for user access?	AWS Identity and Access Management (IAM) enables customers to securely control access to AWS services and resources for their users. Additional information about IAM can be found on website at https://aws.amazon.com/iam/ . AWS SOC reports provides details on the specific control activities executed by AWS.
	IAM -12.7	Do you allow tenants to use third-party identity assurance services?	
	IAM -12.8	Do you support password (minimum length, age, history, complexity) and account lockout (lockout threshold, lockout duration) policy enforcement?	
	IAM -12.9	Do you allow tenants/customers to define password and account lockout policies for their accounts?	
	IAM -12.10	Do you support the ability to force password changes upon first logon?	
	IAM -12.11	Do you have mechanisms in place for unlocking accounts that have been locked out (e.g., self-service via email, defined challenge questions, manual unlock)?	
Identity & Access Management <i>Utility Programs Access</i>	IAM -13.1	Are utilities that can significantly manage virtualized partitions (e.g., shutdown, clone, etc.) appropriately restricted and monitored?	In alignment with ISO 27001 standards, system utilities are appropriately restricted and monitored. AWS SOC reports provides details on the specific control activities executed by AWS. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IAM -13.2	Do you have a capability to detect attacks that target the virtual infrastructure directly (e.g., shimming, Blue Pill, Hyper jumping, etc.)?	
	IAM -13.3	Are attacks that target the virtual infrastructure prevented with technical controls?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>Audit Logging / Intrusion Detection</i>	IVS-01.1	Are file integrity (host) and network intrusion detection (IDS) tools implemented to help facilitate timely detection, investigation by root cause analysis and response to incidents?	AWS Incident response program (detection, investigation and response to incidents) has been developed in alignment with ISO 27001 standards, system utilities are appropriately restricted and monitored. AWS SOC reports provides additional details on controls in place to restrict system access. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IVS-01.2	Is physical and logical user access to audit logs restricted to authorized personnel?	
	IVS-01.3	Can you provide evidence that due diligence mapping of regulations and standards to your controls/architecture/processes has been done?	
	IVS-01.4	Are audit logs centrally stored and retained?	In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-01.5	Are audit logs reviewed on a regular basis for security events (e.g., with automated tools)?	AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Infrastructure & Virtualization Security <i>Change Detection</i>	IVS-02.1	Do you log and alert any changes made to virtual machine images regardless of their running state (e.g. dormant, off or running)?	Virtual Machines are assigned to customers as a part of the EC2 service. Customers retain control over what resources are being used and where resources reside. Refer to the AWS website for additional details - http://aws.amazon.com .
	IVS-02.2	Are changes made to virtual machines, or moving of an image and subsequent validation of the image's integrity, made immediately available to customers through electronic methods (e.g. portals or alerts)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>Clock Synchronization</i>	IVS-03.1	Do you use a synchronized time-service protocol (e.g., NTP) to ensure all systems have a common time reference?	In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-04.1	Do you provide documentation regarding what levels of system (network, storage, memory, I/O, etc.) oversubscription you maintain and under what circumstances/scenarios?	Details regarding AWS Service Limits and how to request an increase for specific services is available on the AWS website at http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html . AWS manages capacity and utilization data in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-04.2	Do you restrict use of the memory oversubscription capabilities present in the hypervisor?	
	IVS-04.3	Do your system capacity requirements take into account current, projected and anticipated capacity needs for all systems used to provide services to the tenants?	
	IVS-04.4	Is system performance monitored and tuned in order to continuously meet regulatory, contractual and business requirements for all the systems used to provide services to the tenants?	
Infrastructure & Virtualization Security <i>Management - Vulnerability Management</i>	IVS-05.1	Do security vulnerability assessment tools or services accommodate the virtualization technologies being used (e.g. virtualization aware)?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. Regular internal and external vulnerability scans are performed on the host operating system, web application and databases in the AWS environment utilizing a variety of tools. Vulnerability scanning and remediation practices are regularly reviewed as a part of AWS continued compliance with PCI DSS and FedRAMP.
Infrastructure & Virtualization Security <i>Network Security</i>	IVS-06.1	For your IaaS offering, do you provide customers with guidance on how to create a layered security architecture equivalence using your virtualized solution?	AWS website provides guidance on creating a layered security architecture in a number of white papers available via the AWS public website - http://aws.amazon.com/documentation/ .
	IVS-06.2	Do you regularly update network architecture diagrams that include data flows between security domains/zones?	Boundary protection devices that employ rule sets, access control lists (ACL), and configurations enforce the flow of information between network fabrics. Several network fabrics exist at Amazon, each separated by devices that

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-06.3	Do you regularly review for appropriateness the allowed access/connectivity (e.g., firewall rules) between security domains/zones within the network?	control the flow of information between fabrics. The flow of information between fabrics is established by approved authorizations, which exist as access control lists (ACL) which reside on these devices. These devices control the flow of information between fabrics as mandated by these ACLs. ACLs are defined, approved by appropriate personnel, managed and deployed using AWS ACL-manage tool.
	IVS-06.4	Are all firewall access control lists documented with business justification?	Amazon's Information Security team approves these ACLs. Approved firewall rule sets and access control lists between network fabrics restrict the flow of information to specific information system services. Access control lists and rule sets are reviewed and approved, and are automatically pushed to boundary protection devices on a periodic basis (at least every 24 hours) to ensure rule-sets and access control lists are up-to-date.
Infrastructure & Virtualization Security <i>OS Hardening and Base Controls</i>	IVS-07.1	Are operating systems hardened to provide only the necessary ports, protocols and services to meet business needs using technical controls (i.e antivirus, file integrity monitoring and logging) as part of their baseline build standard or template?	AWS Network Management is regularly reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm. AWS implements least privilege throughout its infrastructure components. AWS prohibits all ports and protocols that do not have a specific business purpose. AWS follows a rigorous approach to minimal implementation of only those features and functions that are essential to use of the device. Network scanning is performed and any unnecessary ports or protocols in use are corrected. Regular internal and external vulnerability scans are performed on the host operating system, web application and databases in the AWS environment utilizing a variety of tools. Vulnerability scanning and remediation practices are regularly reviewed as a part of AWS continued compliance with PCI DSS and FedRAMP.
	IVS-08.1	For your SaaS or PaaS offering, do you provide tenants with separate environments for production and test processes?	AWS Customers retain the ability and the responsibility to create and maintain production and test environments. AWS website provides guidance on creating an environment utilizing the AWS services - http://aws.amazon.com/documentation/ .
Infrastructure & Virtualization Security <i>Production / Nonproduction Environments</i>	IVS-08.2	For your IaaS offering, do you provide tenants with guidance on how to create suitable production and test environments?	
	IVS-08.3	Do you logically and physically segregate production and non-production environments?	AWS Customers retain responsibility to manage their own network segmentation in adherence with their defined requirements. Internally, AWS network segmentation is aligned with ISO 27001 standards. Refer to ISO 27001 standard, Annex A. domain 13 for further detail. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Infrastructure & Virtualization Security <i>Segmentation</i>	IVS-09.1	Are system and network environments protected by a firewall or virtual firewall to ensure business and customer security requirements?	
	IVS-09.2	Are system and network environments protected by a firewall or virtual firewall to ensure compliance with legislative, regulatory and contractual requirements?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-09.3	Are system and network environments protected by a firewall or virtual firewall to ensure separation of production and non-production environments?	
	IVS-09.4	Are system and network environments protected by a firewall or virtual firewall to ensure protection and isolation of sensitive data?	
Infrastructure & Virtualization Security <i>VM Security - vMotion Data Protection</i>	IVS-10.1	Are secured and encrypted communication channels used when migrating physical servers, applications or data to virtual servers?	AWS provides customers the ability to use their own encryption mechanism for nearly all services including S3, EBS and EC2. VPC sessions are also encrypted.
	IVS-10.2	Do you use a network segregated from production-level networks when migrating physical servers, applications or data to virtual servers?	AWS Customers retain control and ownership of their own data. AWS provides customers the ability to maintain and develop production and non-production environments. It is the responsibility of the customer to ensure that their production data is not replicated to non-production environments.
Infrastructure & Virtualization Security <i>VMM Security - Hypervisor Hardening</i>	IVS-11.1	Do you restrict personnel access to all hypervisor management functions or administrative consoles for systems hosting virtualized systems based on the principle of least privilege and supported through technical controls (e.g. two-factor authentication, audit trails, IP address filtering, firewalls and TLS-encapsulated communications to the administrative consoles)?	AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. When user accounts are created, user accounts are created to have minimal access. Access above these least privileges requires appropriate authorization. Refer to AWS SOC reports for more information on Access Controls.
Infrastructure & Virtualization Security <i>Wireless Security</i>	IVS-12.1	Are policies and procedures established and mechanisms configured and implemented to protect the wireless network environment perimeter and to restrict unauthorized wireless traffic?	Policies, procedures and mechanisms to protect AWS network environment are in place. AWS security controls are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
	IVS-12.2	Are policies and procedures established and mechanisms implemented to ensure wireless security settings are enabled with strong encryption for authentication and transmission, replacing vendor default settings? (e.g., encryption keys, passwords, SNMP community strings)	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-12.3	Are policies and procedures established and mechanisms implemented to protect wireless network environments and detect the presence of unauthorized (rogue) network devices for a timely disconnect from the network?	
Infrastructure & Virtualization Security <i>Network Architecture</i>	IVS-13.1	Do your network architecture diagrams clearly identify high-risk environments and data flows that may have legal compliance impacts?	AWS Customers retain responsibility to manage their own network segmentation in adherence with their defined requirements. Internally, AWS network segmentation is aligned with the ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-13.2	Do you implement technical measures and apply defense-in-depth techniques (e.g., deep packet analysis, traffic throttling and black-holing) for detection and timely response to network-based attacks associated with anomalous ingress or egress traffic patterns (e.g., MAC spoofing and ARP poisoning attacks) and/or distributed denial-of-service (DDoS) attacks?	AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment. AWS security controls are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Interoperability & Portability APIs	IPY-01	Do you publish a list of all APIs available in the service and indicate which are standard and which are customized?	Details regarding AWS APIs can be found on the AWS website at https://aws.amazon.com/documentation/. In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources.
Interoperability & Portability <i>Data Request</i>	IPY-02	Is unstructured customer data available on request in an industry-standard format (e.g., .doc, .xls, or .pdf)?	Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
Interoperability & Portability <i>Policy & Legal</i>	IPY-03.1	Do you provide policies and procedures (i.e. service level agreements) governing the use of APIs for interoperability between your service and third-party applications?	
	IPY-03.2	Do you provide policies and procedures (i.e. service level agreements) governing the migration of application data to and from your service?	Customer retain control and ownership of their content. Customers can choose how they migrate applications and content both on and off the AWS platform at their discretion.

Control Group	CID	Consensus Assessment Questions	AWS Response
Interoperability & Portability <i>Standardized Network Protocols</i>	IPY-04.1	Can data import, data export and service management be conducted over secure (e.g., non-clear text and authenticated), industry accepted standardized network protocols?	AWS allows customers to move data as needed on and off AWS storage. Refer to http://aws.amazon.com/choosing-a-cloud-platform for more information on Storage options.
	IPY-04.2	Do you provide consumers (tenants) with documentation detailing the relevant interoperability and portability network protocol standards that are involved?	
Interoperability & Portability <i>Virtualization</i>	IPY-05.1	Do you use an industry-recognized virtualization platform and standard virtualization formats (e.g., OVF) to help ensure interoperability?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	IPY-05.2	Do you have documented custom changes made to any hypervisor in use, and all solution-specific virtualization hooks available for customer review?	
Mobile Security <i>Anti-Malware</i>	MOS-01	Do you provide anti-malware training specific to mobile devices as part of your information security awareness training?	AWS' program, processes and procedures to managing antivirus / malicious software is in alignment with ISO 27001 standards. Refer to ISO 27001 standard, Annex A, domain 12 for additional information.
Mobile Security <i>Application Stores</i>	MOS-02	Do you document and make available lists of approved application stores for mobile devices accessing or storing company data and/or company systems?	AWS has established an information security framework and policies and has effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.1 and the National Institute of Standards and Technology (NIST) Publication 800-53 (Recommended Security Controls for Federal Information Systems). Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
Mobile Security <i>Approved Applications</i>	MOS-03	Do you have a policy enforcement capability (e.g., XACML) to ensure that only approved applications and those from approved application stores be loaded onto a mobile device?	
Mobile Security <i>Approved Software for BYOD</i>	MOS-04	Does your BYOD policy and training clearly state which applications and applications stores are approved for use on BYOD devices?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Mobile Security <i>Awareness and Training</i>	MOS-05	Do you have a documented mobile device policy in your employee training that clearly defines mobile devices and the accepted usage and requirements for mobile devices?	
Mobile Security <i>Cloud Based Services</i>	MOS-06	Do you have a documented list of pre-approved cloud based services that are allowed to be used for use and storage of company business data via a mobile device?	
Mobile Security <i>Compatibility</i>	MOS-07	Do you have a documented application validation process for testing device, operating system and application compatibility issues?	
Mobile Security <i>Device Eligibility</i>	MOS-08	Do you have a BYOD policy that defines the device(s) and eligibility requirements allowed for BYOD usage?	
Mobile Security <i>Device Inventory</i>	MOS-09	Do you maintain an inventory of all mobile devices storing and accessing company data which includes device status (os system and patch levels, lost or decommissioned, device assignee)?	
Mobile Security <i>Device Management</i>	MOS-10	Do you have a centralized mobile device management solution deployed to all mobile devices that are permitted to store, transmit, or process company data?	
Mobile Security <i>Encryption</i>	MOS-11	Does your mobile device policy require the use of encryption for either the entire device or for data identified as sensitive enforceable through technology controls for all mobile devices?	
Mobile Security <i>Jailbreaking and Rooting</i>	MOS-12.1	Does your mobile device policy prohibit the circumvention of built-in security controls on mobile devices (e.g., jailbreaking or rooting)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	MOS -12.2	Do you have detective and preventative controls on the device or via a centralized device management system which prohibit the circumvention of built-in security controls?	
Mobile Security <i>Legal</i>	MOS -13.1	Does your BYOD policy clearly define the expectation of privacy, requirements for litigation, e-discovery and legal holds?	Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
	MOS -13.2	Do you have detective and preventative controls on the device or via a centralized device management system which prohibit the circumvention of built-in security controls?	
Mobile Security <i>Lockout Screen</i>	MOS -14	Do you require and enforce via technical controls an automatic lockout screen for BYOD and company owned devices?	
Mobile Security <i>Operating Systems</i>	MOS -15	Do you manage all changes to mobile device operating systems, patch levels and applications via your company's change management processes?	
Mobile Security <i>Passwords</i>	MOS -16.1	Do you have password policies for enterprise issued mobile devices and/or BYOD mobile devices?	
	MOS -16.2	Are your password policies enforced through technical controls (i.e. MDM)?	
	MOS -16.3	Do your password policies prohibit the changing of authentication requirements (i.e. password/PIN length) via a mobile device?	
Mobile Security <i>Policy</i>	MOS -17.1	Do you have a policy that requires BYOD users to perform backups of specified corporate data?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	MOS -17.2	Do you have a policy that requires BYOD users to prohibit the usage of unapproved application stores?	
	MOS -17.3	Do you have a policy that requires BYOD users to use anti-malware software (where supported)?	
Mobile Security <i>Remote Wipe</i>	MOS -18.1	Does your IT provide remote wipe or corporate data wipe for all company-accepted BYOD devices?	
	MOS -18.2	Does your IT provide remote wipe or corporate data wipe for all company-assigned mobile devices?	
Mobile Security <i>Security Patches</i>	MOS -19.1	Do your mobile devices have the latest available security-related patches installed upon general release by the device manufacturer or carrier?	
	MOS -19.2	Do your mobile devices allow for remote validation to download the latest security patches by company IT personnel?	
Mobile Security <i>Users</i>	MOS -20.1	Does your BYOD policy clarify the systems and servers allowed for use or access on the BYOD-enabled device?	AWS maintains contacts with industry bodies, risk and compliance organizations, local authorities and regulatory bodies as required by the ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. The AWS SOC reports provides details on the specific control activities executed by AWS. All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities.
	MOS -20.2	Does your BYOD policy specify the user roles that are allowed access via a BYOD-enabled device?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Contact / Authority Maintenance</i>	SEF-01.1	Do you maintain liaisons and points of contact with local authorities in accordance with contracts and appropriate regulations?	
	SEF-02.1	Do you have a documented security incident response plan?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Management</i>	SEF-02.2	Do you integrate customized tenant requirements into your security incident response plans?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	SEF-02.3	Do you publish a roles and responsibilities document specifying what you vs. your tenants are responsible for during security incidents?	The AWS Cloud Security Whitepaper (available at http://aws.amazon.com/security) provides additional details.
	SEF-02.4	Have you tested your security incident response plans in the last year?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Reporting</i>	SEF-03.1	Does your security information and event management (SIEM) system merge data sources (app logs, firewall logs, IDS logs, physical access logs, etc.) for granular analysis and alerting?	
	SEF-03.2	Does your logging and monitoring framework allow isolation of an incident to specific tenants?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Response Legal Preparation</i>	SEF-04.1	Does your incident response plan comply with industry standards for legally admissible chain-of-custody management processes and controls?	
	SEF-04.2	Does your incident response capability include the use of legally admissible forensic data collection and analysis techniques?	
	SEF-04.3	Are you capable of supporting litigation holds (freeze of data from a specific point in time) for a specific tenant without freezing other tenant data?	
	SEF-04.4	Do you enforce and attest to tenant data separation when producing data in response to legal subpoenas?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Response Metrics</i>	SEF-05.1	Do you monitor and quantify the types, volumes and impacts on all information security incidents?	
	SEF-05.2	Will you share statistical information for security incident data with your tenants upon request?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Supply Chain Management, Transparency and Accountability <i>Data Quality and Integrity</i>	STA-01.1	Do you inspect and account for data quality errors and associated risks, and work with your cloud supply-chain partners to correct them?	Customers retain control and ownership over the quality of their data and potential quality errors that may arise through their usage of AWS services. Refer to AWS SOC report for specific details in relation to Data Integrity and Access Management (including least privilege access)
	STA-01.2	Do you design and implement controls to mitigate and contain data security risks through proper separation of duties, role-based access, and least-privileged access for all personnel within your supply chain?	
Supply Chain Management, Transparency and Accountability <i>Incident Reporting</i>	STA-02.1	Do you make security incident information available to all affected customers and providers periodically through electronic methods (e.g. portals)?	AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC reports provides details on the specific control activities executed by AWS. The AWS Cloud Security Whitepaper (available at http://aws.amazon.com/security) provides additional details.
Supply Chain Management, Transparency and Accountability <i>Network / Infrastructure Services</i>	STA-03.1	Do you collect capacity and use data for all relevant components of your cloud service offering?	AWS manages capacity and utilization data in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	STA-03.2	Do you provide tenants with capacity planning and use reports?	
Supply Chain Management, Transparency and Accountability <i>Provider Internal Assessments</i>	STA-04.1	Do you perform annual internal assessments of conformance and effectiveness of your policies, procedures, and supporting measures and metrics?	AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 15 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Supply Chain Management, Transparency and Accountability <i>Third Party Agreements</i>	STA-05.1	Do you select and monitor outsourced providers in compliance with laws in the country where the data is processed, stored and transmitted?	Personnel security requirements for third-party providers supporting AWS systems and devices are established in a Mutual Non-Disclosure Agreement between AWS' parent organization, Amazon.com, and the respective third-party provider. The Amazon Legal Counsel and the AWS Procurement team define AWS third party provider personnel security requirements in contract agreements with the third party provider. All persons working with AWS information must at a minimum, meet the screening process for pre-employment background checks and sign a Non-Disclosure Agreement (NDA) prior to being granted access to AWS information. AWS does not generally outsource development of AWS services to subcontractors.
	STA-05.2	Do you select and monitor outsourced providers in compliance with laws in the country where the data originates?	
	STA-05.3	Does legal counsel review all third-party agreements?	
	STA-05.4	Do third-party agreements include provision for the security and protection of information and assets?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	STA-05.5	Do you provide the client with a list and copies of all sub processing agreements and keep this updated?	
Supply Chain Management, Transparency and Accountability <i>Supply Chain Governance Reviews</i>	STA-06.1	Do you review the risk management and governance processes of partners to account for risks inherited from other members of that partner's supply chain?	AWS maintains formal agreements with key third party suppliers and implements appropriate relationship management mechanisms in line with their relationship to the business. AWS' third party management processes are reviewed by independent auditors as part of AWS ongoing compliance with SOC and ISO 27001.
Supply Chain Management, Transparency and Accountability <i>Supply Chain Metrics</i>	STA-07.1	Are policies and procedures established, and supporting business processes and technical measures implemented, for maintaining complete, accurate and relevant agreements (e.g., SLAs) between providers and customers (tenants)?	
	STA-07.2	Do you have the ability to measure and address non-conformance of provisions and/or terms across the entire supply chain (upstream/downstream)?	
	STA-07.3	Can you manage service-level conflicts or inconsistencies resulting from disparate supplier relationships?	
	STA-07.4	Do you review all agreements, policies and processes at least annually?	
Supply Chain Management, Transparency and Accountability <i>Third Party Assessment</i>	STA-08.1	Do you assure reasonable information security across your information supply chain by performing an annual review?	
	STA-8.2	Does your annual review include all partners/third-party providers upon which your information supply chain depends?	
Supply Chain Management, Transparency	STA-09.1	Do you permit tenants to perform independent vulnerability assessments?	

Control Group	CID	Consensus Assessment Questions	AWS Response
and Accountability <i>Third Party Audits</i>	STA-09.2	Do you have external third party services conduct vulnerability scans and periodic penetration tests on your applications and networks?	these types of scans can be initiated by submitting a request via the AWS Vulnerability / Penetration Testing Request Form . AWS Security regularly engages independent security firms to perform external vulnerability threat assessments. The AWS SOC reports provides additional details on the specific control activities executed by AWS.
Threat and Vulnerability Management <i>Antivirus / Malicious Software</i>	TVM-01.1	Do you have anti-malware programs that support or connect to your cloud service offerings installed on all of your systems?	AWS' program, processes and procedures to managing antivirus / malicious software is in alignment with ISO 27001 standards. Refer to AWS SOC reports provides further details. In addition, refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	TVM-01.2	Do you ensure that security threat detection systems using signatures, lists or behavioral patterns are updated across all infrastructure components within industry accepted time frames?	
Threat and Vulnerability Management <i>Vulnerability / Patch Management</i>	TVM-02.1	Do you conduct network-layer vulnerability scans regularly as prescribed by industry best practices?	Customers retain control of their own guest operating systems, software and applications and are responsible for performing vulnerability scans and patching of their own systems. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. AWS Security regularly scans all Internet-facing service endpoint IP addresses for vulnerabilities. AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. AWS' own maintenance and system patching generally do not impact customers. Refer to AWS Cloud Security Whitepaper for further information - available at http://aws.amazon.com/security . Refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	TVM-02.2	Do you conduct application-layer vulnerability scans regularly as prescribed by industry best practices?	
	TVM-02.3	Do you conduct local operating system-layer vulnerability scans regularly as prescribed by industry best practices?	
	TVM-02.4	Will you make the results of vulnerability scans available to tenants at their request?	
	TVM-02.5	Do you have a capability to rapidly patch vulnerabilities across all of your computing devices, applications and systems?	
	TVM-02.6	Will you provide your risk-based systems patching time frames to your tenants upon request?	
Threat and Vulnerability Management <i>Mobile Code</i>	TVM-03.1	Is mobile code authorized before its installation and use, and the code configuration checked, to ensure that the authorized mobile code operates according to a clearly defined security policy?	AWS allows customers to manage client and mobile applications to their own requirements.

Control Group	CID	Consensus Assessment Questions	AWS Response
	TVM - 03.2	Is all unauthorized mobile code prevented from executing?	

Appendix B: AWS alignment with the Australian Signals Directorate (ASD) Cloud Computing Security Considerations

The Cloud Computing Security Considerations was created to assist agencies in performing a risk assessment of services offered by Cloud Service Providers. The following provides AWS alignment to the Security Considerations, published on September 2012. For additional details refer to:

http://www.asd.gov.au/publications/csocprotect/Cloud_Computing_Security_Considerations.pdf

Key Area	Questions	AWS RESPONSE
Maintaining Availability and Business Functionality	a. Business criticality of data or functionality. Am I moving business critical data or functionality to the cloud?	AWS customers retain control and ownership of their content. Customers are responsible for the classification and use of their content.
	b. Vendor's business continuity and disaster recovery plan. Can I thoroughly review a copy of the vendor's business continuity and disaster recovery plan that covers the availability and restoration of both my data and the vendor's services that I use? How much time does it take for my data and the services that I use to be recovered after a disaster, and do the vendor's other customers that are larger and pay more money than me get prioritization?	AWS customers retain control and ownership of their data. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS SOC 1 Type 2 report provides further details. ISO 27001 standard Annex A, domain 11 provides additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification. Customers utilize AWS to enable faster disaster recovery of their critical IT systems without incurring the infrastructure expense of a second physical site. The AWS cloud supports many popular disaster recovery (DR) architectures from "pilot light" environments that are ready to scale up at a moment's notice to "hot standby" environments that enable rapid failover. To learn more about Disaster Recovery on AWS visit https://aws.amazon.com/disaster-recovery/. AWS provides customers with the capability to implement a robust continuity plan, including the utilization of frequent server instance back-ups, data redundancy replication, and multi-region/availability zone deployment architectures. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 9 and the AWS SOC 1 Type II report for additional information.

Key Area	Questions	AWS RESPONSE
	c. My data backup plan. Will I spend additional money to maintain an up to date backup copy of my data located either at my agency's premises, or stored with a second vendor that has no common points of failure with the first vendor?	AWS customers retain control and ownership of their content and it is the customer's responsibility to manage their data backup plans. AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 service is designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy. For information on data durability and redundancy, please refer to the AWS website. AWS offers a range of cloud computing services to support Disaster Recovery. To learn more about Disaster Recovery on AWS visit https://aws.amazon.com/disaster-recovery/.
	d. My business continuity and disaster recovery plan. Will I spend additional money to replicate my data or business functionality with a second vendor that uses a different data center and ideally has no common points of failure with the first vendor? This replication should preferably be configured to automatically "failover", so that if one vendor's services become unavailable, control is automatically and smoothly transitioned to the other vendor.	Customers retain control and ownership of their data. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security. AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS. AWS data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. This means that Availability Zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by Region). In addition to discrete uninterruptable power supply (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability Zones are all redundantly connected to multiple tier-1 transit providers. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Distributing applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes including natural disasters or system failures. AWS SOC 1 Type 2 report provides further details. ISO 27001 standard Annex A, domain 11 provides additional details. AWS has been validated

Key Area	Questions	AWS RESPONSE
		and certified by an independent auditor to confirm alignment with ISO 27001 certification.
	e. My network connectivity to the cloud. Is the network connectivity between my agency's users and the vendor's network adequate in terms of availability, traffic throughput (bandwidth), delays (latency) and packet loss?	Customers can also choose their network path to AWS facilities, including multiple VPN endpoints in each AWS Region. In addition, AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security.

Key Area	Questions	AWS RESPONSE
	f. Vendor's guarantee of availability. Does the Service Level Agreement (SLA guarantee that the vendor will provide adequate system availability and quality of service, using their robust system architecture and business processes?	AWS does commit to high levels of availability in its service level agreements (SLAs). For example, Amazon EC2 commits to annual uptime percentage of at least 99.95% during the service year. Amazon S3 commits to monthly uptime percentage of at least 99.99% Service credits are provided in the case these availability metrics are not met. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Distributing applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes including natural disasters or system failures. AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel. AWS Network Management is regularly reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm.
	g. Impact of outages. Can I tolerate the maximum possible downtime of the SLA? Are the scheduled outage windows acceptable both in duration and time of day, or will scheduled outages interfere with my critical business processes?	AWS does not require systems to be brought offline to perform regular maintenance and system patching. AWS' own maintenance and system patching generally do not impact customers. Maintenance of instances themselves is controlled by the customer.
	h. SLA inclusion of scheduled outages. Does the SLA guaranteed availability percentage include scheduled outages?	AWS does not operate an environment with scheduled outage as AWS provides customers the ability to architect their environment to take advantage of multiple Availability Zones and regions.
	i. SLA compensation. Does the SLA adequately reflect the actual damage caused by a breach of the SLA such as unscheduled downtime or data loss?	AWS provides customer remuneration for losses they may incur due to outages in alignment with AWS' Service Level Agreement.

Key Area	Questions	AWS RESPONSE
	j. Data integrity and availability. How does the vendor implement mechanisms such as redundancy and offsite backups to prevent corruption or loss of my data, and guarantee both the integrity and the availability of my data?	AWS data integrity controls as described in AWS SOC 1 Type II report provides reasonable assurance that data integrity is maintained through all phases including transmission, storage and processing. In addition, refer to ISO 27001 standard, Annex A, domain 12 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. Data centers are built in clusters in various global regions. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. You choose where to store your data by specifying a region (for Amazon S3) or an availability zone within a region (for EBS). Data stored in Amazon Elastic Block Store (Amazon EBS) is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones. Amazon S3 provides a highly durable storage infrastructure. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 Region. Once stored, Amazon S3 maintains the durability of objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. Data stored in S3 is designed to provide 99.99999999% durability and 99.99% availability of objects over a given year. Refer to AWS Overview of Security Processes whitepaper for additional details - available at http://aws.amazon.com/security
	k. Data restoration. If I accidentally delete a file, email or other data, how much time does it take for my data to be partially or fully restored from backup, and is the maximum acceptable time captured in the SLA?	AWS customers retain control and ownership of their data. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region.
	l. Scalability. How much available spare computing resources does the vendor provide to enable my usage of the vendor's services to scale at short notice?	The AWS cloud is distributed, highly secure and resilient, giving customers large scaling potential. Customers may scale up or down, paying for only what they use.

Key Area	Questions	AWS RESPONSE
	m. Changing vendor. If I want to move my data to my agency or to a different vendor, or if the vendor suddenly becomes bankrupt or otherwise quits the cloud business, how do I get access to my data in a vendor-neutral format to avoid vendor lock-in? How cooperative will the vendor be? How do I ensure that my data is permanently deleted from the vendor's storage media? For Platform as a Service, which standards does the vendor use that facilitate portability and interoperability to easily move my application to a different vendor or to my agency?	Customers retain control and ownership of their data. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security. AWS allows customers to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport. AWS allows customers to perform their own backups to tapes using their own tape backup service provider. However, a tape backup is not a service provided by AWS.
Protecting Data from Unauthorized Access by a Third Party	a. Choice of cloud deployment model. Am I considering using a potentially less secure public cloud, a potentially more secure hybrid cloud or community cloud, or a potentially most secure private cloud?	AWS' Compliance and Security teams have established an information security framework and policies based on the Control Objectives for Information and related Technology (COBIT) framework. The AWS security framework integrates the ISO 27002 best practices and the PCI Data Security Standard. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security. AWS provides third-party attestations, certifications, Service Organization Controls 1 (SOC 1) Type II report and other relevant compliance reports directly to our customers under NDA. Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can easily customize the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your web servers that has access to the Internet, and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet. Additionally, you can create a Hardware Virtual Private Network (VPN) connection between your corporate data center and your VPC and leverage the AWS cloud as an extension of your corporate data center

Key Area	Questions	AWS RESPONSE
	b. Sensitivity of my data. Is my data to be stored or processed in the cloud classified, sensitive, private, or data that is publicly available such as information from my public web site? Does the aggregation of my data make it more sensitive than any individual piece of data? For example, the sensitivity may increase if storing a significant amount of data, or storing a variety of data that if compromised would facilitate identity theft. If there is a data compromise, could I demonstrate my due diligence to senior management, government officials and the public?	AWS customers retain control and ownership of their data and may implement a structured data-classification program to meet their requirements.
	c. Legislative obligations. What obligations do I have to protect and manage my data under various legislation, for example the Privacy Act, the Archives Act, as well as other legislation specific to the type of data? Will the vendor contractually accept adhering to these obligations to help me ensure that the obligations are met to the satisfaction of the Australian Government?	AWS customers retain responsibility to ensure their usage of AWS is within compliance of applicable laws and regulations. AWS communicates its security and control environment to customers through industry certifications and third-party attestations, white papers (available at http://aws.amazon.com/security) and providing certifications, reports and other relevant documentation directly to AWS customers. AWS has published a whitepaper on using AWS in the context of Australian privacy considerations, available here.

Key Area	Questions	AWS RESPONSE
	d. Countries with access to my data. In which countries is my data stored, backed up and processed? Which foreign countries does my data transit? In which countries is the failover or redundant data centers? Will the vendor notify me if the answers to these questions change?	AWS customers choose the AWS Region or regions in which their content and servers will be located. This allows customers with geographic specific requirements to establish environments in a location of their choice. AWS customers in Australia can choose to deploy their AWS services exclusively in the Asia Pacific (Sydney) region and store their content onshore in Australia. If the customer makes this choice, their content will be located in Australia unless the customer chooses to move the data. Customers can replicate and back up content in more than one region, but AWS does not move or replicate customer content outside of the customer's chosen region or regions. AWS is vigilant about customers' security and does not disclose or move data in response to a request from the Australian, U.S. or other government unless legally required to do so in order to comply with a legally valid and binding order, such as a subpoena or a court order, or as is otherwise required by applicable law. Non-U.S. governmental or regulatory bodies typically must use recognized international processes, such as Mutual Legal Assistance Treaties with the U.S. government, to obtain valid and binding orders. Additionally, our practice is to notify customers where practicable before disclosing their content so they can seek protection from disclosure, unless we are legally prevented from doing so.

Key Area	Questions	AWS RESPONSE
	e. Data encryption technologies. Are hash algorithms, encryption algorithms and key lengths deemed appropriate by the DSD ISM used to protect my data when it is in transit over a network, and stored on both the vendor's computers and on backup media? The ability to encrypt data while it is being processed by the vendor's computers is still an emerging technology and is an area of current research by industry and academia. Is the encryption deemed strong enough to protect my data for the duration of time that my data is sensitive?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. VPC sessions are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Customers may also use third-party encryption technologies. Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications. AWS cryptographic processes are reviewed by independent third-party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm. The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance. The AWS CloudHSM service works with Amazon Virtual Private Cloud (VPC). CloudHSMs are provisioned inside your VPC with an IP address that you specify, providing simple and private network connectivity to your Amazon Elastic Compute Cloud (EC2) instances. Placing CloudHSMs near your EC2 instances decreases network latency, which can improve application performance. AWS provides dedicated and exclusive access to CloudHSMs, isolated from other AWS customers. Available in multiple Regions and Availability Zones (AZs), AWS CloudHSM allows you to add secure and durable key storage to your Amazon EC2 applications.
	f. Media sanitization. What processes are used to sanitize the storage media storing my data at its end of life, and are the processes deemed appropriate by the DSD ISM?	When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security.

Key Area	Questions	AWS RESPONSE
	g. Vendor's remote monitoring and management. Does the vendor monitor, administer or manage the computers that store or process my data? If yes, is this performed remotely from foreign countries or from Australia? Can the vendor provide patch compliance reports and other details about the security of workstations used to perform this work, and what controls prevent the vendor's employees from using untrustworthy personally owned laptops?	Moving IT infrastructure to AWS services creates a model of shared responsibility between the customer and AWS. This shared model can help relieve customer's operational burden as AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall.
	h. My monitoring and management. Can I use my existing tools for integrity checking, compliance checking, security monitoring and network management, to obtain visibility of all my systems regardless of whether these systems are located locally or in the cloud? Do I have to learn to use additional tools provided by the vendor? Does the vendor even provide such a mechanism for me to perform monitoring?	AWS Cloudwatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com. The AWS Trusted Advisor inspects your AWS environment and makes recommendations when opportunities exist to save money, improve system performance and reliability, or help close security gap.
	i. Data ownership. Do I retain legal ownership of my data, or does it belong to the vendor and may be considered an asset for sale by liquidators if the vendor declares bankruptcy?	AWS customers retain ownership and control of their data. AWS only uses each customer's content to provide the AWS services selected by each customer to that customer and does not use customer content for any secondary purposes. AWS treats all customer content the same and has no insight as to what type of content the customer chooses to store in AWS. AWS simply makes available the compute, storage, database and networking services selected by customer – AWS does not require access to customer content to provide its services.

Key Area	Questions	AWS RESPONSE
	j. Gateway technologies. What technologies does the vendor use to create a secure gateway environment? Examples include firewalls, traffic flow filters, content filters, and antivirus software and data diodes where appropriate.	The AWS network provides significant protection against traditional network security issues and customers can implement further protection. Refer to the AWS Overview of Security whitepaper (available at http://aws.amazon.com/security) for additional details. Amazon assets (e.g. laptops) are configured with anti-virus software that includes e-mail filtering and malware detection. AWS Network Firewall management and Amazon's anti-virus program are reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMPsm.
	k. Gateway certification. Is the vendor's gateway environment certified against government security standards and regulations?	AWS obtains certain industry certifications and independent third-party attestations which include the AWS Gateway environment.
	l. Email content filtering. For email Software as a Service, does the vendor provide customizable email content filtering that can enforce my agency's email content policy?	A Customer can utilize a system to host e-mail capabilities, however in that case it is the Customer's responsibility to employ the appropriate levels of spam and malware protection at e-mail entry and exit points and update spam and malware definitions when new releases are made available.

Key Area	Questions	AWS RESPONSE
	m. Policies and processes supporting the vendor's IT security posture. Can I have details of how the vendor's computer and network security posture is supported by policies and processes including threat and risk assessments, ongoing vulnerability management, a change management process that incorporates security, penetration testing, logging and regular log analysis, use of security products endorsed by the Australian Government, and compliance with Australian government security standards and regulations?	Policies and procedures have been established by AWS Information Security based upon the COBIT framework, ISO 27001 standards and the PCI DSS requirements. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. In addition AWS publishes a SOC 1 Type II report. Refer to the SOC 1 report for further details. The AWS Risk and Compliance whitepaper for additional details - available at http://aws.amazon.com/security. AWS customers are able to identify key controls managed by AWS. Key controls are critical to the customer's control environment and require an external attestation of the operating effectiveness of these key controls in order to comply with compliance requirements—such as the annual financial audit. For this purpose, AWS publishes a wide range of specific IT controls in its Service Organization Controls 1 (SOC 1) Type II report. The SOC 1 report, formerly the Statement on Auditing Standards (SAS) No. 70, Service Organizations report and formerly referred to as the Statement on Standards for Attestation Engagements No. 16 (SSAE 16) report, is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). The SOC 1 audit is an in-depth audit of both the design and operating effectiveness of AWS' defined control objectives and control activities (which include control objectives and control activities over the part of the infrastructure AWS manages). "Type II" refers to the fact that each of the controls described in the report are not only evaluated for adequacy of design, but are also tested for operating effectiveness by the external auditor. Because of the independence and competence of AWS' external auditor, controls identified in the report should provide customers with a high level of confidence in AWS' control environment.
	n. Technologies supporting the vendor's IT security posture. Can I have details of how the vendor's computer and network security posture is supported by direct technical controls including timely application of security patches, regularly updated antivirus software, defense in depth mechanisms to protect against unknown vulnerabilities, hardened operating systems and software applications configured with the strongest possible security settings, intrusion detection and prevention systems,	AWS provides third-party attestations, certifications, Service Organization Controls 1 (SOC 1) Type II report and other relevant compliance reports directly to our customers under NDA. AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment.

Key Area	Questions	AWS RESPONSE
	and data loss prevention mechanisms?	
	o. Auditing the vendor's IT security posture. Can I audit the vendor' implementation of security measures, including performing scans and other penetration testing of the environment provided to me? If there is justifiable reason why auditing is not possible, which reputable third party has performed audits and other vulnerability assessments? What sort of internal audits does the vendor perform, and which compliance standards and other recommended practices from organization's such as the Cloud Security Alliance are used for these assessments? Can I thoroughly review a copy of recent resulting reports?	AWS provides third-party attestations, certifications, Service Organization Controls 1 (SOC 1) Type II report and other relevant compliance reports directly to our customers under NDA. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these types of scans can be initiated by submitting a request via the AWS Vulnerability / Penetration Testing Request Form. AWS Security regularly engages independent security firms to perform external vulnerability threat assessments. The AWS SOC 1 Type 2 report provides additional details on the specific control activities executed by AWS.

Key Area	Questions	AWS RESPONSE
	p. User authentication. What identity and access management systems does the vendor support for users to log in to use Software as a Service?	AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources. AWS supports identity federation that makes it easier to manage users by maintaining their identities in a single place. AWS IAM includes support for the Security Assertion Markup Language (SAML) 2.0, an open standard used by many identity providers. This new feature enables federated single sign-on, or SSO, empowering users to log into the AWS Management Console or make programmatic calls to AWS APIs, by using assertions from a SAML-compliant identity provider, such as Shibboleth and Windows Active Directory Federation Services.
	q. Centralized control of data. What user training, policies and technical controls prevent my agency's users from using unapproved or insecure computing devices without a trusted operating environment to store or process sensitive data accessed using Software as a Service?	N/A

Key Area	Questions	AWS RESPONSE
	r. Vendor's physical security posture. Does the vendor use physical security products and devices that are endorsed by the Australian Government? How is the vendor's physical data center designed to prevent the tampering or theft of servers, infrastructure and the data stored thereon? Is the vendor's physical data center accredited by an authoritative third party?	The definition of AWS-defined logical and physical controls is documented in the SOC 1 Type II report, and the report is available for review by audit and compliance teams. AWS ISO 27001 and other certifications are also available for auditors to review. Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. Physical access is strictly controlled both at the perimeter and at building ingress points and includes, but is not limited to, professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. Physical access points to server locations are recorded by closed circuit television camera (CCTV) as defined in the AWS Data Center Physical Security Policy. Images are retained for 90 days, unless limited to 30 days by legal or contractual obligations AWS provides data center physical access and information to approved employees and contractors who have a legitimate business need for such privileges. All visitors are required to present identification and are signed in and escorted by authorized staff. See the SOC 1 Type II report for specific controls related to physical access, data center access authorization, and other related controls. Refer to ISO 27001 standard, Annex A, domain 9 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	s. Software and hardware procurement. What procurement process is used to ensure that cloud infrastructure software and hardware has been supplied by a legitimate source and has not been maliciously modified in transit?	In alignment with ISO 27001 standards, AWS Hardware assets are assigned an owner, tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standard, Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Key Area	Questions	AWS RESPONSE
Protecting Data from Unauthorized Access by the Vendor's Customers	a. Customer segregation. What assurance do I have that the virtualization and "multi-tenancy" mechanisms guarantee adequate logical and network segregation between multiple tenants, so that a malicious customer using the same physical computer as me cannot access my data?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers retain control and ownership of their data, thus it is their responsibility to choose to encrypt the data. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. VPC sessions are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security.
	b. Weakening my security posture. How would using the vendor's cloud infrastructure weaken my agency's existing network security posture? Would the vendor advertise me as one of their customers without my explicit consent, thereby assisting an adversary that is specifically targeting me?	AWS customers are considered confidential and would not advertise customer details without explicit consent. Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.
	c. Dedicated servers. Do I have some control over which physical computer runs my virtual machines? Can I pay extra to ensure that no other customer can use the same physical computer as me e.g. dedicated servers or virtual private cloud?	VPC allows customers to launch Amazon EC2 instances that are physically isolated at the host hardware level; they will run on single tenant hardware. A VPC can be created with 'dedicated' tenancy, in which case all instances launched into the VPC will utilize this feature. Alternatively, a VPC may be created with 'default' tenancy, but customers may specify 'dedicated' tenancy for particular instances launched into the VPC.
	d. Media sanitization. When I delete portions of my data, what processes are used to sanitize the storage media before it is made available to another customer, and are the processes deemed appropriate by the DSD ISM?	Customers retain ownership and control of their content and provide customers with the ability to delete their data. When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security.

Key Area	Questions	AWS RESPONSE
Protecting Data from Unauthorized Access by Rogue Vendor Employees	a. Data encryption key management. Does the vendor know the password or key used to decrypt my data, or do I encrypt and decrypt the data on my computer so the vendor only ever has encrypted data?	AWS Customers manage their own encryption unless they are utilizing AWS server side encryption service. In this case, AWS does create a unique encryption key per tenant. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security .
	b. Vetting of vendor's employees. What personnel employment checks and vetting processes does the vendor perform to ensure that employees are trustworthy?	AWS conducts criminal background checks, as permitted by applicable law, as part of pre-employment screening practices for employees commensurate with the employee's position and level of access to AWS facilities.
	c. Auditing vendor's employees. What robust identity and access management system do the vendor's employees use? What auditing process is used to log and review the actions performed by the vendor's employees?	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC 1 Type 2 report outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Overview of Security Processes whitepaper for additional details - available at http://aws.amazon.com/security .
	d. Visitors to data center. Are visitors to data centers escorted at all times, and is the name and other personal details of every visitor verified and recorded?	All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff. AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to data centers by AWS employees is routinely logged and audited.
	e. Physical tampering by vendor's employees. Is network cabling professionally installed to Australian standards or internationally acceptable standards, to help avoid the vendor's employees from accidentally connecting cables to the wrong computers, and to help readily highlight any deliberate attempts by the vendor's employees to tamper with the cabling?	Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. This includes appropriate protection for network cables. The AWS SOC 1 Type 2 report provides additional details on the specific control activities executed by AWS. Refer to ISO 27001 standard, Annex A, domain 9 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Key Area	Questions	AWS RESPONSE
	f. Vendor's subcontractors. Do the answers to these questions apply equally to all of the vendor's subcontractors?	Provisioning contractor / vendor access is managed the same for both employees and contractors, with responsibility shared across Human Resources (HR), Corporate Operations and Service Owners. Vendors are subject to the same access requirements as employees.
Handling Security Incidents	a. Timely vendor support. Is the vendor readily contactable and responsive to requests for support, and is the maximum acceptable response time captured in the SLA or simply a marketing claim that the vendor will try their best? Is the support provided locally, or from a foreign country, or from several foreign countries using an approach that follows the sun? What mechanism does the vendor use to obtain a real-time understanding of the security posture of my use of the vendor's services so that the vendor can provide support?	AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced and technical support engineers. The service helps customers of all sizes and technical abilities to successfully utilize the products and features provided by Amazon Web Services. All AWS Support tiers offer customers of AWS Infrastructure Services an unlimited number of support cases with pay-by-the-month pricing and no long-term contracts. The four tiers provide developers and businesses the flexibility to choose the support tiers that meet their specific needs.
	b. Vendor's incident response plan. Does the vendor have a security incident response plan that specifies how to detect and respond to security incidents, in a way that is similar to incident handling procedures detailed in the DSD ISM? Can I thoroughly review a copy?	The Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide 24 x 7 x 365 coverage to detect incidents and to manage the impact and resolution. AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC 1 Type 2 report provides details on the specific control activities executed by AWS. The AWS Overview of Security Processes whitepaper (available at http://aws.amazon.com/security) provides additional details.
	c. Training of vendor's employees. What qualifications, certifications and regular information security awareness training do the vendor's employees require, to know how to use the vendor's systems in a secure manner and to identify potential security incidents?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at http://aws.amazon.com/security .

Key Area	Questions	AWS RESPONSE
	d. Notification of security incidents. Will the vendor notify me via secure communications of security incidents that are more serious than an agreed threshold, especially in cases where the vendor might be liable? Will the vendor automatically notify law enforcement or other authorities, who may confiscate computing equipment used to store or process my data?	Notification of security incidents are handled on a case-by-case basis and as required by applicable law. Any notification is performed via secure communications.
	e. Extent of vendor support. How much assistance will the vendor provide me with investigations if there is a security breach such as an unauthorized disclosure of my data, or if there is a need to perform legal electronic discovery of evidence?	AWS provides infrastructure and customers manage everything else, including the operating system, the network configuration and the installed applications. Customers are responsible for responding appropriately to legal procedures involving the identification, collection, processing, analysis and production of electronic documents they store or process using AWS. Upon request, AWS may work with customers who require AWS' assistance in legal proceedings.
	f. My access to logs. How do I obtain access to time synchronized audit logs and other logs to perform a forensic investigation, and how are the logs created and stored to be suitable evidence for a court of law?	Customers retain control of their own guest operating systems, software and applications and are responsible for developing logical monitoring of the conditions of these systems. In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS CloudTrail provides a simple solution to log user activity that helps alleviate the burden of running a complex logging system. Refer to aws.amazon.com/cloudtrail for additional details. AWS Cloudwatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com.
	g. Security incident compensation. How will the vendor adequately compensate me if the vendor's actions, faulty software or hardware contributed to a security breach?	AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC 1 Type 2 report provides details on the specific control activities executed by AWS. The AWS Overview of Security Processes whitepaper (available at http://aws.amazon.com/security) provides additional details.

Key Area	Questions	AWS RESPONSE
	h. Data spills. If data that I consider is too sensitive to be stored in the cloud is accidentally placed into the cloud, referred to as a data spill, how can the spilled data be deleted using forensic sanitization techniques? Is the relevant portion of physical storage media zeroed whenever data is deleted? If not, how long does it take for deleted data to be overwritten by customers as part of normal operation, noting that clouds typically have significant spare unused storage capacity? Can the spilled data be forensically deleted from the vendor's backup media? Where else is the spilled data stored, and can it be forensically deleted?	Customers retain ownership and control of their content. All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPSec tunnels to VPC are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/security.

Appendix C: Glossary of Terms

Authentication: Authentication is the process of determining whether someone or something is, in fact, who or what it is declared to be.

Availability Zone: Amazon EC2 locations are composed of regions and Availability Zones. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same region.

DSS: The Payment Card Industry Data Security Standard (DSS) is a worldwide information security standard assembled and managed by the Payment Card Industry Security Standards Council.

EBS: Amazon Elastic Block Store (EBS) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance.

FedRAMPsm: The Federal Risk and Authorization Management Program (FedRAMPsm) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMPsm is mandatory for Federal Agency cloud deployments and service models at the low and moderate risk impact levels.

FISMA: The Federal Information Security Management Act of 2002. The act requires each federal agency to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including those provided or managed by another agency, contractor, or other source.

FIPS 140-2: The Federal Information Processing Standard (FIPS) Publication 140-2 is a US government security standard that specifies the security requirements for cryptographic modules protecting sensitive information.

GLBA: The Gramm–Leach–Bliley Act (GLB or GLBA), also known as the Financial Services Modernization Act of 1999, sets forth requirements for financial institutions with regard to, among other things, the disclosure of nonpublic customer information and the protection of threats in security and data integrity.

HIPAA: The Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires the establishment of national standards for electronic health care transactions and national identifiers for providers, health insurance plans, and employers. The Administration Simplification provisions also address the security and privacy of health data. The standards are meant to improve the efficiency and effectiveness of the nation's health care system by encouraging the widespread use of electronic data interchange in the U.S. health care system.

Hypervisor: A hypervisor, also called Virtual Machine Monitor (VMM), is software/hardware platform virtualization software that allows multiple operating systems to run on a host computer concurrently.

IAM: AWS Identity and Access Management (IAM) enables a customer to create multiple Users and manage the permissions for each of these Users within their AWS Account.

ITAR: International Traffic in Arms Regulations (ITAR) is a set of United States government regulations that control the export and import of defense-related articles and services on the United States Munitions List (USML). Government agencies and contractors must comply with ITAR and restrict access to protected data.



ISAE 3402: The International Standards for Assurance Engagements No. 3402 (ISAE 3402) is the international standard on assurance engagements. It was put forth by the International Auditing and Assurance Standards Board (IAASB), a standard-setting board within the International Federation of Accountants (IFAC). ISAE 3402 is now the new globally recognized standard for assurance reporting on service organizations.

ISO 9001: AWS' ISO 9001 certification directly supports customers who develop, migrate and operate their quality-controlled IT systems in the AWS cloud. Customers can leverage AWS' compliance reports as evidence for their own ISO 9001 programs and industry-specific quality programs, such as GxP in life sciences, ISO 13485 in medical devices, AS9100 in aerospace, and ISO/TS 16949 in automotive. AWS customers who don't have quality system requirements will still benefit from the additional assurance and transparency that an ISO 9001 certification provides.

ISO 27001: ISO/IEC 27001 is an Information Security Management System (ISMS) standard published by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). ISO 27001 formally specifies a management system that is intended to bring information security under explicit management control. Being a formal specification means that it mandates specific requirements. Organizations that claim to have adopted ISO/IEC 27001 can therefore be audited and certified compliant with the standard.

NIST: National Institute of Standards and Technology. This agency sets detailed security standards as needed by industry or government programs. Compliance with FISMA requires agencies to adhere to NIST standards.

Object: The fundamental entities stored in Amazon S3. Objects consist of object data and metadata. The data portion is opaque to Amazon S3. The metadata is a set of name-value pairs that describe the object. These include some default metadata such as the date last modified and standard HTTP metadata such as Content-Type. The developer can also specify custom metadata at the time the Object is stored.

PCI: Refers to the Payment Card Industry Security Standards Council, an independent council originally formed by American Express, Discover Financial Services, JCB, MasterCard Worldwide and Visa International, with the goal of managing the ongoing evolution of the Payment Card Industry Data Security Standard.

QSA: The Payment Card Industry (PCI) Qualified Security Assessor (QSA) designation is conferred by the PCI Security Standards Council to those individuals that meet specific qualification requirements and are authorized to perform PCI compliance assessments.

SAS 70: Statement on Auditing Standards No. 70: Service Organizations is an auditing statement issued by the Auditing Standards Board of the American Institute of Certified Public Accountants (AICPA). SAS 70 provides guidance to service auditors when assessing the internal controls of a service organization (such as AWS) and issuing a service auditor's report. SAS 70 also provides guidance to auditors of financial statements of an entity that uses one or more service organizations. The SAS 70 report has been replaced by the Service Organization Controls 1 report.

Service: Software or computing ability provided across a network (e.g., EC2, S3, VPC, etc.).

Service Level Agreement (SLA): A service level agreement is a part of a service contract where the level of service is formally defined. The SLA is used to refer to the contracted delivery time (of the service) or performance.

SOC 1: Service Organization Controls 1 (SOC 1) Type II report, formerly the Statement on Auditing Standards (SAS) No. 70, Service Organizations report (formerly referred to as the SSAE 16 report), is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). The



international standard is referenced as the International Standards for Assurance Engagements No. 3402 (ISAE 3402).

SSAE 16 [deprecated]: The Statement on Standards for Attestation Engagements No. 16 (SSAE 16) is an attestation standard published by the Auditing Standards Board (ASB) of the American Institute of Certified Public Accountants (AICPA). The standard addresses engagements undertaken by a service auditor for reporting on controls at organizations that provide services to user entities, for which a service organization's controls are likely to be relevant to a user entities internal control over financial reporting (ICFR). SSAE 16 effectively replaces Statement on Auditing Standards No. 70 (SAS 70) for service auditor's reporting periods ending on or after June 15, 2011.

SOC 2: Service Organization Controls 2 (SOC 2) reports are intended to meet the needs of a broad range of users that need to understand internal control at a service organization as it relates to security, availability, processing integrity, confidentiality and privacy. These reports are performed using the AICPA Guide: Reporting on Controls at a Service Organizations Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy and are intended for use by stakeholders (e.g., customers, regulators, business partners, suppliers, directors) of the service organization that have a thorough understanding of the service organization and its internal controls.

SOC 3: Service Organization Controls 3 (SOC 3) reports are designed to meet the needs of users who want assurance on the controls at a service organization related to security, availability, processing integrity, confidentiality, or privacy but do not have the need for or the knowledge necessary to make effective use of a SOC 2 Report. These reports are prepared using the AICPA/Canadian Institute of Chartered Accountants (CICA) Trust Services Principles, Criteria, and Illustrations for Security, Availability, Processing Integrity, Confidentiality, and Privacy. Because they are general use reports, SOC 3 Reports can be freely distributed or posted on a website as a seal.

Virtual Instance: Once an AMI has been launched, the resulting running system is referred to as an instance. All instances based on the same AMI start out identical and any information on them is lost when the instances are terminated or fail.

Version History

January 2016

- Added GxP Compliance Program
- Twelfth region added (Asia Pacific - Seoul)

December 2015

- Updates to certifications and third-party attestations summaries
- Added ISO 27017 certification
- Added ISO 27018 certification
- Eleventh region added (China - Beijing)

November 2015

- Update to CSA v3.0.1

August 2015

- Updates to in-scope services for PCI 3.1
- Updates to regions in-scope for PCI 3.1

May 2015

- Tenth region added (EU - Frankfurt)
- Updates to in-scope services for SOC 3
- SSAE 16 language deprecated

Apr 2015

- Updates to in-scope services for: FedRAMPsm, HIPAA, SOC 1, ISO 27001, ISO 9001

Feb 2015

- Updates to FIPS 140-2 VPN endpoints and SSL-terminating load balancers
- Updates to PCI DSS verbiage

Dec 2014

- Updates to certifications and third-party attestations summaries

Nov 2013 version

- Edits to IPsec tunnel encryption verbiage

Jun 2013 version

- Updates to certifications and third-party attestations summaries
- Updates to Appendix C: Glossary of Terms
- Minor changes to formatting

Jan 2013 version

- Edits to certifications and third-party attestations summaries

Nov 2012 version

- Edits to content and updated certification scope
- Added reference to the SOC 2 and MPAA

Jul 2012 version

- Edits to content and updated certification scope
- Addition of the CSA Consensus Assessments Initiative Questionnaire (Appendix A)

Jan 2012 version

- Minor edits to content based on updated certification scope



- Minor grammatical edits

Dec 2011 version

- Change to Certifications and Third-party Attestation section to reflect SOC 1/SSAE 16, FISMA Moderate, International Traffic in Arms Regulations, and FIPS 140-2
- Addition of S3 Server Side Encryption
- Added additional cloud computing issue topics

May 2011 version

- Initial release

Notices

© 2010-2016 Amazon.com, Inc., or its affiliates. This document is provided for informational purposes only. It represents AWS' current product offerings as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS' products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.



Amazon Web Services: Overview of Security Processes

August 2015

(Please consult <http://aws.amazon.com/security/> for the latest version of this paper)

Table of Contents

Introduction	5
Shared Security Responsibility Model	5
AWS Security Responsibilities	6
Customer Security Responsibilities	6
AWS Global Infrastructure Security	7
AWS Compliance Program	7
Physical and Environmental Security	8
Fire Detection and Suppression	8
Power	8
Climate and Temperature	8
Management	8
Storage Device Decommissioning	8
Business Continuity Management	9
Availability	9
Incident Response	9
Company-Wide Executive Review	9
Communication	9
Network Security	10
Secure Network Architecture	10
Secure Access Points	10
Transmission Protection	10
Amazon Corporate Segregation	10
Fault-Tolerant Design	11
Network Monitoring and Protection	12
AWS Access	14
Account Review and Audit	14
Background Checks	14
Credentials Policy	14
Secure Design Principles	14
Change Management	15
Software	15
Infrastructure	15
AWS Account Security Features	16
AWS Credentials	16

Passwords	17
AWS Multi-Factor Authentication (AWS MFA)	17
Access Keys	18
Key Pairs.....	18
X.509 Certificates	18
Individual User Accounts.....	19
Secure HTTPS Access Points.....	19
Security Logs	19
AWS Trusted Advisor Security Checks	20
AWS Service-Specific Security	20
Compute Services.....	20
Amazon Elastic Compute Cloud (Amazon EC2) Security.....	20
Auto Scaling Security	24
Networking Services.....	25
Amazon Elastic Load Balancing Security	25
Amazon Virtual Private Cloud (Amazon VPC) Security	26
Amazon Route 53 Security.....	31
Amazon CloudFront Security	32
AWS Direct Connect Security.....	34
Storage Services	35
Amazon Simple Storage Service (Amazon S3) Security	35
AWS Glacier Security.....	37
AWS Storage Gateway Security	38
AWS Import/Export Security.....	39
Database Services	41
Amazon DynamoDB Security	41
Amazon Relational Database Service (Amazon RDS) Security.....	42
Amazon Redshift Security	46
Amazon ElastiCache Security	48
Application Services	50
Amazon CloudSearch Security	50
Amazon Simple Queue Service (Amazon SQS) Security.....	51
Amazon Simple Notification Service (Amazon SNS) Security	51
Amazon Simple Workflow Service (Amazon SWF) Security.....	52
Amazon Simple Email Service (Amazon SES) Security	52
Amazon Elastic Transcoder Service Security.....	53
Amazon AppStream Security	54
Analytics Services	55

Amazon Elastic MapReduce (Amazon EMR) Security	55
Amazon Kinesis Security	56
AWS Data Pipeline Security	56
Deployment and Management Services	57
AWS Identity and Access Management (AWS IAM)	57
Amazon CloudWatch Security	58
AWS CloudHSM Security	59
AWS CloudTrail Security	60
Mobile Services	60
Amazon Cognito	60
Amazon Mobile Analytics	62
Applications	62
Amazon WorkSpaces	62
Amazon WorkDocs	63
Appendix – Glossary of Terms	65

Introduction

Amazon Web Services (AWS) delivers a scalable cloud computing platform with high availability and dependability, providing the tools that enable customers to run a wide range of applications. Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence. This document is intended to answer questions such as, "How does AWS help me protect my data?" Specifically, AWS physical and operational security processes are described for the network and server infrastructure under AWS's management, as well as service-specific security implementations.

Shared Security Responsibility Model

Before we go into the details of how AWS secures its resources, we should talk about how security in the cloud is slightly different than security in your on-premises data centers. When you move computer systems and data to the cloud, security responsibilities become shared between you and your cloud service provider. In this case, AWS is responsible for securing the underlying infrastructure that supports the cloud, and you're responsible for anything you put on the cloud or connect to the cloud. This shared security responsibility model can reduce your operational burden in many ways, and in some cases may even improve your default security posture without additional action on your part.



Figure 1: AWS Shared Security Responsibility Model

The amount of security configuration work you have to do varies depending on which services you select and how sensitive your data is. However, there are certain security features—such as individual user accounts and credentials, SSL/TLS for data transmissions, and user activity logging—that you should configure no matter which AWS service you use. For more information about these security features, see the "AWS Account Security Features" section below.

AWS Security Responsibilities

Amazon Web Services is responsible for protecting the global infrastructure that runs all of the services offered in the AWS cloud. This infrastructure is comprised of the hardware, software, networking, and facilities that run AWS services. Protecting this infrastructure is AWS's number one priority, and while you can't visit our data centers or offices to see this protection firsthand, we provide several reports from third-party auditors who have verified our compliance with a variety of computer security standards and regulations (for more information, visit aws.amazon.com/compliance).

Note that in addition to protecting this global infrastructure, AWS is responsible for the security configuration of its products that are considered managed services. Examples of these types of services include Amazon DynamoDB, Amazon RDS, Amazon Redshift, Amazon Elastic MapReduce, Amazon WorkSpaces, and several other services. These services provide the scalability and flexibility of cloud-based resources with the additional benefit of being managed. For these services, AWS will handle basic security tasks like guest operating system (OS) and database patching, firewall configuration, and disaster recovery. For most of these managed services, all you have to do is configure logical access controls for the resources and protect your account credentials. A few of them may require additional tasks, such as setting up database user accounts, but overall the security configuration work is performed by the service.

Customer Security Responsibilities

With the AWS cloud, you can provision virtual servers, storage, databases, and desktops in minutes instead of weeks. You can also use cloud-based analytics and workflow tools to process your data as you need it, and then store it in your own data centers or in the cloud. Which AWS services you use will determine how much configuration work you have to perform as part of your security responsibilities.

AWS products that fall into the well-understood category of Infrastructure as a Service (IaaS)—such as Amazon EC2, Amazon VPC, and Amazon S3—are completely under your control and require you to perform all of the necessary security configuration and management tasks. For example, for EC2 instances, you're responsible for management of the guest OS (including updates and security patches), any application software or utilities you install on the instances, and the configuration of the AWS-provided firewall (called a security group) on each instance. These are basically the same security tasks that you're used to performing no matter where your servers are located.

AWS managed services like Amazon RDS or Amazon Redshift provide all of the resources you need in order to perform a specific task—but without the configuration work that can come with them. With managed services, you don't have to worry about launching and maintaining instances, patching the guest OS or database, or replicating databases—AWS handles that for you. But as with all services, you should protect your AWS Account credentials and set up individual user accounts with Amazon Identity and Access Management (IAM) so that each of your users has their own credentials and you can implement segregation of duties. We also recommend using multi-factor authentication (MFA) with each account, requiring the use of SSL/TLS to communicate with your AWS resources, and setting up API/user activity logging with AWS CloudTrail. For more information about additional measures you can take, refer to the [AWS Security Best Practices whitepaper](#) and recommended reading on the [AWS Security Resources](#) webpage.

AWS Global Infrastructure Security

AWS operates the global cloud infrastructure that you use to provision a variety of basic computing resources such as processing and storage. The AWS global infrastructure includes the facilities, network, hardware, and operational software (e.g., host OS, virtualization software, etc.) that support the provisioning and use of these resources. The AWS global infrastructure is designed and managed according to security best practices as well as a variety of security compliance standards. As an AWS customer, you can be assured that you're building web architectures on top of some of the most secure computing infrastructure in the world.

AWS Compliance Program

Amazon Web Services Compliance enables customers to understand the robust controls in place at AWS to maintain security and data protection in the cloud. As systems are built on top of [AWS cloud infrastructure](#), compliance responsibilities will be [shared](#). By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, AWS [Compliance enablers](#) build on traditional programs; helping customers to establish and operate in an AWS security control environment. The IT infrastructure that AWS provides to its customers is designed and managed in alignment with security best practices and a variety of IT security standards, including:

- SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70)
- SOC 2
- SOC 3
- FISMA, DIACAP, and FedRAMP
- DOD CSM Levels 1-5
- PCI DSS Level 1
- ISO 9001 / ISO 27001
- ITAR
- FIPS 140-2
- MTCS Level 3

In addition, the flexibility and control that the AWS platform provides allows customers to deploy solutions that meet several industry-specific standards, including:

- Criminal Justice Information Services (CJIS)
- Cloud Security Alliance (CSA)
- Family Educational Rights and Privacy Act (FERPA)
- Health Insurance Portability and Accountability Act (HIPAA)
- Motion Picture Association of America (MPAA)

AWS provides a wide range of information regarding its IT control environment to customers through white papers, reports, certifications, accreditations, and other third-party attestations. More information is available in the Risk and Compliance whitepaper available on the website: <http://aws.amazon.com/compliance/>.



Physical and Environmental Security

AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches. Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to data centers by AWS employees is logged and audited routinely.

Fire Detection and Suppression

Automatic fire detection and suppression equipment has been installed to reduce risk. The fire detection system utilizes smoke detection sensors in all data center environments, mechanical and electrical infrastructure spaces, chiller rooms and generator equipment rooms. These areas are protected by either wet-pipe, double-interlocked pre-action, or gaseous sprinkler systems.

Power

The data center electrical power systems are designed to be fully redundant and maintainable without impact to operations, 24 hours a day, and seven days a week. Uninterruptible Power Supply (UPS) units provide back-up power in the event of an electrical failure for critical and essential loads in the facility. Data centers use generators to provide back-up power for the entire facility.

Climate and Temperature

Climate control is required to maintain a constant operating temperature for servers and other hardware, which prevents overheating and reduces the possibility of service outages. Data centers are conditioned to maintain atmospheric conditions at optimal levels. Personnel and systems monitor and control temperature and humidity at appropriate levels.

Management

AWS monitors electrical, mechanical, and life support systems and equipment so that any issues are immediately identified. Preventative maintenance is performed to maintain the continued operability of equipment.

Storage Device Decommissioning

When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.



Business Continuity Management

Amazon's infrastructure has a high level of availability and provides customers the features to deploy a resilient IT architecture. AWS has designed its systems to tolerate system or hardware failures with minimal customer impact. Data center Business Continuity Management at AWS is under the direction of the Amazon Infrastructure Group.

Availability

Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.

AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by Region). In addition to discrete uninterruptable power supply (UPS) and onsite backup generation facilities, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability zones are all redundantly connected to multiple tier-1 transit providers.

You should architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure modes, including natural disasters or system failures.

Incident Response

The Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide 24x7x365 coverage to detect incidents and to manage the impact and resolution.

Company-Wide Executive Review

Amazon's Internal Audit group has recently reviewed the AWS services resiliency plans, which are also periodically reviewed by members of the Senior Executive management team and the Audit Committee of the Board of Directors.

Communication

AWS has implemented various methods of internal communication at a global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employees; regular management meetings for updates on business performance and other matters; and electronics means such as video conferencing, electronic mail messages, and the posting of information via the Amazon intranet.

AWS has also implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact the customer experience. A "[Service Health Dashboard](#)" is available and maintained by the customer support team to alert customers to any issues that may be of broad impact. The "AWS [Security Center](#)" is available to provide you with security and compliance details about AWS. You can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.



Network Security

The AWS network has been architected to permit you to select the level of security and resiliency appropriate for your workload. To enable you to build geographically dispersed, fault-tolerant web architectures with cloud resources, AWS has implemented a world-class network infrastructure that is carefully monitored and managed.

Secure Network Architecture

Network devices, including firewall and other boundary devices, are in place to monitor and control communications at the external boundary of the network and at key internal boundaries within the network. These boundary devices employ rule sets, access control lists (ACL), and configurations to enforce the flow of information to specific information system services.

ACLs, or traffic flow policies, are established on each managed interface, which manage and enforce the flow of traffic. ACL policies are approved by Amazon Information Security. These policies are automatically pushed using AWS's ACL-Manage tool, to help ensure these managed interfaces enforce the most up-to-date ACLs.

Secure Access Points

AWS has strategically placed a limited number of access points to the cloud to allow for a more comprehensive monitoring of inbound and outbound communications and network traffic. These customer access points are called API endpoints, and they allow secure HTTP access (HTTPS), which allows you to establish a secure communication session with your storage or compute instances within AWS. To support customers with FIPS cryptographic requirements, the SSL-terminating load balancers in AWS GovCloud (US) are FIPS 140-2-compliant.

In addition, AWS has implemented network devices that are dedicated to managing interfacing communications with Internet service providers (ISPs). AWS employs a redundant connection to more than one communication service at each Internet-facing edge of the AWS network. These connections each have dedicated network devices.

Transmission Protection

You can connect to an AWS access point via HTTP or HTTPS using Secure Sockets Layer (SSL), a cryptographic protocol that is designed to protect against eavesdropping, tampering, and message forgery.

For customers who require additional layers of network security, AWS offers the Amazon Virtual Private Cloud (VPC), which provides a private subnet within the AWS cloud, and the ability to use an IPsec Virtual Private Network (VPN) device to provide an encrypted tunnel between the Amazon VPC and your data center. For more information about VPC configuration options, refer to the [Amazon Virtual Private Cloud \(Amazon VPC\) Security](#) section below.

Amazon Corporate Segregation

Logically, the AWS Production network is segregated from the Amazon Corporate network by means of a complex set of network security / segregation devices. AWS developers and administrators on the corporate network who need to access AWS cloud components in order to maintain them must explicitly request access through the AWS ticketing system. All requests are reviewed and approved by the applicable service owner.

Approved AWS personnel then connect to the AWS network through a bastion host that restricts access to network devices and other cloud components, logging all activity for security review. Access to bastion hosts require SSH public-key authentication for all user accounts on the host. For more information on AWS developer and administrator logical access, see *AWS Access* below.



Fault-Tolerant Design

Amazon's infrastructure has a high level of availability and provides you with the capability to deploy a resilient IT architecture. AWS has designed its systems to tolerate system or hardware failures with minimal customer impact.

Data centers are built in clusters in various global *regions*. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.

AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple availability zones within each region. Each availability zone is designed as an independent failure zone. This means that availability zones are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by region). In addition to utilizing discrete uninterruptable power supply (UPS) and onsite backup generators, they are each fed via different grids from independent utilities to further reduce single points of failure. Availability zones are all redundantly connected to multiple tier-1 transit providers.

You should architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures. However, you should be aware of location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between regions unless proactively done so by the customer, thus allowing customers with these types of data placement and privacy requirements the ability to establish compliant environments. It should be noted that all communications between regions is across public Internet infrastructure; therefore, appropriate encryption methods should be used to protect sensitive data.

As of this writing, there are eleven regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), South America (Sao Paulo), and China (Beijing).

AWS GovCloud (US) is an isolated AWS Region designed to allow US government agencies and customers to move workloads into the cloud by helping them meet certain regulatory and compliance requirements. The AWS GovCloud (US) framework allows US government agencies and their contractors to comply with U.S. International Traffic in Arms Regulations (ITAR) regulations as well as the Federal Risk and Authorization Management Program (FedRAMP) requirements. AWS GovCloud (US) has received an Agency Authorization to Operate (ATO) from the US Department of Health and Human Services (HHS) utilizing a FedRAMP accredited Third Party Assessment Organization (3PAO) for several AWS services.

The AWS GovCloud (US) Region provides the same fault-tolerant design as other regions, with two Availability Zones. In addition, the AWS GovCloud (US) region is a mandatory AWS Virtual Private Cloud (VPC) service by default to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses. More information about GovCloud is available on the AWS website: <http://aws.amazon.com/govcloud-us/>






Figure 2: Regions and Availability Zones

Note that the number of Availability Zones may change.

Network Monitoring and Protection

AWS utilizes a wide variety of automated monitoring systems to provide a high level of service performance and availability. AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools have the ability to set custom performance metrics thresholds for unusual activity.

Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Documentation is maintained to aid and inform operations personnel in handling incidents or issues. If the resolution of an issue requires collaboration, a conferencing system is used which supports communication and logging capabilities. Trained call leaders facilitate communication and progress during the handling of operational issues that require collaboration. Post-mortems are convened after any significant operational issue, regardless of external impact, and Cause of Error (COE) documents are drafted so the root cause is captured and preventative actions are taken in the future. Implementation of the preventative measures is tracked during weekly operations meetings.

AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant telecommunication providers at each region as well as additional capacity protect against the possibility of DoS attacks.

The AWS network provides significant protection against traditional network security issues, and you can implement further protection. The following are a few examples:

- **Distributed Denial Of Service (DDoS) Attacks.** AWS API endpoints are hosted on large, Internet-scale, world-class infrastructure that benefits from the same engineering expertise that has built Amazon into the world's largest online retailer. Proprietary DDoS mitigation techniques are used. Additionally, AWS's networks are multi-homed across a number of providers to achieve Internet access diversity.
- **Man in the Middle (MITM) Attacks.** All of the AWS APIs are available via SSL-protected endpoints which provide server authentication. Amazon EC2 AMIs automatically generate new SSH host certificates on first boot and log them to the instance's console. You can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. We encourage you to use SSL for all of your interactions with AWS.
- **IP Spoofing.** Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.
- **Port Scanning.** Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on our website at: <http://aws.amazon.com/contact-us/report-abuse/>. When unauthorized port scanning is detected by AWS, it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by you. Your strict management of security groups can further mitigate the threat of port scans. If you configure the security group to allow traffic from any source to a specific port, then that specific port will be vulnerable to a port scan. In these cases, you must use appropriate security measures to protect listening services that may be essential to their application from being discovered by an unauthorized port scan. For example, a web server must clearly have port 80 (HTTP) open to the world, and the administrator of this server is responsible for the security of the HTTP server software, such as Apache. You may request permission to conduct vulnerability scans as required to meet your specific compliance requirements. These scans must be limited to your own instances and must not violate the AWS Acceptable Use Policy. Advanced approval for these types of scans can be initiated by submitting a request via the website at: <https://aws-portal.amazon.com/gp/aws/html-forms-controller/contactus/AWSecurityPenTestRequest>
- **Packet sniffing by other tenants.** It is not possible for a virtual instance running in promiscuous mode to receive or "sniff" traffic that is intended for a different virtual instance. While you can place your interfaces into promiscuous mode, the hypervisor will not deliver any traffic to them that is not addressed to them. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic. Attacks such as ARP cache poisoning do not work within Amazon EC2 and Amazon VPC. While Amazon EC2 does provide ample protection against one customer inadvertently or maliciously attempting to view another's data, as a standard practice you should encrypt sensitive traffic.

In addition to monitoring, regular vulnerability scans are performed on the host operating system, web application, and databases in the AWS environment using a variety of tools. Also, AWS Security teams subscribe to newsfeeds for applicable vendor flaws and proactively monitor vendors' websites and other relevant outlets for new patches. AWS customers also have the ability to report issues to AWS via the AWS Vulnerability Reporting website at:

<http://aws.amazon.com/security/vulnerability-reporting/>

AWS Access

The AWS Production network is segregated from the Amazon Corporate network and requires a separate set of credentials for logical access. The Amazon Corporate network relies on user IDs, passwords, and Kerberos, while the AWS Production network requires SSH public-key authentication through a bastion host.

AWS developers and administrators on the Amazon Corporate network who need to access AWS cloud components must explicitly request access through the AWS access management system. All requests are reviewed and approved by the appropriate owner or manager.

Account Review and Audit

Accounts are reviewed every 90 days; explicit re-approval is required or access to the resource is automatically revoked. Access is also automatically revoked when an employee's record is terminated in Amazon's Human Resources system. Windows and UNIX accounts are disabled and Amazon's permission management system removes the user from all systems.

Requests for changes in access are captured in the Amazon permissions management tool audit log. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked.

Background Checks

AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS conducts criminal background checks, as permitted by law, as part of pre-employment screening practices for employees and commensurate with the employee's position and level of access. The policies also identify functional responsibilities for the administration of logical access and security.

Credentials Policy

AWS Security has established a credentials policy with required configurations and expiration intervals. Passwords must be complex and are forced to be changed every 90 days.

Secure Design Principles

AWS's development process follows secure software development best practices, which include formal design reviews by the AWS Security Team, threat modeling, and completion of a risk assessment. Static code analysis tools are run as a part of the standard build process, and all deployed software undergoes recurring penetration testing performed by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through launch to ongoing operations.



Change Management

Routine, emergency, and configuration changes to existing AWS infrastructure are authorized, logged, tested, approved, and documented in accordance with industry norms for similar systems. Updates to AWS's infrastructure are done to minimize any impact on the customer and their use of the services. AWS will communicate with customers, either via email, or through the AWS Service Health Dashboard (<http://status.aws.amazon.com/>) when service use is likely to be adversely affected.

Software

AWS applies a systematic approach to managing change so that changes to customer-impacting services are thoroughly reviewed, tested, approved, and well-communicated. The AWS change management process is designed to avoid unintended service disruptions and to maintain the integrity of service to the customer. Changes deployed into production environments are:

- Reviewed: Peer reviews of the technical aspects of a change are required.
- Tested: Changes being applied are tested to help ensure they will behave as expected and not adversely impact performance.
- Approved: All changes must be authorized in order to provide appropriate oversight and understanding of business impact.

Changes are typically pushed into production in a phased deployment starting with lowest impact areas. Deployments are tested on a single system and closely monitored so impacts can be evaluated. Service owners have a number of configurable metrics that measure the health of the service's upstream dependencies. These metrics are closely monitored with thresholds and alarming in place. Rollback procedures are documented in the Change Management (CM) ticket.

When possible, changes are scheduled during regular change windows. Emergency changes to production systems that require deviations from standard change management procedures are associated with an incident and are logged and approved as appropriate.

Periodically, AWS performs self-audits of changes to key services to monitor quality, maintain high standards, and facilitate continuous improvement of the change management process. Any exceptions are analyzed to determine the root cause, and appropriate actions are taken to bring the change into compliance or roll back the change if necessary. Actions are then taken to address and remediate the process or people issue.

Infrastructure

Amazon's Corporate Applications team develops and manages software to automate IT processes for UNIX/Linux hosts in the areas of third-party software delivery, internally developed software, and configuration management. The Infrastructure team maintains and operates a UNIX/Linux configuration management framework to address hardware scalability, availability, auditing, and security management. By centrally managing hosts through the use of automated processes that manage change, Amazon is able to achieve its goals of high availability, repeatability, scalability, security, and disaster recovery. Systems and network engineers monitor the status of these automated tools on a continuous basis, reviewing reports to respond to hosts that fail to obtain or update their configuration and software.



Internally developed configuration management software is installed when new hardware is provisioned. These tools are run on all UNIX hosts to validate that they are configured and that software is installed in compliance with standards determined by the role assigned to the host. This configuration management software also helps to regularly update packages that are already installed on the host. Only approved personnel enabled through the permissions service may log in to the central configuration management servers.

AWS Account Security Features

AWS provides a variety of tools and features that you can use to keep your AWS Account and resources safe from unauthorized use. This includes credentials for access control, HTTPS endpoints for encrypted data transmission, the creation of separate IAM user accounts, user activity logging for security monitoring, and Trusted Advisor security checks. You can take advantage of all of these security tools no matter which AWS services you select.

AWS Credentials

To help ensure that only authorized users and processes access your AWS Account and resources, AWS uses several types of credentials for authentication. These include passwords, cryptographic keys, digital signatures, and certificates. We also provide the option of requiring multi-factor authentication (MFA) to log into your AWS Account or IAM user accounts. The following table highlights the various AWS credentials and their uses.

Credential Type	Use	Description
Passwords	AWS root account or IAM user account login to the AWS Management Console	A string of characters used to log into your AWS account or IAM account. AWS passwords must be a minimum of 6 characters and may be up to 128 characters.
Multi-Factor Authentication (MFA)	AWS root account or IAM user account login to the AWS Management Console	A six-digit single-use code that is required in addition to your password to log in to your AWS Account or IAM user account.
Access Keys	Digitally signed requests to AWS APIs (using the AWS SDK, CLI, or REST/Query APIs)	Includes an access key ID and a secret access key. You use access keys to digitally sign programmatic requests that you make to AWS.
Key Pairs	- SSH login to EC2 instances - CloudFront signed URLs	A key pair is required to connect to an EC2 instance launched from a public AMI. The keys that Amazon EC2 uses are 1024-bit SSH-2 RSA keys. You can have a key pair generated automatically for you when you launch the instance or you can upload your own.
X.509 Certificates	- Digitally signed SOAP requests to AWS APIs - SSL server certificates for HTTPS	X.509 certificates are only used to sign SOAP-based requests (currently used only with Amazon S3). You can have AWS create an X.509 certificate and private key that you can download, or you can upload your own certificate by using the Security Credentials page.

You can download a Credential Report for your account at any time from the Security Credentials page. This report lists all of your account's users and the status of their credentials—whether they use a password, whether their password

expires and must be changed regularly, the last time they changed their password, the last time they rotated their access keys, and whether they have MFA enabled.

For security reasons, if your credentials have been lost or forgotten, you cannot recover them or re-download them. However, you can create new credentials and then disable or delete the old set of credentials.

In fact, AWS recommends that you change (rotate) your access keys and certificates on a regular basis. To help you do this without potential impact to your application's availability, AWS supports multiple concurrent access keys and certificates. With this feature, you can rotate keys and certificates into and out of operation on a regular basis without any downtime to your application. This can help to mitigate risk from lost or compromised access keys or certificates. The AWS IAM API enables you to rotate the access keys of your AWS Account as well as for IAM user accounts.

Passwords

Passwords are required to access your AWS Account, individual IAM user accounts, AWS Discussion Forums, and the AWS Support Center. You specify the password when you first create the account, and you can change it at any time by going to the Security Credentials page. AWS passwords can be up to 128 characters long and contain special characters, so we encourage you to create a strong password that cannot be easily guessed.

You can set a password policy for your IAM user accounts to ensure that strong passwords are used and that they are changed often. A password policy is a set of rules that define the type of password an IAM user can set. For more information about password policies, go to [Managing Passwords](#) in Using IAM.

AWS Multi-Factor Authentication (AWS MFA)

AWS Multi-Factor Authentication (AWS MFA) is an additional layer of security for accessing AWS services. When you enable this optional feature, you will need to provide a six-digit single-use code in addition to your standard user name and password credentials before access is granted to your AWS Account settings or AWS services and resources. You get this single-use code from an authentication device that you keep in your physical possession. This is called multi-factor authentication because more than one authentication factor is checked before access is granted: a password (something you know) and the precise code from your authentication device (something you have). You can enable MFA devices for your AWS Account as well as for the users you have created under your AWS Account with AWS IAM. In addition, you add MFA protection for access across AWS Accounts, for when you want to allow a user you've created under one AWS Account to use an IAM role to access resources under another AWS Account. You can require the user to use MFA before assuming the role as an additional layer of security.

AWS MFA supports the use of both hardware tokens and virtual MFA devices. Virtual MFA devices use the same protocols as the physical MFA devices, but can run on any mobile hardware device, including a smartphone. A virtual MFA device uses a software application that generates six-digit authentication codes that are compatible with the Time-Based One-Time Password (TOTP) standard, as described in [RFC 6238](#). Most virtual MFA applications allow you to host more than one virtual MFA device, which makes them more convenient than hardware MFA devices. However, you should be aware that because a virtual MFA might be run on a less secure device such as a smartphone, a virtual MFA might not provide the same level of security as a hardware MFA device.

You can also enforce MFA authentication for AWS service APIs in order to provide an extra layer of protection over powerful or privileged actions such as terminating Amazon EC2 instances or reading sensitive data stored in Amazon S3. You do this by adding an MFA-authentication requirement to an IAM access policy. You can attach these access policies to IAM users, IAM groups, or resources that support Access Control Lists (ACLs) like Amazon S3 buckets, SQS queues, and SNS topics.



It is easy to obtain hardware tokens from a participating third-party provider or virtual MFA applications from an AppStore and to set it up for use via the AWS website. More information about AWS MFA is available on the AWS website: <http://aws.amazon.com/mfa/>

Access Keys

AWS requires that all API requests be signed—that is, they must include a digital signature that AWS can use to verify the identity of the requestor. You calculate the digital signature using a cryptographic hash function. The input to the hash function in this case includes the text of your request and your secret access key. If you use any of the AWS SDKs to generate requests, the digital signature calculation is done for you; otherwise, you can have your application calculate it and include it in your REST or Query requests by following the directions [in our documentation](#).

Not only does the signing process help protect message integrity by preventing tampering with the request while it is in transit, it also helps protect against potential replay attacks. A request must reach AWS within 15 minutes of the time stamp in the request. Otherwise, AWS denies the request.

The most recent version of the digital signature calculation process is Signature Version 4, which calculates the signature using the HMAC-SHA256 protocol. Version 4 provides an additional measure of protection over previous versions by requiring that you sign the message using a key that is derived from your secret access key rather than using the secret access key itself. In addition, you derive the signing key based on *credential scope*, which facilitates cryptographic isolation of the signing key.

Because access keys can be misused if they fall into the wrong hands, we encourage you to save them in a safe place and not embed them in your code. For customers with large fleets of elastically scaling EC2 instances, the use of IAM roles can be a more secure and convenient way to manage the distribution of access keys. IAM roles provide temporary credentials, which not only get automatically loaded to the target instance, but are also automatically rotated multiple times a day.

Key Pairs

Amazon EC2 instances created from a public AMI use a public/private key pair rather than a password for signing in via Secure Shell (SSH). The public key is embedded in your instance, and you use the private key to sign in securely without a password. After you create your own AMIs, you can choose other mechanisms to securely log in to your new instances.

You can have a key pair generated automatically for you when you launch the instance or you can upload your own. Save the private key in a safe place on your system, and record the location where you saved it.

For Amazon CloudFront, you use key pairs to create signed URLs for private content, such as when you want to distribute restricted content that someone paid for. You create Amazon CloudFront key pairs by using the Security Credentials page. CloudFront key pairs can be created only by the root account and cannot be created by IAM users.

X.509 Certificates

X.509 certificates are used to sign SOAP-based requests. X.509 certificates contain a public key and additional metadata (like an expiration date that AWS verifies when you upload the certificate), and is associated with a private key. When you create a request, you create a digital signature with your private key and then include that signature in the request, along with your certificate. AWS verifies that you're the sender by decrypting the signature with the public key that is in your certificate. AWS also verifies that the certificate you sent matches the certificate that you uploaded to AWS.

For your AWS Account, you can have AWS create an X.509 certificate and private key that you can download, or you can upload your own certificate by using the Security Credentials page. For IAM users, you must create the X.509 certificate (signing certificate) by using third-party software. In contrast with root account credentials, AWS cannot create an X.509 certificate for IAM users. After you create the certificate, you attach it to an IAM user by using IAM.

In addition to SOAP requests, X.509 certificates are used as SSL/TLS server certificates for customers who want to use HTTPS to encrypt their transmissions. To use them for HTTPS, you can use an open-source tool like OpenSSL to create a unique private key. You'll need the private key to create the Certificate Signing Request (CSR) that you submit to a certificate authority (CA) to obtain the server certificate. You'll then use the AWS CLI to upload the certificate, private key, and certificate chain to IAM.

You'll also need an X.509 certificate to create a customized Linux AMI for EC2 instances. The certificate is only required to create an instance-backed AMI (as opposed to an EBS-backed AMI). You can have AWS create an X.509 certificate and private key that you can download, or you can upload your own certificate by using the Security Credentials page.

Individual User Accounts

AWS provides a centralized mechanism called AWS Identity and Access Management (IAM) for creating and managing individual users within your AWS Account. A user can be any individual, system, or application that interacts with AWS resources, either programmatically or through the AWS Management Console or AWS Command Line Interface (CLI). Each user has a unique name within the AWS Account, and a unique set of security credentials not shared with other users. AWS IAM eliminates the need to share passwords or keys, and enables you to minimize the use of your AWS Account credentials.

With IAM, you define policies that control which AWS services your users can access and what they can do with them. You can grant users only the minimum permissions they need to perform their jobs. See the AWS Identity and Access Management (AWS IAM) section below for more information.

Secure HTTPS Access Points

For greater communication security when accessing AWS resources, you should use HTTPS instead of HTTP for data transmissions. HTTPS uses the SSL/TLS protocol, which uses public-key cryptography to prevent eavesdropping, tampering, and forgery. All AWS services provide secure customer access points (also called API endpoints) that allow you to establish secure HTTPS communication sessions.

Several services also now offer more advanced cipher suites that use the Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) protocol. ECDHE allows SSL/TLS clients to provide Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This helps prevent the decoding of captured data by unauthorized third parties, even if the secret long-term key itself is compromised.

Security Logs

As important as credentials and encrypted endpoints are for preventing security problems, logs are just as crucial for understanding events after a problem has occurred. And to be effective as a security tool, a log must include not just a list of what happened and when, but also identify the source. To help you with your after-the-fact investigations and near-realtime intrusion detection, AWS CloudTrail provides a log of all requests for AWS resources within your account. For each event, you can see what service was accessed, what action was performed, and who made the request. CloudTrail captures information about every API call to every AWS resource you use, including sign-in events.



Once you have enabled CloudTrail, event logs are delivered every 5 minutes. You can configure CloudTrail so that it aggregates log files from multiple regions into a single Amazon S3 bucket. From there, you can then upload them to your favorite log management and analysis solutions to perform security analysis and detect user behavior patterns. By default, log files are stored securely in Amazon S3, but you can also archive them to Amazon Glacier to help meet audit and compliance requirements.

In addition to CloudTrail's user activity logs, you can use the Amazon CloudWatch Logs feature to collect and monitor system, application, and custom log files from your EC2 instances and other sources in near-real time. For example, you can monitor your web server's log files for invalid user messages to detect unauthorized login attempts to your guest OS.

AWS Trusted Advisor Security Checks

The AWS Trusted Advisor customer support service not only monitors for cloud performance and resiliency, but also cloud security. Trusted Advisor inspects your AWS environment and makes recommendations when opportunities may exist to save money, improve system performance, or close security gaps. It provides alerts on several of the most common security misconfigurations that can occur, including leaving certain ports open that make you vulnerable to hacking and unauthorized access, neglecting to create IAM accounts for your internal users, allowing public access to Amazon S3 buckets, not turning on user activity logging (AWS CloudTrail), or not using MFA on your root AWS Account. You also have the option for a Security contact at your organization to automatically receive a weekly email with an updated status of your Trusted Advisor security checks.

The AWS Trusted Advisor service provides four checks at no additional charge to all users, including three important security checks: specific ports unrestricted, IAM use, and MFA on root account. And when you sign up for Business- or Enterprise-level AWS Support, you receive full access to all Trusted Advisor checks.

AWS Service-Specific Security

Not only is security built into every layer of the AWS infrastructure, but also into each of the services available on that infrastructure. AWS services are architected to work efficiently and securely with all AWS networks and platforms. Each service provides extensive security features to enable you to protect sensitive data and applications.

Compute Services

Amazon Web Services provides a variety of cloud-based computing services that include a wide selection of compute instances that can scale up and down automatically to meet the needs of your application or enterprise.

Amazon Elastic Compute Cloud (Amazon EC2) Security

Amazon Elastic Compute Cloud (EC2) is a key component in Amazon's Infrastructure as a Service (IaaS), providing resizable computing capacity using server instances in AWS's data centers. Amazon EC2 is designed to make web-scale computing easier by enabling you to obtain and configure capacity with minimal friction. You create and launch *instances*, which are collections of platform hardware and software.

Multiple Levels of Security

Security within Amazon EC2 is provided on multiple levels: the operating system (OS) of the host platform, the virtual instance OS or guest OS, a firewall, and signed API calls. Each of these items builds on the capabilities of the others. The goal is to prevent data contained within Amazon EC2 from being intercepted by unauthorized systems or users and to



provide Amazon EC2 instances themselves that are as secure as possible without sacrificing the flexibility in configuration that customers demand.

The Hypervisor

Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor, taking advantage of paravirtualization (in the case of Linux guests). Because paravirtualized guests rely on the hypervisor to provide support for operations that normally require privileged access, the guest OS has no elevated access to the CPU. The CPU provides four separate privilege modes: 0-3, called *rings*. Ring 0 is the most privileged and 3 the least. The host OS executes in Ring 0. However, rather than executing in Ring 0 as most operating systems do, the guest OS runs in a lesser-privileged Ring 1 and applications in the least privileged Ring 3. This explicit virtualization of the physical resources leads to a clear separation between guest and hypervisor, resulting in additional security separation between the two.

Instance Isolation

Different instances running on the same physical machine are isolated from each other via the Xen hypervisor. Amazon is active in the Xen community, which provides awareness of the latest developments. In addition, the AWS firewall resides within the hypervisor layer, between the physical network interface and the instance's virtual interface. All packets must pass through this layer, thus an instance's neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. The physical RAM is separated using similar mechanisms.

Customer instances have no access to raw disk devices, but instead are presented with virtualized disks. The AWS proprietary disk virtualization layer automatically resets every block of storage used by the customer, so that one customer's data is never unintentionally exposed to another. In addition, memory allocated to guests is scrubbed (set to zero) by the hypervisor when it is unallocated to a guest. The memory is not returned to the pool of free memory available for new allocations until the memory scrubbing is complete.

AWS recommends customers further protect their data using appropriate means. One common solution is to run an encrypted file system on top of the virtualized disk device.




Figure 3: Amazon EC2 Multiple Layers of Security

Host Operating System: Administrators with a business need to access the management plane are required to use multi-factor authentication to gain access to purpose-built administration hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane of the cloud. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems can be revoked.

Guest Operating System: Virtual instances are completely controlled by you, the customer. You have full root access or administrative control over accounts, services, and applications. AWS does not have any access rights to your instances or the guest OS. AWS recommends a base set of security best practices to include disabling password-only access to your guests, and utilizing some form of multi-factor authentication to gain access to your instances (or at a minimum certificate-based SSH Version 2 access). Additionally, you should employ a privilege escalation mechanism with logging on a per-user basis. For example, if the guest OS is Linux, after hardening your instance you should utilize certificate-based SSHv2 to access the virtual instance, disable remote root login, use command-line logging, and use 'sudo' for privilege escalation. You should generate your own key pairs in order to guarantee that they are unique, and not shared with other customers or with AWS.

AWS also supports the use of the Secure Shell (SSH) network protocol to enable you to log in securely to your UNIX/Linux EC2 instances. Authentication for SSH used with AWS is via a public/private key pair to reduce the risk of unauthorized access to your instance. You can also connect remotely to your Windows instances using Remote Desktop Protocol (RDP) by utilizing an RDP certificate generated for your instance.

You also control the updating and patching of your guest OS, including security updates. Amazon-provided Windows and Linux-based AMIs are updated regularly with the latest patches, so if you do not need to preserve data or customizations on your running Amazon AMI instances, you can simply relaunch new instances with the latest updated AMI. In addition, updates are provided for the Amazon Linux AMI via the Amazon Linux yum repositories.

Firewall: Amazon EC2 provides a complete firewall solution; this mandatory inbound firewall is configured in a default deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic. The traffic may be restricted by protocol, by service port, as well as by source IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

The firewall can be configured in groups permitting different classes of instances to have different rules. Consider, for example, the case of a traditional three-tiered web application. The group for the web servers would have port 80 (HTTP) and/or port 443 (HTTPS) open to the Internet. The group for the application servers would have port 8000 (application specific) accessible only to the web server group. The group for the database servers would have port 3306 (MySQL) open only to the application server group. All three groups would permit administrative access on port 22 (SSH), but only from the customer's corporate network. Highly secure applications can be deployed using this expressive mechanism. See diagram below:




Figure 4: Amazon EC2 Security Group Firewall

The firewall isn't controlled through the guest OS; rather it requires your X.509 certificate and key to authorize changes, thus adding an extra layer of security. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling you to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports you open, and for what duration and purpose. The default state is to deny all incoming traffic, and you should plan carefully what you will open when building and securing your applications. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages you to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall and VPNs. This can restrict both inbound and outbound traffic.

API Access: API calls to launch and terminate instances, change firewall parameters, and perform other functions are all signed by your Amazon Secret Access Key, which could be either the AWS Accounts Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to your Secret Access Key, Amazon EC2 API calls cannot be

made on your behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints.

Permissions: AWS IAM also enables you to further control what APIs a user has permissions to call.

Elastic Block Storage (Amazon EBS) Security

Amazon Elastic Block Storage (EBS) allows you to create storage volumes from 1 GB to 16 TB that can be mounted as devices by Amazon EC2 instances. Storage volumes behave like raw, unformatted block devices, with user supplied device names and a block device interface. You can create a file system on top of Amazon EBS volumes, or use them in any other way you would use a block device (like a hard drive). Amazon EBS volume access is restricted to the AWS Account that created the volume, and to the users under the AWS Account created with AWS IAM if the user has been granted access to the EBS operations, thus denying all other AWS Accounts and users the permission to view or access the volume.

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability. For customers who have architected complex transactional databases using EBS, it is recommended that backups to Amazon S3 be performed through the database management system so that distributed transactions and logs can be checkpointed. AWS does not perform backups of data that are maintained on virtual disks attached to running instances on Amazon EC2.

You can make Amazon EBS volume snapshots publicly available to other AWS Accounts to use as the basis for creating your own volumes. Sharing Amazon EBS volume snapshots does not provide other AWS Accounts with the permission to alter or delete the original snapshot, as that right is explicitly reserved for the AWS Account that created the volume. An EBS snapshot is a block-level view of an entire EBS volume. Note that data that is not visible through the file system on the volume, such as files that have been deleted, may be present in the EBS snapshot. If you want to create shared snapshots, you should do so carefully. If a volume has held sensitive data or has had files deleted from it, a new EBS volume should be created. The data to be contained in the shared snapshot should be copied to the new volume, and the snapshot created from the new volume.

Amazon EBS volumes are presented to you as raw unformatted block devices that have been wiped prior to being made available for use. Wiping occurs immediately before reuse so that you can be assured that the wipe process completed. If you have procedures requiring that all data be wiped via a specific method, such as those detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization"), you have the ability to do so on Amazon EBS. You should conduct a specialized wipe procedure prior to deleting the volume for compliance with your established requirements.

Encryption of sensitive data is generally a good security practice, and AWS provides the ability to encrypt EBS volumes and their snapshots with AES-256. The encryption occurs on the servers that host the EC2 instances, providing encryption of data as it moves between EC2 instances and EBS storage. In order to be able to do this efficiently and with low latency, the EBS encryption feature is only available on EC2's more powerful instance types (e.g., M3, C3, R3, G2).

Auto Scaling Security

Auto Scaling allows you to automatically scale your Amazon EC2 capacity up or down according to conditions you define, so that the number of Amazon EC2 instances you are using scales up seamlessly during demand spikes to maintain performance, and scales down automatically during demand lulls to minimize costs.



Like all AWS services, Auto Scaling requires that every request made to its control API be authenticated so only authenticated users can access and manage Auto Scaling. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. However, getting credentials out to new EC2 instances launched with Auto Scaling can be challenging for large or elastically scaling fleets. To simplify this process, you can use *roles* within IAM, so that any new instances launched with a role will be given credentials automatically. When you launch an EC2 instance with an IAM role, temporary AWS security credentials with permissions specified by the role will be securely provisioned to the instance and will be made available to your application via the Amazon EC2 Instance Metadata Service. The Metadata Service will make new temporary security credentials available prior to the expiration of the current active credentials, so that valid credentials are always available on the instance. In addition, the temporary security credentials are automatically rotated multiple times per day, providing enhanced security. You can further control access to Auto Scaling by creating users under your AWS Account using AWS IAM, and controlling what Auto Scaling APIs these users have permission to call. More information about using roles when launching instances is available in the Amazon EC2 User Guide on the AWS website: <http://docs.amazonwebservices.com/AWSEC2/latest/UserGuide/UsingIAM>

Networking Services

Amazon Web Services provides a range of networking services that enable you to create a logically isolated network that you define, establish a private network connection to the AWS cloud, use a highly available and scalable DNS service and deliver content to your end users with low latency at high data transfer speeds with a content delivery web service.

Amazon Elastic Load Balancing Security

Amazon Elastic Load Balancing is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

- Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer
- Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network
- When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options
- Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance.

HTTPS/TLS uses a long-term secret key to generate a short-term session key to be used between the server and the browser to create the ciphered (encrypted) message. Amazon Elastic Load Balancing configures your load balancer with a pre-defined cipher set that is used for TLS negotiation when a connection is established between a client and your load balancer. The pre-defined cipher set provides compatibility with a broad range of clients and uses strong cryptographic algorithms. However, some customers may have requirements for allowing only specific ciphers and protocols (such as PCI, SOX, etc.) from clients to ensure that standards are met. In these cases, Amazon Elastic Load Balancing provides options for selecting different configurations for TLS protocols and ciphers. You can choose to enable or disable the ciphers depending on your specific requirements.

To help ensure the use of newer and stronger cipher suites when establishing a secure connection, you can configure the load balancer to have the final say in the cipher suite selection during the client-server negotiation. When the Server Order Preference option is selected, the load balancer will select a cipher suite based on the server's prioritization of cipher suites rather than the client's. This gives you more control over the level of security that clients use to connect to your load balancer.

For even greater communication privacy, Amazon Elastic Load Balancer allows the use of Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This prevents the decoding of captured data, even if the secret long-term key itself is compromised.

Amazon Elastic Load Balancing allows you to identify the originating IP address of a client connecting to your servers, whether you're using HTTPS or TCP load balancing. Typically, client connection information, such as IP address and port, is lost when requests are proxied through a load balancer. This is because the load balancer sends requests to the server on behalf of the client, making your load balancer appear as though it is the requesting client. Having the originating client IP address is useful if you need more information about visitors to your applications in order to gather connection statistics, analyze traffic logs, or manage whitelists of IP addresses.

Amazon Elastic Load Balancing access logs contain information about each HTTP and TCP request processed by your load balancer. This includes the IP address and port of the requesting client, the backend IP address of the instance that processed the request, the size of the request and response, and the actual request line from the client (for example, GET http://www.example.com: 80/HTTP/1.1). All requests sent to the load balancer are logged, including requests that never made it to back-end instances.

Amazon Virtual Private Cloud (Amazon VPC) Security

Normally, each Amazon EC2 instance you launch is randomly assigned a public IP address in the Amazon EC2 address space. Amazon VPC enables you to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses in the range of your choice (e.g., 10.0.0.0/16). You can define subnets within your VPC, grouping similar kinds of instances based on IP address range, and then set up routing and security to control the flow of traffic in and out of the instances and subnets.

AWS offers a variety of VPC architecture templates with configurations that provide varying levels of public access:

- **VPC with a single public subnet only.** Your instances run in a private, isolated section of the AWS cloud with direct access to the Internet. Network ACLs and security groups can be used to provide strict control over inbound and outbound network traffic to your instances.
- **VPC with public and private subnets.** In addition to containing a public subnet, this configuration adds a private subnet whose instances are not addressable from the Internet. Instances in the private subnet can establish outbound connections to the Internet via the public subnet using Network Address Translation (NAT).
- **VPC with public and private subnets and hardware VPN access.** This configuration adds an IPsec VPN connection between your Amazon VPC and your data center, effectively extending your data center to the cloud while also providing direct access to the Internet for public subnet instances in your Amazon VPC. In this configuration, customers add a VPN appliance on their corporate data center side.
- **VPC with private subnet only and hardware VPN access.** Your instances run in a private, isolated section of the AWS cloud with a private subnet whose instances are not addressable from the Internet. You can connect this

private subnet to your corporate data center via an IPsec VPN tunnel.

You can also connect two VPCs using a private IP address, which allows instances in the two VPCs to communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another AWS account within a single region.

Security features within Amazon VPC include security groups, network ACLs, routing tables, and external gateways. Each of these items is complementary to providing a secure, isolated network that can be extended through selective enabling of direct Internet access or private connectivity to another network. Amazon EC2 instances running within an Amazon VPC inherit all of the benefits described below related to the guest OS and protection against packet sniffing. Note, however, that you must create VPC security groups specifically for your Amazon VPC; any Amazon EC2 security groups you have created will not work inside your Amazon VPC. Also, Amazon VPC security groups have additional capabilities that Amazon EC2 security groups do not have, such as being able to change the security group after the instance is launched and being able to specify any protocol with a standard protocol number (as opposed to just TCP, UDP, or ICMP).

Each Amazon VPC is a distinct, isolated network within the cloud; network traffic within each Amazon VPC is isolated from all other Amazon VPCs. At creation time, you select an IP address range for each Amazon VPC. You may create and attach an Internet gateway, virtual private gateway, or both to establish external connectivity, subject to the controls below.

API Access: Calls to create and delete Amazon VPCs, change routing, security group, and network ACL parameters, and perform other functions are all signed by your Amazon Secret Access Key, which could be either the AWS Account's Secret Access Key or the Secret Access key of a user created with AWS IAM. Without access to your Secret Access Key, Amazon VPC API calls cannot be made on your behalf. In addition, API calls can be encrypted with SSL to maintain confidentiality. Amazon recommends always using SSL-protected API endpoints. AWS IAM also enables a customer to further control what APIs a newly created user has permissions to call.

Subnets and Route Tables: You create one or more subnets within each Amazon VPC; each instance launched in the Amazon VPC is connected to one subnet. Traditional Layer 2 security attacks, including MAC spoofing and ARP spoofing, are blocked.

Each subnet in an Amazon VPC is associated with a routing table, and all network traffic leaving the subnet is processed by the routing table to determine the destination.

Firewall (Security Groups): Like Amazon EC2, Amazon VPC supports a complete firewall solution enabling filtering on both ingress and egress traffic from an instance. The default group enables inbound communication from other members of the same group and outbound communication to any destination. Traffic can be restricted by any IP protocol, by service port, as well as source/destination IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

The firewall isn't controlled through the guest OS; rather, it can be modified only through the invocation of Amazon VPC APIs. AWS supports the ability to grant granular access to different administrative functions on the instances and the firewall, therefore enabling you to implement additional security through separation of duties. The level of security afforded by the firewall is a function of which ports you open, and for what duration and purpose. Well-informed traffic management and security design are still required on a per-instance basis. AWS further encourages you to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall.






Figure 5: Amazon VPC Network Architecture

Network Access Control Lists: To add a further layer of security within Amazon VPC, you can configure network ACLs. These are stateless traffic filters that apply to all traffic inbound or outbound from a subnet within Amazon VPC. These ACLs can contain ordered rules to allow or deny traffic based upon IP protocol, by service port, as well as source/destination IP address.

Like security groups, network ACLs are managed through Amazon VPC APIs, adding an additional layer of protection and enabling additional security through separation of duties. The diagram below depicts how the security controls above inter-relate to enable flexible network topologies while providing complete control over network traffic flows.




Figure 6: Flexible Network Topologies

Virtual Private Gateway: A virtual private gateway enables private connectivity between the Amazon VPC and another network. Network traffic within each virtual private gateway is isolated from network traffic within all other virtual private gateways. You can establish VPN connections to the virtual private gateway from gateway devices at your premises. Each connection is secured by a pre-shared key in conjunction with the IP address of the customer gateway device.

Internet Gateway: An Internet gateway may be attached to an Amazon VPC to enable direct connectivity to Amazon S3, other AWS services, and the Internet. Each instance desiring this access must either have an Elastic IP associated with it or route traffic through a NAT instance. Additionally, network routes are configured (see above) to direct traffic to the Internet gateway. AWS provides reference NAT AMIs that you can extend to perform network logging, deep packet inspection, application-layer filtering, or other security controls.

This access can only be modified through the invocation of Amazon VPC APIs. AWS supports the ability to grant granular access to different administrative functions on the instances and the Internet gateway, therefore enabling you to implement additional security through separation of duties.

Dedicated Instances: Within a VPC, you can launch Amazon EC2 instances that are physically isolated at the host hardware level (i.e., they will run on single-tenant hardware). An Amazon VPC can be created with ‘dedicated’ tenancy, so that all instances launched into the Amazon VPC will utilize this feature. Alternatively, an Amazon VPC may be created with ‘default’ tenancy, but you can specify dedicated tenancy for particular instances launched into it.

Elastic Network Interfaces: Each Amazon EC2 instance has a default network interface that is assigned a private IP address on your Amazon VPC network. You can create and attach an additional network interface, known as an elastic network interface (ENI), to any Amazon EC2 instance in your Amazon VPC for a total of two network interfaces per instance. Attaching more than one network interface to an instance is useful when you want to create a management network, use network and security appliances in your Amazon VPC, or create dual-homed instances with workloads/roles on distinct subnets. An ENI's attributes, including the private IP address, elastic IP addresses, and MAC address, will follow the ENI as it is attached or detached from an instance and reattached to another instance. More information about Amazon VPC is available on the AWS website: <http://aws.amazon.com/vpc/>

Additional Network Access Control with EC2-VPC

If you launch instances in a region where you did not have instances before AWS launched the new EC2-VPC feature (also called Default VPC), all instances are automatically provisioned in a ready-to-use default VPC. You can choose to create additional VPCs, or you can create VPCs for instances in regions where you already had instances before we launched EC2-VPC.

If you create a VPC later, using regular VPC, you specify a CIDR block, create subnets, enter the routing and security for those subnets, and provision an Internet gateway or NAT instance if you want one of your subnets to be able to reach the Internet. When you launch EC2 instances into an EC2-VPC, most of this work is automatically performed for you. When you launch an instance into a default VPC using EC2-VPC, we do the following to set it up for you:

- Create a default subnet in each Availability Zone
- Create an Internet gateway and connect it to your default VPC
- Create a main route table for your default VPC with a rule that sends all traffic destined for the Internet to the Internet gateway
- Create a default security group and associate it with your default VPC
- Create a default network access control list (ACL) and associate it with your default VPC
- Associate the default DHCP options set for your AWS account with your default VPC

In addition to the default VPC having its own private IP range, EC2 instances launched in a default VPC can also receive a public IP.

The following table summarizes the differences between instances launched into EC2-Classic, instances launched into a default VPC, and instances launched into a non-default VPC.

Characteristic	EC2-Classic	EC2-VPC (Default VPC)	Regular VPC
Public IP address	Your instance receives a public IP address.	Your instance launched in a default subnet receives a public	Your instance doesn't receive a public IP address by default,

Characteristic	EC2-Classic	EC2-VPC (Default VPC)	Regular VPC
		IP address by default, unless you specify otherwise during launch.	unless you specify otherwise during launch.
Private IP address	Your instance receives a private IP address from the EC2-Classic range each time it's started.	Your instance receives a static private IP address from the address range of your default VPC.	Your instance receives a static private IP address from the address range of your VPC.
Multiple private IP addresses	We select a single IP address for your instance. Multiple IP addresses are not supported.	You can assign multiple private IP addresses to your instance.	You can assign multiple private IP addresses to your instance.
Elastic IP address	An EIP is disassociated from your instance when you stop it.	An EIP remains associated with your instance when you stop it.	An EIP remains associated with your instance when you stop it.
DNS hostnames	DNS hostnames are enabled by default.	DNS hostnames are enabled by default.	DNS hostnames are disabled by default.
Security group	A security group can reference security groups that belong to other AWS accounts.	A security group can reference security groups for your VPC only.	A security group can reference security groups for your VPC only.
Security group association	You must terminate your instance to change its security group.	You can change the security group of your running instance.	You can change the security group of your running instance.
Security group rules	You can add rules for inbound traffic only.	You can add rules for inbound and outbound traffic.	You can add rules for inbound and outbound traffic.
Tenancy	Your instance runs on shared hardware; you cannot run an instance on single-tenant hardware.	You can run your instance on shared hardware or single-tenant hardware.	You can run your instance on shared hardware or single-tenant hardware.

Note that security groups for instances in EC2-Classic are slightly different than security groups for instances in EC2-VPC. For example, you can add rules for inbound traffic for EC2-Classic, but you can add rules for both inbound and outbound traffic to EC2-VPC. In EC2-Classic, you can't change the security groups assigned to an instance after it's launched, but in EC2-VPC, you can change security groups assigned to an instance after it's launched. In addition, you can't use the security groups that you've created for use with EC2-Classic with instances in your VPC. You must create security groups specifically for use with instances in your VPC. The rules you create for use with a security group for a VPC can't reference a security group for EC2-Classic, and vice versa.

Amazon Route 53 Security

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) service that answers DNS queries, translating domain names into IP addresses so computers can communicate with each other. Route 53 can be used to connect user requests to infrastructure running in AWS – such as an Amazon EC2 instance or an Amazon S3 bucket – or to infrastructure outside of AWS.

Amazon Route 53 lets you manage the IP addresses (records) listed for your domain names and it answers requests (queries) to translate specific domain names into their corresponding IP addresses. Queries for your domain are automatically routed to a nearby DNS server using anycast in order to provide the lowest latency possible. Route 53



makes it possible for you to manage traffic globally through a variety of routing types, including Latency Based Routing (LBR), Geo DNS, and Weighted Round-Robin (WRR) —all of which can be combined with DNS Failover in order to help create a variety of low-latency, fault-tolerant architectures. The failover algorithms implemented by Amazon Route 53 are designed not only to route traffic to endpoints that are healthy, but also to help avoid making disaster scenarios worse due to misconfigured health checks and applications, endpoint overloads, and partition failures.

Route 53 also offers Domain Name Registration – you can purchase and manage domain names such as example.com and Route 53 will automatically configure default DNS settings for your domains. You can buy, manage, and transfer (both in and out) domains from a wide selection of generic and country-specific top-level domains (TLDs). During the registration process, you have the option to enable privacy protection for your domain. This option will hide most of your personal information from the public Whois database in order to help thwart scraping and spamming.

Amazon Route 53 is built using AWS's highly available and reliable infrastructure. The distributed nature of the AWS DNS servers helps ensure a consistent ability to route your end users to your application. Route 53 also helps ensure the availability of your website by providing health checks and DNS failover capabilities. You can easily configure Route 53 to check the health of your website on a regular basis (even secure web sites that are available only over SSL), and to switch to a backup site if the primary one is unresponsive.

Like all AWS Services, Amazon Route 53 requires that every request made to its control API be authenticated so only authenticated users can access and manage Route 53. API requests are signed with an HMAC-SHA1 or HMAC-SHA256 signature calculated from the request and the user's AWS Secret Access key. Additionally, the Amazon Route 53 control API is only accessible via SSL-encrypted endpoints. It supports both IPv4 and IPv6 routing.

You can control access to Amazon Route 53 DNS management functions by creating users under your AWS Account using AWS IAM, and controlling which Route 53 operations these users have permission to perform.

Amazon CloudFront Security

Amazon CloudFront gives customers an easy way to distribute content to end users with low latency and high data transfer speeds. It delivers dynamic, static, and streaming content using a global network of edge locations. Requests for customers' objects are automatically routed to the nearest edge location, so content is delivered with the best possible performance. Amazon CloudFront is optimized to work with other AWS services, like Amazon S3, Amazon EC2, Amazon Elastic Load Balancing, and Amazon Route 53. It also works seamlessly with any non-AWS origin server that stores the original, definitive versions of your files.

Amazon CloudFront requires every request made to its control API be authenticated so only authorized users can create, modify, or delete their own Amazon CloudFront distributions. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Additionally, the Amazon CloudFront control API is only accessible via SSL-enabled endpoints.

There is no guarantee of durability of data held in Amazon CloudFront edge locations. The service may from time to time remove objects from edge locations if those objects are not requested frequently. Durability is provided by Amazon S3, which works as the origin server for Amazon CloudFront holding the original, definitive copies of objects delivered by Amazon CloudFront.

If you want control over who is able to download content from Amazon CloudFront, you can enable the service's private content feature. This feature has two components: the first controls how content is delivered from the Amazon CloudFront edge location to viewers on the Internet. The second controls how the Amazon CloudFront edge locations



access objects in Amazon S3. CloudFront also supports Geo Restriction, which restricts access to your content based on the geographic location of your viewers.

To control access to the original copies of your objects in Amazon S3, Amazon CloudFront allows you to create one or more “Origin Access Identities” and associate these with your distributions. When an Origin Access Identity is associated with an Amazon CloudFront distribution, the distribution will use that identity to retrieve objects from Amazon S3. You can then use Amazon S3’s ACL feature, which limits access to that Origin Access Identity so the original copy of the object is not publicly readable.

To control who is able to download objects from Amazon CloudFront edge locations, the service uses a signed-URL verification system. To use this system, you first create a public-private key pair, and upload the public key to your account via the AWS Management Console. Second, you configure your Amazon CloudFront distribution to indicate which accounts you would authorize to sign requests – you can indicate up to five AWS Accounts you trust to sign requests. Third, as you receive requests you will create policy documents indicating the conditions under which you want Amazon CloudFront to serve your content. These policy documents can specify the name of the object that is requested, the date and time of the request, and the source IP (or CIDR range) of the client making the request. You then calculate the SHA1 hash of your policy document and sign this using your private key. Finally, you include both the encoded policy document and the signature as query string parameters when you reference your objects. When Amazon CloudFront receives a request, it will decode the signature using your public key. Amazon CloudFront will only serve requests that have a valid policy document and matching signature.

Note that private content is an optional feature that must be enabled when you set up your CloudFront distribution. Content delivered without this feature enabled will be publicly readable.

Amazon CloudFront provides the option to transfer content over an encrypted connection (HTTPS). By default, CloudFront will accept requests over both HTTP and HTTPS protocols. However, you can also configure CloudFront to require HTTPS for all requests or have CloudFront redirect HTTP requests to HTTPS. You can even configure CloudFront distributions to allow HTTP for some objects but require HTTPS for other objects.




Figure 7: Amazon CloudFront Encrypted Transmission

You can configure one or more CloudFront origins to require CloudFront fetch objects from your origin using the protocol that the viewer used to request the objects. For example, when you use this CloudFront setting and the viewer uses HTTPS to request an object from CloudFront, CloudFront also uses HTTPS to forward the request to your origin.

Amazon CloudFront uses the SSLv3 or TLSv1 protocols and a selection of cipher suites that includes the Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) protocol on connections to both viewers and the origin. ECDHE allows SSL/TLS clients

to provide Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This helps prevent the decoding of captured data by unauthorized third parties, even if the secret long-term key itself is compromised.

Note that if you're using your own server as your origin, and you want to use HTTPS both between viewers and CloudFront and between CloudFront and your origin, you must install a valid SSL certificate on the HTTP server that is signed by a third-party certificate authority, for example, VeriSign or DigiCert.

By default, you can deliver content to viewers over HTTPS by using your CloudFront distribution domain name in your URLs; for example, <https://dxxxxx.cloudfront.net/image.jpg>. If you want to deliver your content over HTTPS using your own domain name and your own SSL certificate, you can use SNI Custom SSL or Dedicated IP Custom SSL. With Server Name Identification (SNI) Custom SSL, CloudFront relies on the SNI extension of the TLS protocol, which is supported by most modern web browsers. However, some users may not be able to access your content because some older browsers do not support SNI. (For a list of supported browsers, visit <http://aws.amazon.com/cloudfront/faqs/>.) With Dedicated IP Custom SSL, CloudFront dedicates IP addresses to your SSL certificate at each CloudFront edge location so that CloudFront can associate the incoming requests with the proper SSL certificate.

Amazon CloudFront access logs contain a comprehensive set of information about requests for content, including the object requested, the date and time of the request, the edge location serving the request, the client IP address, the referrer, and the user agent. To enable access logs, just specify the name of the Amazon S3 bucket to store the logs in when you configure your Amazon CloudFront distribution.

AWS Direct Connect Security

With AWS Direct Connect, you can provision a direct link between your internal network and an AWS region using a high-throughput, dedicated connection. Doing this may help reduce your network costs, improve throughput, or provide a more consistent network experience. With this dedicated connection in place, you can then create virtual interfaces directly to the AWS cloud (for example, to Amazon EC2 and Amazon S3) and Amazon VPC.

With Direct Connect, you bypass Internet service providers in your network path. You can procure rack space within the facility housing the AWS Direct Connect location and deploy your equipment nearby. Once deployed, you can connect this equipment to AWS Direct Connect using a cross-connect. Each AWS Direct Connect location enables connectivity to the geographically nearest AWS region as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).

Using industry standard 802.1q VLANs, the dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon VPC using private IP space, while maintaining network separation between the public and private environments.

Amazon Direct Connect requires the use of the Border Gateway Protocol (BGP) with an Autonomous System Number (ASN). To create a virtual interface, you use an MD5 cryptographic key for message authorization. MD5 creates a keyed hash using your secret key. You can have AWS automatically generate a BGP MD5 key or you can provide your own.

Storage Services

Amazon Web Services provides low-cost data storage with high durability and availability. AWS offers storage choices for backup, archiving, and disaster recovery, as well as block and object storage.

Amazon Simple Storage Service (Amazon S3) Security

Amazon Simple Storage Service (S3) allows you to upload and retrieve data at any time, from anywhere on the web. Amazon S3 stores data as *objects* within *buckets*. An object can be any kind of file: a text file, a photo, a video, etc. When you add a file to Amazon S3, you have the option of including metadata with the file and setting permissions to control access to the file. For each bucket, you can control access to the bucket (who can create, delete, and list objects in the bucket), view access logs for the bucket and its objects, and choose the geographical region where Amazon S3 will store the bucket and its contents.

Data Access

Access to data stored in Amazon S3 is restricted by default; only bucket and object owners have access to the Amazon S3 resources they create (note that a bucket/object owner is the AWS Account owner, not the user who created the bucket/object). There are multiple ways to control access to buckets and objects:

- **Identity and Access Management (IAM) Policies.** AWS IAM enables organizations with many employees to create and manage multiple users under a single AWS Account. IAM policies are attached to the users, enabling centralized control of permissions for users under your AWS Account to access buckets or objects. With IAM policies, you can only grant *users within your own AWS account* permission to access your Amazon S3 resources.
- **Access Control Lists (ACLs).** Within Amazon S3, you can use ACLs to give read or write access on buckets or objects to groups of users. With ACLs, you can only grant *other AWS accounts* (not specific users) access to your Amazon S3 resources.
- **Bucket Policies.** Bucket policies in Amazon S3 can be used to add or deny permissions across some or all of the objects within a single bucket. Policies can be attached to users, groups, or Amazon S3 buckets, enabling centralized management of permissions. With bucket policies, you can grant users within your AWS Account *or* other AWS Accounts access to your Amazon S3 resources.

Type of Access Control	AWS Account-Level Control?	User-Level Control?
IAM Policies	No	Yes
ACLs	Yes	No
Bucket Policies	Yes	Yes

You can further restrict access to specific resources based on certain conditions. For example, you can restrict access based on request time (Date Condition), whether the request was sent using SSL (Boolean Conditions), a requester's IP address (IP Address Condition), or based on the requester's client application (String Conditions). To identify these conditions, you use *policy keys*. For more information about action-specific policy keys available within Amazon S3, refer to the [Amazon Simple Storage Service Developer Guide](#).

Amazon S3 also gives developers the option to use *query string authentication*, which allows them to share Amazon S3 objects through URLs that are valid for a predefined period of time. Query string authentication is useful for giving HTTP



or browser access to resources that would normally require authentication. The signature in the query string secures the request.

Data Transfer

For maximum security, you can securely upload/download data to Amazon S3 via the SSL encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data is transferred securely both within AWS and to and from sources outside of AWS.

Data Storage

Amazon S3 provides multiple options for protecting data at rest. For customers who prefer to manage their own encryption, they can use a client encryption library like the [Amazon S3 Encryption Client](#) to encrypt data before uploading to Amazon S3. Alternatively, you can use Amazon S3 Server Side Encryption (SSE) if you prefer to have Amazon S3 manage the encryption process for you. Data is encrypted with a key generated by AWS or with a key you supply, depending on your requirements. With Amazon S3 SSE, you can encrypt data on upload simply by adding an additional request header when writing the object. Decryption happens automatically when data is retrieved.

Note that metadata, which you can include with your object, is not encrypted. Therefore, AWS recommends that customers not place sensitive information in Amazon S3 metadata.

Amazon S3 SSE uses one of the strongest block ciphers available – 256-bit Advanced Encryption Standard (AES-256). With Amazon S3 SSE, every protected object is encrypted with a unique encryption key. This object key itself is then encrypted with a regularly rotated master key. Amazon S3 SSE provides additional security by storing the encrypted data and encryption keys in different hosts. Amazon S3 SSE also makes it possible for you to enforce encryption requirements. For example, you can create and apply bucket policies that require that only encrypted data can be uploaded to your buckets.

For long-term storage, you can automatically archive the contents of your Amazon S3 buckets to AWS's archival service called Glacier. You can have data transferred at specific intervals to Glacier by creating lifecycle rules in Amazon S3 that describe which objects you want to be archived to Glacier and when. As part of your data management strategy, you can also specify how long Amazon S3 should wait after the objects are put into Amazon S3 to delete them.

When an object is deleted from Amazon S3, removal of the mapping from the public name to the object starts immediately, and is generally processed across the distributed system within several seconds. Once the mapping is removed, there is no remote access to the deleted object. The underlying storage area is then reclaimed for use by the system.

Data Durability and Reliability

Amazon S3 is designed to provide 99.999999999% durability and 99.99% availability of objects over a given year. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 region. To help provide durability, Amazon S3 PUT and COPY operations synchronously store customer data across multiple facilities before returning SUCCESS. Once stored, Amazon S3 helps maintain the durability of the objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. In addition, Amazon S3 calculates checksums on all network traffic to detect corruption of data packets when storing or retrieving data.

Amazon S3 provides further protection via Versioning. You can use Versioning to preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket. With Versioning, you can easily recover from both unintended



user actions and application failures. By default, requests will retrieve the most recently written version. Older versions of an object can be retrieved by specifying a version in the request. You can further protect versions using Amazon S3 Versioning's MFA Delete feature. Once enabled for an Amazon S3 bucket, each version deletion request must include the six-digit code and serial number from your multi-factor authentication device.

Access Logs

An Amazon S3 bucket can be configured to log access to the bucket and objects within it. The access log contains details about each access request including request type, the requested resource, the requestor's IP, and the time and date of the request. When logging is enabled for a bucket, log records are periodically aggregated into log files and delivered to the specified Amazon S3 bucket.

Cross-Origin Resource Sharing (CORS)

AWS customers who use Amazon S3 to host static web pages or store objects used by other web pages can load content securely by configuring an Amazon S3 bucket to explicitly enable cross-origin requests. Modern browsers use the Same Origin policy to block JavaScript or HTML5 from allowing requests to load content from another site or domain as a way to help ensure that malicious content is not loaded from a less reputable source (such as during cross-site scripting attacks). With the Cross-Origin Resource Sharing (CORS) policy enabled, assets such as web fonts and images stored in an Amazon S3 bucket can be safely referenced by external web pages, style sheets, and HTML5 applications.

AWS Glacier Security

Like Amazon S3, the Amazon Glacier service provides low-cost, secure, and durable storage. But where Amazon S3 is designed for rapid retrieval, Glacier is meant to be used as an archival service for data that is not accessed often and for which retrieval times of several hours are suitable.

Amazon Glacier stores files as *archives* within *vaults*. Archives can be any data such as a photo, video, or document, and can contain one or several files. You can store an unlimited number of archives in a single vault and can create up to 1,000 vaults per region. Each archive can contain up to 40 TB of data.

Data Upload

To transfer data into Amazon Glacier vaults, you can upload an archive in a single upload operation or a multipart operation. In a single upload operation, you can upload archives up to 4 GB in size. However, customers can achieve better results using the Multipart Upload API to upload archives greater than 100 MB. Using the Multipart Upload API allows you to upload large archives, up to about 40,000 GB. The Multipart Upload API call is designed to improve the upload experience for larger archives; it enables the parts to be uploaded independently, in any order, and in parallel. If a multipart upload fails, you only need to upload the failed part again and not the entire archive.

When you upload data to Glacier, you must compute and supply a tree hash. Glacier checks the hash against the data to help ensure that it has not been altered en route. A tree hash is generated by computing a hash for each megabyte-sized segment of the data, and then combining the hashes in tree fashion to represent ever-growing adjacent segments of the data.

As an alternate to using the Multipart Upload feature, customers with very large uploads to Amazon Glacier may consider using the AWS Import/Export service instead to transfer the data. AWS Import/Export facilitates moving large amounts of data into AWS using portable storage devices for transport. AWS transfers your data directly off of storage devices using Amazon's high-speed internal network, bypassing the Internet.



You can also set up Amazon S3 to transfer data at specific intervals to Glacier. You can create lifecycle rules in Amazon S3 that describe which objects you want to be archived to Glacier and when. You can also specify how long Amazon S3 should wait after the objects are put into Amazon S3 to delete them.

To achieve even greater security, you can securely upload/download data to Amazon Glacier via the SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2, so that data is transferred securely both within AWS and to and from sources outside of AWS.

Data Retrieval

Retrieving archives from Amazon Glacier requires the initiation of a retrieval job, which is generally completed in 3 to 5 hours. You can then access the data via HTTP GET requests. The data will remain available to you for 24 hours.

You can retrieve an entire archive or several files from an archive. If you want to retrieve only a subset of an archive, you can use one retrieval request to specify the range of the archive that contains the files you are interested or you can initiate multiple retrieval requests, each with a range for one or more files. You can also limit the number of vault inventory items retrieved by filtering on an archive creation date range or by setting a maximum items limit. Whichever method you choose, when you retrieve portions of your archive, you can use the supplied checksum to help ensure the integrity of the files provided that the range that is retrieved is aligned with the tree hash of the overall archive.

Data Storage

Amazon Glacier automatically encrypts the data using AES-256 and stores it durably in an immutable form. Amazon Glacier is designed to provide average annual durability of 99.999999999% for an archive. It stores each archive in multiple facilities and multiple devices. Unlike traditional systems which can require laborious data verification and manual repair, Glacier performs regular, systematic data integrity checks and is built to be automatically self-healing.

Data Access

Only your account can access your data in Amazon Glacier. To control access to your data in Amazon Glacier, you can use AWS IAM to specify which users within your account have rights to operations on a given vault.

AWS Storage Gateway Security

The AWS Storage Gateway service connects your on-premises software appliance with cloud-based storage to provide seamless and secure integration between your IT environment and AWS's storage infrastructure. The service enables you to securely upload data to AWS' scalable, reliable, and secure Amazon S3 storage service for cost-effective backup and rapid disaster recovery.

AWS Storage Gateway transparently backs up data off-site to Amazon S3 in the form of Amazon EBS snapshots. Amazon S3 redundantly stores these snapshots on multiple devices across multiple facilities, detecting and repairing any lost redundancy. The Amazon EBS snapshot provides a point-in-time backup that can be restored on-premises or used to instantiate new Amazon EBS volumes. Data is stored within a single region that you specify.

AWS Storage Gateway offers three options:

- **Gateway-Stored Volumes (where the cloud is backup).** In this option, your volume data is stored locally and then pushed to Amazon S3, where it is stored in redundant, encrypted form, and made available in the form of Elastic Block Storage (EBS) snapshots. When you use this model, the on-premises storage is primary, delivering low-latency access to your entire dataset, and the cloud storage is the backup.



- **Gateway-Cached Volumes (where the cloud is primary).** In this option, your volume data is stored encrypted in Amazon S3, visible within your enterprise's network via an iSCSI interface. Recently accessed data is cached on-premises for low-latency local access. When you use this model, the cloud storage is primary, but you get low-latency access to your active working set in the cached volumes on premises.
- **Gateway-Virtual Tape Library (VTL).** In this option, you can configure a Gateway-VTL with up to 10 virtual tape drives per gateway, 1 media changer and up to 1500 virtual tape cartridges. Each virtual tape drive responds to the SCSI command set, so your existing on-premises backup applications (either disk-to-tape or disk-to-disk-to-tape) will work without modification.

No matter which option you choose, data is asynchronously transferred from your on-premises storage hardware to AWS over SSL. The data is stored encrypted in Amazon S3 using Advanced Encryption Standard (AES) 256, a symmetric-key encryption standard using 256-bit encryption keys. The AWS Storage Gateway only uploads data that has changed, minimizing the amount of data sent over the Internet.

The AWS Storage Gateway runs as a virtual machine (VM) that you deploy on a host in your data center running VMware ESXi Hypervisor v 4.1 or v 5 or Microsoft Hyper-V (you download the VMware software during the setup process). You can also run within EC2 using a gateway AMI. During the installation and configuration process, you can create up to 12 stored volumes, 20 Cached volumes, or 1500 virtual tape cartridges per gateway. Once installed, each gateway will automatically download, install, and deploy updates and patches. This activity takes place during a maintenance window that you can set on a per-gateway basis.

The iSCSI protocol supports authentication between targets and initiators via CHAP (Challenge-Handshake Authentication Protocol). CHAP provides protection against man-in-the-middle and playback attacks by periodically verifying the identity of an iSCSI initiator as authenticated to access a storage volume target. To set up CHAP, you must configure it in both the AWS Storage Gateway console and in the iSCSI initiator software you use to connect to the target.

After you deploy the AWS Storage Gateway VM, you must activate the gateway using the AWS Storage Gateway console. The activation process associates your gateway with your AWS Account. Once you establish this connection, you can manage almost all aspects of your gateway from the console. In the activation process, you specify the IP address of your gateway, name your gateway, identify the AWS region in which you want your snapshot backups stored, and specify the gateway time zone.

AWS Import/Export Security

AWS Import/Export is a simple, secure method for physically transferring large amounts of data to Amazon S3, EBS, or Glacier storage. This service is typically used by customers who have over 100 GB of data and/or slow connection speeds that would result in very slow transfer rates over the Internet. With AWS Import/Export, you prepare a portable storage device that you ship to a secure AWS facility. AWS transfers the data directly off of the storage device using Amazon's high-speed internal network, thus bypassing the Internet. Conversely, data can also be exported from AWS to a portable storage device.

Like all other AWS services, the AWS Import/Export service requires that you securely identify and authenticate your storage device. In this case, you will submit a job request to AWS that includes your Amazon S3 bucket, Amazon EBS region, AWS Access Key ID, and return shipping address. You then receive a unique identifier for the job, a digital signature for authenticating your device, and an AWS address to ship the storage device to. For Amazon S3, you place



the signature file on the root directory of your device. For Amazon EBS, you tape the signature barcode to the exterior of the device. The signature file is used only for authentication and is not uploaded to Amazon S3 or EBS.

For transfers to Amazon S3, you specify the specific buckets to which the data should be loaded and ensure that the account doing the loading has write permission for the buckets. You should also specify the access control list to be applied to each object loaded to Amazon S3.

For transfers to EBS, you specify the target region for the EBS import operation. If the storage device is less than or equal to the maximum volume size of 1 TB, its contents are loaded directly into an Amazon EBS snapshot. If the storage device's capacity exceeds 1 TB, a device image is stored within the specified S3 log bucket. You can then create a RAID of Amazon EBS volumes using software such as Logical Volume Manager, and copy the image from S3 to this new volume.

For added protection, you can encrypt the data on your device before you ship it to AWS. For Amazon S3 data, you can use a PIN-code device with hardware encryption or TrueCrypt software to encrypt your data before sending it to AWS. For EBS and Glacier data, you can use any encryption method you choose, including a PIN-code device. AWS will decrypt your Amazon S3 data before importing using the PIN code and/or TrueCrypt password you supply in your import manifest. AWS uses your PIN to access a PIN-code device, but does not decrypt software-encrypted data for import to Amazon EBS or Amazon Glacier. The following table summarizes your encryption options for each type of import/export job.

Import to Amazon S3		
Source	Target	Result
- Files on a device file system - Encrypt data using PIN-code device and/or TrueCrypt before shipping device	- Objects in an existing Amazon S3 bucket - AWS decrypts the data before performing the import	- One object for each file. - AWS erases your device after every import job prior to shipping
Export from Amazon S3		
Source	Target	Result
- Objects in one or more Amazon S3 buckets - Provide a PIN code and/or password that AWS will use to encrypt your data	- Files on your storage device - AWS formats your device - AWS copies your data to an encrypted file container on your device	- One file for each object - AWS encrypts your data prior to shipping - Use PIN-code device and/or TrueCrypt to decrypt the files
Import to Amazon Glacier		
Source	Target	Result
- Entire device - Encrypt the data using the encryption method of your choice before shipping	- One archive in an existing Amazon Glacier vault - AWS does not decrypt your device	- Device image stored as a single archive - AWS erases your device after every import job prior to shipping
Import to Amazon EBS (Device Capacity < 1 TB)		
Source	Target	Result

- Entire device - Encrypt the data using the encryption method of your choice before shipping	- One Amazon EBS snapshot - AWS does not decrypt your device	- Device image is stored as a single snapshot - If the device was encrypted, the image is encrypted - AWS erases your device after every import job prior to shipping
Import to Amazon EBS (Device Capacity > 1 TB)		
Source	Target	Result
- Entire device - Encrypt the data using the encryption method of your choice before shipping	- Multiple objects in an existing Amazon S3 bucket - AWS does not decrypt your device	- Device image chunked into series of 1 TB snapshots stored as objects in Amazon S3 bucket specified in manifest file - If the device was encrypted, the image is encrypted - AWS erases your device after every import job prior to shipping

After the import is complete, AWS Import/Export will erase the contents of your storage device to safeguard the data during return shipment. AWS overwrites all writable blocks on the storage device with zeroes. You will need to repartition and format the device after the wipe. If AWS is unable to erase the data on the device, it will be scheduled for destruction and our support team will contact you using the email address specified in the manifest file you ship with the device.

When shipping a device internationally, the customs option and certain required subfields are required in the manifest file sent to AWS. AWS Import/Export uses these values to validate the inbound shipment and prepare the outbound customs paperwork. Two of these options are whether the data on the device is encrypted or not and the encryption software's classification. When shipping encrypted data to or from the United States, the encryption software must be classified as 5D992 under the United States Export Administration Regulations.

Database Services

Amazon Web Services provides a number of database solutions for developers and businesses—from managed relational and NoSQL database services, to in-memory caching as a service and petabyte-scale data-warehouse service.

Amazon DynamoDB Security


Amazon DynamoDB is a managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB enables you to offload the administrative burdens of operating and scaling distributed databases to AWS, so you don't have to worry about hardware provisioning, setup and configuration, replication, software patching, or cluster scaling.

You can create a database table that can store and retrieve any amount of data, and serve any level of request traffic. DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity you specified and the amount of data stored, while maintaining consistent, fast performance. All data items are stored on Solid State Drives (SSDs) and are automatically replicated across multiple availability zones in a region to provide built-in high availability and data durability.



You can set up automatic backups using a special template in AWS Data Pipeline that was created just for copying DynamoDB tables. You can choose full or incremental backups to a table in the same region or a different region. You can use the copy for disaster recovery (DR) in the event that an error in your code damages the original table, or to federate DynamoDB data across regions to support a multi-region application.

To control who can use the DynamoDB resources and API, you set up permissions in AWS IAM. In addition to controlling access at the resource-level with IAM, you can also control access at the database level—you can create database-level permissions that allow or deny access to items (rows) and attributes (columns) based on the needs of your application. These database-level permissions are called *fine-grained access controls*, and you create them using an IAM policy that specifies under what circumstances a user or application can access a DynamoDB table. The IAM policy can restrict access to individual items in a table, access to the attributes in those items, or both at the same time.



You can optionally use web identity federation to control access by application users who are authenticated by Login with Amazon, Facebook, or Google. Web identity federation removes the need for creating individual IAM users; instead, users can sign in to an identity provider and then obtain temporary security credentials from AWS Security Token Service (AWS STS). AWS STS returns temporary AWS credentials to the application and allows it to access the specific DynamoDB table.

In addition to requiring database and user permissions, each request to the DynamoDB service must contain a valid HMAC-SHA256 signature, or the request is rejected. The AWS SDKs automatically sign your requests; however, if you want to write your own HTTP POST requests, you must provide the signature in the header of your request to Amazon DynamoDB. To calculate the signature, you must request temporary security credentials from the AWS Security Token Service. Use the temporary security credentials to sign your requests to Amazon DynamoDB.

Amazon DynamoDB is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2.

Amazon Relational Database Service (Amazon RDS) Security

Amazon RDS allows you to quickly create a relational database (DB) instance and flexibly scale the associated compute resources and storage capacity to meet application demand. Amazon RDS manages the database instance on your behalf by performing backups, handling failover, and maintaining the database software. Currently, Amazon RDS is available for MySQL, Oracle, Microsoft SQL Server, and PostgreSQL database engines.

Amazon RDS has multiple features that enhance reliability for critical production databases, including DB security groups, permissions, SSL connections, automated backups, DB snapshots, and multi-AZ deployments. DB instances can also be deployed in an Amazon VPC for additional network isolation.

Access Control

When you first create a DB Instance within Amazon RDS, you will create a master user account, which is used only within the context of Amazon RDS to control access to your DB Instance(s). The master user account is a native database user account that allows you to log on to your DB Instance with all database privileges. You can specify the master user name and password you want associated with each DB Instance when you create the DB Instance. Once you have created your DB Instance, you can connect to the database using the master user credentials. Subsequently, you can create additional user accounts so that you can restrict who can access your DB Instance.

You can control Amazon RDS DB Instance access via DB Security Groups, which are similar to Amazon EC2 Security Groups but not interchangeable. DB Security Groups act like a firewall controlling network access to your DB Instance. Database Security Groups default to a “deny all” access mode and customers must specifically authorize network ingress. There are two ways of doing this: authorizing a network IP range or authorizing an existing Amazon EC2 Security Group. DB Security Groups only allow access to the database server port (all others are blocked) and can be updated without restarting the Amazon RDS DB Instance, which allows a customer seamless control of their database access. Using AWS IAM, you can further control access to your RDS DB instances. AWS IAM enables you to control what RDS operations each individual AWS IAM user has permission to call.

Network Isolation

For additional network access control, you can run your DB Instances in an Amazon VPC. Amazon VPC enables you to isolate your DB Instances by specifying the IP range you wish to use, and connect to your existing IT infrastructure through industry-standard encrypted IPsec VPN. Running Amazon RDS in a VPC enables you to have a DB instance within a private subnet. You can also set up a virtual private gateway that extends your corporate network into your VPC, and allows access to the RDS DB instance in that VPC. Refer to the [Amazon VPC User Guide](#) for more details.

For Multi-AZ deployments, defining a subnet for all availability zones in a region will allow Amazon RDS to create a new standby in another availability zone should the need arise. You can create DB Subnet Groups, which are collections of subnets that you may want to designate for your RDS DB Instances in a VPC. Each DB Subnet Group should have at least one subnet for every availability zone in a given region. In this case, when you create a DB Instance in a VPC, you select a DB Subnet Group; Amazon RDS then uses that DB Subnet Group and your preferred availability zone to select a subnet and an IP address within that subnet. Amazon RDS creates and associates an Elastic Network Interface to your DB Instance with that IP address.

DB Instances deployed within an Amazon VPC can be accessed from the Internet or from Amazon EC2 Instances outside the VPC via VPN or bastion hosts that you can launch in your public subnet. To use a bastion host, you will need to set up a public subnet with an EC2 instance that acts as a SSH Bastion. This public subnet must have an Internet gateway and routing rules that allow traffic to be directed via the SSH host, which must then forward requests to the private IP address of your Amazon RDS DB instance.

DB Security Groups can be used to help secure DB Instances within an Amazon VPC. In addition, network traffic entering and exiting each subnet can be allowed or denied via network ACLs. All network traffic entering or exiting your Amazon VPC via your IPsec VPN connection can be inspected by your on-premises security infrastructure, including network firewalls and intrusion detection systems.



Encryption

You can encrypt connections between your application and your DB Instance using SSL. For MySQL and SQL Server, RDS creates an SSL certificate and installs the certificate on the DB instance when the instance is provisioned. For MySQL, you launch the mysql client using the `--ssl_ca` parameter to reference the public key in order to encrypt connections. For SQL Server, download the public key and import the certificate into your Windows operating system. Oracle RDS uses Oracle native network encryption with a DB instance. You simply add the native network encryption option to an option group and associate that option group with the DB instance. Once an encrypted connection is established, data transferred between the DB Instance and your application will be encrypted during transfer. You can also require your DB instance to only accept encrypted connections.

Amazon RDS supports Transparent Data Encryption (TDE) for SQL Server (SQL Server Enterprise Edition) and Oracle (part of the Oracle Advanced Security option available in Oracle Enterprise Edition). The TDE feature automatically encrypts data before it is written to storage and automatically decrypts data when it is read from storage. If you require your MySQL data to be encrypted while “at rest” in the database, your application must manage the encryption and decryption of data.

Note that SSL support within Amazon RDS is for encrypting the connection between your application and your DB Instance; it should not be relied on for authenticating the DB Instance itself.

While SSL offers security benefits, be aware that SSL encryption is a compute intensive operation and will increase the latency of your database connection. To learn more about how SSL works with MySQL, you can refer directly to the MySQL documentation found [here](#). To learn how SSL works with SQL Server, you can read more in the [RDS User Guide](#).

Automated Backups and DB Snapshots

Amazon RDS provides two different methods for backing up and restoring your DB Instance(s): automated backups and database snapshots (DB Snapshots).

Turned on by default, the automated backup feature of Amazon RDS enables point-in-time recovery for your DB Instance. Amazon RDS will back up your database and transaction logs and store both for a user-specified retention period. This allows you to restore your DB Instance to any second during your retention period, up to the last 5 minutes. Your automatic backup retention period can be configured to up to 35 days.

During the backup window, storage I/O may be suspended while your data is being backed up. This I/O suspension typically lasts a few minutes. This I/O suspension is avoided with Multi-AZ DB deployments, since the backup is taken from the standby.

DB Snapshots are user-initiated backups of your DB Instance. These full database backups are stored by Amazon RDS until you explicitly delete them. You can copy DB snapshots of any size and move them between any of AWS’s public regions, or copy the same snapshot to multiple regions simultaneously. You can then create a new DB Instance from a DB Snapshot whenever you desire.

DB Instance Replication

Amazon cloud computing resources are housed in highly available data center facilities in different regions of the world, and each region contains multiple distinct locations called Availability Zones. Each Availability Zone is engineered to be isolated from failures in other Availability Zones, and to provide inexpensive, low-latency network connectivity to other Availability Zones in the same region.



To architect for high availability of your Oracle, PostgreSQL, or MySQL databases, you can run your RDS DB instance in several Availability Zones, an option called a Multi-AZ deployment. When you select this option, Amazon automatically provisions and maintains a synchronous standby replica of your DB instance in a different Availability Zone. The primary DB instance is synchronously replicated across Availability Zones to the standby replica. In the event of DB instance or Availability Zone failure, Amazon RDS will automatically failover to the standby so that database operations can resume quickly without administrative intervention.

For customers who use MySQL and need to scale beyond the capacity constraints of a single DB Instance for read-heavy database workloads, Amazon RDS provides a Read Replica option. Once you create a read replica, database updates on the source DB instance are replicated to the read replica using MySQL's native, asynchronous replication. You can create multiple read replicas for a given source DB instance and distribute your application's read traffic among them. Read replicas can be created with Multi-AZ deployments to gain read scaling benefits in addition to the enhanced database write availability and data durability provided by Multi-AZ deployments.

Automatic Software Patching

Amazon RDS will make sure that the relational database software powering your deployment stays up-to-date with the latest patches. When necessary, patches are applied during a maintenance window that you can control. You can think of the Amazon RDS maintenance window as an opportunity to control when DB Instance modifications (such as scaling DB Instance class) and software patching occur, in the event either are requested or required. If a "maintenance" event is scheduled for a given week, it will be initiated and completed at some point during the 30-minute maintenance window you identify.

The only maintenance events that require Amazon RDS to take your DB Instance offline are scale compute operations (which generally take only a few minutes from start-to-finish) or required software patching. Required patching is automatically scheduled only for patches that are security and durability related. Such patching occurs infrequently (typically once every few months) and should seldom require more than a fraction of your maintenance window. If you do not specify a preferred weekly maintenance window when creating your DB Instance, a 30-minute default value is assigned. If you wish to modify when maintenance is performed on your behalf, you can do so by modifying your DB Instance in the [AWS Management Console](#) or by using the `ModifyDBInstance` API. Each of your DB Instances can have different preferred maintenance windows, if you so choose.

Running your DB Instance as a Multi-AZ deployment can further reduce the impact of a maintenance event, as Amazon RDS will conduct maintenance via the following steps: 1) Perform maintenance on standby, 2) Promote standby to primary, and 3) Perform maintenance on old primary, which becomes the new standby.

When an Amazon RDS DB Instance deletion API (`DeleteDBInstance`) is run, the DB Instance is marked for deletion. Once the instance no longer indicates 'deleting' status, it has been removed. At this point the instance is no longer accessible and unless a final snapshot copy was asked for, it cannot be restored and will not be listed by any of the tools or APIs.

Event Notification

You can receive notifications of a variety of important events that can occur on your RDS instance, such as whether the instance was shut down, a backup was started, a failover occurred, the security group was changed, or your storage space is low. The Amazon RDS service groups events into categories that you can subscribe to so that you can be notified when an event in that category occurs. You can subscribe to an event category for a DB instance, DB snapshot, DB security group, or for a DB parameter group. RDS events are published via AWS SNS and sent to you as an email or text message. For more information about RDS notification event categories, refer to the [RDS User Guide](#).



Amazon Redshift Security

Amazon Redshift is a petabyte-scale SQL data warehouse service that runs on highly optimized and managed AWS compute and storage resources. The service has been architected to not only scale up or down rapidly, but to significantly improve query speeds—even on extremely large datasets. To increase performance, Redshift uses techniques such as columnar storage, data compression, and zone maps to reduce the amount of IO needed to perform queries. It also has a massively parallel processing (MPP) architecture, parallelizing and distributing SQL operations to take advantage of all available resources.

When you create a Redshift data warehouse, you provision a single-node or multi-node cluster, specifying the type and number of nodes that will make up the cluster. The node type determines the storage size, memory, and CPU of each node. Each multi-node cluster includes a leader node and two or more compute nodes. A leader node manages connections, parses queries, builds execution plans, and manages query execution in the compute nodes. The compute nodes store data, perform computations, and run queries as directed by the leader node. The leader node of each cluster is accessible through ODBC and JDBC endpoints, using standard PostgreSQL drivers. The compute nodes run on a separate, isolated network and are never accessed directly.

After you provision a cluster, you can upload your dataset and perform data analysis queries by using common SQL-based tools and business intelligence applications.

Cluster Access

By default, clusters that you create are closed to everyone. Amazon Redshift enables you to configure firewall rules (security groups) to control network access to your data warehouse cluster. You can also run Redshift inside an Amazon VPC to isolate your data warehouse cluster in your own virtual network and connect it to your existing IT infrastructure using industry-standard encrypted IPsec VPN.

The AWS account that creates the cluster has full access to the cluster. Within your AWS account, you can use AWS IAM to create user accounts and manage permissions for those accounts. By using IAM, you can grant different users permission to perform only the cluster operations that are necessary for their work.

Like all databases, you must grant permission in Redshift at the database level in addition to granting access at the resource level. Database users are named user accounts that can connect to a database and are authenticated when they login to Amazon Redshift. In Redshift, you grant database user permissions on a per-cluster basis instead of on a per-table basis. However, a user can see data only in the table rows that were generated by his own activities; rows generated by other users are not visible to him.

The user who creates a database object is its owner. By default, only a superuser or the owner of an object can query, modify, or grant permissions on the object. For users to use an object, you must grant the necessary permissions to the user or the group that contains the user. And only the owner of an object can modify or delete it.

Data Backups

Amazon Redshift distributes your data across all compute nodes in a cluster. When you run a cluster with at least two compute nodes, data on each node will always be mirrored on disks on another node, reducing the risk of data loss. In addition, all data written to a node in your cluster is continuously backed up to Amazon S3 using snapshots. Redshift stores your snapshots for a user-defined period, which can be from one to thirty-five days. You can also take your own snapshots at any time; these snapshots leverage all existing system snapshots and are retained until you explicitly delete them.



Amazon Redshift continuously monitors the health of the cluster and automatically re-replicates data from failed drives and replaces nodes as necessary. All of this happens without any effort on your part, although you may see a slight performance degradation during the re-replication process.

You can use any system or user snapshot to restore your cluster using the AWS Management Console or the Amazon Redshift APIs. Your cluster is available as soon as the system metadata has been restored and you can start running queries while user data is spooled down in the background.

Data Encryption

When creating a cluster, you can choose to encrypt it in order to provide additional protection for your data at rest. When you enable encryption in your cluster, Amazon Redshift stores all data in user-created tables in an encrypted format using hardware-accelerated AES-256 block encryption keys. This includes all data written to disk as well as any backups.

Amazon Redshift uses a four-tier, key-based architecture for encryption. These keys consist of data encryption keys, a database key, a cluster key, and a master key:

- *Data encryption keys* encrypt data blocks in the cluster. Each data block is assigned a randomly-generated AES-256 key. These keys are encrypted by using the database key for the cluster.
- The *database key* encrypts data encryption keys in the cluster. The database key is a randomly-generated AES-256 key. It is stored on disk in a separate network from the Amazon Redshift cluster and encrypted by a master key. Amazon Redshift passes the database key across a secure channel and keeps it in memory in the cluster.
- The *cluster key* encrypts the database key for the Amazon Redshift cluster. You can use either AWS or a hardware security module (HSM) to store the cluster key. HSMs provide direct control of key generation and management, and make key management separate and distinct from the application and the database.
- The *master key* encrypts the cluster key if it is stored in AWS. The master key encrypts the cluster-key-encrypted database key if the cluster key is stored in an HSM.

You can have Redshift rotate the encryption keys for your encrypted clusters at any time. As part of the rotation process, keys are also updated for all of the cluster's automatic and manual snapshots.

Note that enabling encryption in your cluster will impact performance, even though it is hardware accelerated. Encryption also applies to backups. When restoring from an encrypted snapshot, the new cluster will be encrypted as well.

To encrypt your table load data files when you upload them to Amazon S3, you can use Amazon S3 server-side encryption. When you load the data from Amazon S3, the COPY command will decrypt the data as it loads the table.

Database Audit Logging

Amazon Redshift logs all SQL operations, including connection attempts, queries, and changes to your database. You can access these logs using SQL queries against system tables or choose to have them downloaded to a secure Amazon S3 bucket. You can then use these audit logs to monitor your cluster for security and troubleshooting purposes.

Automatic Software Patching

Amazon Redshift manages all the work of setting up, operating, and scaling your data warehouse, including provisioning capacity, monitoring the cluster, and applying patches and upgrades to the Amazon Redshift engine. Patches are applied only during specified maintenance windows.

SSL Connections

To protect your data in transit within the AWS cloud, Amazon Redshift uses hardware-accelerated SSL to communicate with Amazon S3 or Amazon DynamoDB for COPY, UNLOAD, backup, and restore operations. You can encrypt the connection between your client and the cluster by specifying SSL in the parameter group associated with the cluster. To have your clients also authenticate the Redshift server, you can install the public key (.pem file) for the SSL certificate on your client and use the key to connect to your clusters.

Amazon Redshift offers the newer, stronger cipher suites that use the Elliptic Curve Diffie-Hellman Ephemeral protocol. ECDHE allows SSL clients to provide Perfect Forward Secrecy between the client and the Redshift cluster. Perfect Forward Secrecy uses session keys that are ephemeral and not stored anywhere, which prevents the decoding of captured data by unauthorized third parties, even if the secret long-term key itself is compromised. You do not need to configure anything in Amazon Redshift to enable ECDHE; if you connect from a SQL client tool that uses ECDHE to encrypt communication between the client and server, Amazon Redshift will use the provided cipher list to make the appropriate connection.

Amazon ElastiCache Security

Amazon ElastiCache is a web service that makes it easy to set up, manage, and scale distributed in-memory cache environments in the cloud. The service improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases. It can be used to significantly improve latency and throughput for many read-heavy application workloads (such as social networking, gaming, media sharing, and Q&A portals) or compute-intensive workloads (such as a recommendation engine). Caching improves application performance by storing critical pieces of data in memory for low-latency access. Cached information may include the results of I/O-intensive database queries or the results of computationally-intensive calculations.

The Amazon ElastiCache service automates time-consuming management tasks for in-memory cache environments, such as patch management, failure detection, and recovery. It works in conjunction with other Amazon Web Services (such as Amazon EC2, Amazon CloudWatch, and Amazon SNS) to provide a secure, high-performance, and managed in-memory cache. For example, an application running in Amazon EC2 can securely access an Amazon ElastiCache Cluster in the same region with very low latency.

Using the Amazon ElastiCache service, you create a Cache Cluster, which is a collection of one or more Cache Nodes, each running an instance of the Memcached service. A Cache Node is a fixed-size chunk of secure, network-attached RAM. Each Cache Node runs an instance of the Memcached service, and has its own DNS name and port. Multiple types of Cache Nodes are supported, each with varying amounts of associated memory. A Cache Cluster can be set up with a specific number of Cache Nodes and a Cache Parameter Group that controls the properties for each Cache Node. All Cache Nodes within a Cache Cluster are designed to be of the same Node Type and have the same parameter and security group settings.

Amazon ElastiCache allows you to control access to your Cache Clusters using Cache Security Groups. A Cache Security Group acts like a firewall, controlling network access to your Cache Cluster. By default, network access is turned off to



your Cache Clusters. If you want your applications to access your Cache Cluster, you must explicitly enable access from hosts in specific EC2 security groups. Once ingress rules are configured, the same rules apply to all Cache Clusters associated with that Cache Security Group.

To allow network access to your Cache Cluster, create a Cache Security Group and use the Authorize Cache Security Group Ingress API or CLI command to authorize the desired EC2 security group (which in turn specifies the EC2 instances allowed). IP-range based access control is currently not enabled for Cache Clusters. All clients to a Cache Cluster must be within the EC2 network, and authorized via Cache Security Groups.

ElastiCache for Redis provides backup and restore functionality, where you can create a snapshot of your entire Redis cluster as it exists at a specific point in time. You can schedule automatic, recurring daily snapshots or you can create a manual snapshot at any time. For automatic snapshots, you specify a retention period; manual snapshots are retained until you delete them. The snapshots are stored in Amazon S3 with high durability, and can be used for warm starts, backups, and archiving.

Application Services

Amazon Web Services offers a variety of managed services to use with your applications, including services that provide application streaming, queueing, push notification, email delivery, search, and transcoding.

Amazon CloudSearch Security

Amazon CloudSearch is a managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. It enables you to quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.

An Amazon CloudSearch domain encapsulates a collection of data you want to search, the search instances that process your search requests, and a configuration that controls how your data is indexed and searched. You create a separate search domain for each collection of data you want to make searchable. For each domain, you configure indexing options that describe the fields you want to include in your index and how you want to use them, text options that define domain-specific stopwords, stems, and synonyms, rank expressions that you can use to customize how search results are ranked, and access policies that control access to the domain's document and search endpoints.

Access to your search domain's endpoints is restricted by IP address so that only authorized hosts can submit documents and send search requests. IP address authorization is used only to control access to the document and search endpoints. All Amazon CloudSearch configuration requests must be authenticated using standard AWS authentication.

Amazon CloudSearch provides separate endpoints for accessing the configuration, search, and document services:

- The configuration service is accessed through a general endpoint: `cloudsearch.us-east-1.amazonaws.com`
- The document service endpoint is used to submit documents to the domain for indexing and is accessed through a domain-specific endpoint: <http://doc-domainname-domainid.us-east-1.cloudsearch.amazonaws.com>
- The search endpoint is used to submit search requests to the domain and is accessed through a domain-specific endpoint: <http://search-domainname-domainid.us-east-1.cloudsearch.amazonaws.com>

Note that if you do not have a static IP address, you must re-authorize your computer whenever your IP address changes. If your IP address is assigned dynamically, it is also likely that you're sharing that address with other computers on your network. This means that when you authorize the IP address, all computers that share it will be able to access your search domain's document service endpoint.

Like all AWS Services, Amazon CloudSearch requires that every request made to its control API be authenticated so only authenticated users can access and manage your CloudSearch domain. API requests are signed with an HMAC-SHA1 or HMAC-SHA256 signature calculated from the request and the user's AWS Secret Access key. Additionally, the Amazon CloudSearch control API is accessible via SSL-encrypted endpoints. You can control access to Amazon CloudSearch management functions by creating users under your AWS Account using AWS IAM, and controlling which CloudSearch operations these users have permission to perform.



Amazon Simple Queue Service (Amazon SQS) Security

Amazon SQS is a highly reliable, scalable message queuing service that enables asynchronous message-based communication between distributed components of an application. The components can be computers or Amazon EC2 instances or a combination of both. With Amazon SQS, you can send any number of messages to an Amazon SQS queue at any time from any component. The messages can be retrieved from the same component or a different one right away or at a later time (within 4 days). Messages are highly durable; each message is persistently stored in highly available, highly reliable queues. Multiple processes can read/write from/to an Amazon SQS queue at the same time without interfering with each other.

Amazon SQS access is granted based on an AWS Account or a user created with AWS IAM. Once authenticated, the AWS Account has full access to all user operations. An AWS IAM user, however, only has access to the operations and queues for which they have been granted access via policy. By default, access to each individual queue is restricted to the AWS Account that created it. However, you can allow other access to a queue, using either an SQS-generated policy or a policy you write.

Amazon SQS is accessible via SSL-encrypted endpoints. The encrypted endpoints are accessible from both the Internet and from within Amazon EC2. Data stored within Amazon SQS is not encrypted by AWS; however, the user can encrypt data before it is uploaded to Amazon SQS, provided that the application utilizing the queue has a means to decrypt the message when retrieved. Encrypting messages before sending them to Amazon SQS helps protect against access to sensitive customer data by unauthorized persons, including AWS.

Amazon Simple Notification Service (Amazon SNS) Security

Amazon Simple Notification Service (Amazon SNS) is a web service that makes it easy to set up, operate, and send notifications from the cloud. It provides developers with a highly scalable, flexible, and cost-effective capability to publish messages from an application and immediately deliver them to subscribers or other applications.

Amazon SNS provides a simple web services interface that can be used to create topics that customers want to notify applications (or people) about, subscribe clients to these topics, publish messages, and have these messages delivered over clients' protocol of choice (i.e., HTTP/HTTPS, email, etc.). Amazon SNS delivers notifications to clients using a "push" mechanism that eliminates the need to periodically check or "poll" for new information and updates. Amazon SNS can be leveraged to build highly reliable, event-driven workflows and messaging applications without the need for complex middleware and application management. The potential uses for Amazon SNS include monitoring applications, workflow systems, time-sensitive information updates, mobile applications, and many others. Amazon SNS provides access control mechanisms so that topics and messages are secured against unauthorized access. Topic owners can set policies for a topic that restrict who can publish or subscribe to a topic. Additionally, topic owners can encrypt transmission by specifying that the delivery mechanism must be HTTPS.

Amazon SNS access is granted based on an AWS Account or a user created with AWS IAM. Once authenticated, the AWS Account has full access to all user operations. An AWS IAM user, however, only has access to the operations and topics for which they have been granted access via policy. By default, access to each individual topic is restricted to the AWS Account that created it. However, you can allow other access to SNS, using either an SNS-generated policy or a policy you write.



Amazon Simple Workflow Service (Amazon SWF) Security

The Amazon Simple Workflow Service (SWF) makes it easy to build applications that coordinate work across distributed components. Using Amazon SWF, you can structure the various processing steps in an application as “tasks” that drive work in distributed applications, and Amazon SWF coordinates these tasks in a reliable and scalable manner. Amazon SWF manages task execution dependencies, scheduling, and concurrency based on a developer’s application logic. The service stores tasks, dispatches them to application components, tracks their progress, and keeps their latest state.

Amazon SWF provides simple API calls that can be executed from code written in any language and run on your EC2 instances, or any of your machines located anywhere in the world that can access the Internet. Amazon SWF acts as a coordination hub with which your application hosts interact. You create desired workflows with their associated tasks and any conditional logic you wish to apply and store them with Amazon SWF.

Amazon SWF access is granted based on an AWS Account or a user created with AWS IAM. All actors that participate in the execution of a workflow—deciders, activity workers, workflow administrators—must be IAM users under the AWS Account that owns the Amazon SWF resources. You cannot grant users associated with other AWS Accounts access to your Amazon SWF workflows. An AWS IAM user, however, only has access to the workflows and resources for which they have been granted access via policy.

Amazon Simple Email Service (Amazon SES) Security

Amazon Simple Email Service (SES) is an outbound-only email-sending service built on Amazon’s reliable and scalable infrastructure. Amazon SES helps you maximize email deliverability and stay informed of the delivery status of your emails. Amazon SES integrates with other AWS services, making it easy to send emails from applications being hosted on services such as Amazon EC2.

Unfortunately, with other email systems, it's possible for a spammer to falsify an email header and spoof the originating email address so that it appears as though the email originated from a different source. To mitigate these problems, Amazon SES requires users to verify their email address or domain in order to confirm that they own it and to prevent others from using it. To verify a domain, Amazon SES requires the sender to publish a DNS record that Amazon SES supplies as proof of control over the domain. Amazon SES periodically reviews domain verification status, and revokes verification in cases where it is no longer valid.

Amazon SES takes proactive steps to prevent questionable content from being sent, so that ISPs receive consistently high-quality email from our domains and therefore view Amazon SES as a trusted email origin. Below are some of the features that maximize deliverability and dependability for all of our senders:

- Amazon SES uses content-filtering technologies to help detect and block messages containing viruses or malware before they can be sent.
- Amazon SES maintains complaint feedback loops with major ISPs. Complaint feedback loops indicate which emails a recipient marked as spam. Amazon SES provides you access to these delivery metrics to help guide your sending strategy.
- Amazon SES uses a variety of techniques to measure the quality of each user’s sending. These mechanisms help identify and disable attempts to use Amazon SES for unsolicited mail, and detect other sending patterns that would harm Amazon SES’s reputation with ISPs, mailbox providers, and anti-spam services.

- Amazon SES supports authentication mechanisms such as Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). When you authenticate an email, you provide evidence to ISPs that you own the domain. Amazon SES makes it easy for you to authenticate your emails. If you configure your account to use Easy DKIM, Amazon SES will DKIM-sign your emails on your behalf, so you can focus on other aspects of your email-sending strategy. To ensure optimal deliverability, we recommend that you authenticate your emails.

As with other AWS services, you use security credentials to verify who you are and whether you have permission to interact with Amazon SES. For information about which credentials to use, see [Using Credentials with Amazon SES](#). Amazon SES also integrates with AWS IAM so that you can specify which Amazon SES API actions a user can perform.

If you choose to communicate with Amazon SES through its SMTP interface, you are required to encrypt your connection using TLS. Amazon SES supports two mechanisms for establishing a TLS-encrypted connection: STARTTLS and TLS Wrapper. If you choose to communicate with Amazon SES over HTTP, then all communication will be protected by TLS through Amazon SES's HTTPS endpoint. When delivering email to its final destination, Amazon SES encrypts the email content with opportunistic TLS, if supported by the receiver.

Amazon Elastic Transcoder Service Security

The Amazon Elastic Transcoder service simplifies and automates what is usually a complex process of converting media files from one format, size, or quality to another. The Elastic Transcoder service converts standard-definition (SD) or high-definition (HD) video files as well as audio files. It reads input from an Amazon S3 bucket, transcodes it, and writes the resulting file to another Amazon S3 bucket. You can use the same bucket for input and output, and the buckets can be in any AWS region. The Elastic Transcoder accepts input files in a wide variety of web, consumer, and professional formats. Output file types include the MP3, MP4, OGG, TS, WebM, HLS using MPEG-2 TS, and Smooth Streaming using fmp4 container types, storing H.264 or VP8 video and AAC, MP3, or Vorbis audio.

You'll start with one or more input files, and create transcoding jobs in a type of workflow called a transcoding pipeline for each file. When you create the pipeline you'll specify input and output buckets as well as an IAM role. Each job must reference a media conversion template called a transcoding preset, and will result in the generation of one or more output files. A preset tells the Elastic Transcoder what settings to use when processing a particular input file. You can specify many settings when you create a preset, including the sample rate, bit rate, resolution (output height and width), the number of reference and keyframes, a video bit rate, some thumbnail creation options, etc.

A best effort is made to start jobs in the order in which they're submitted, but this is not a hard guarantee and jobs typically finish out of order since they are worked on in parallel and vary in complexity. You can pause and resume any of your pipelines if necessary.

Elastic Transcoder supports the use of SNS notifications when it starts and finishes each job, and when it needs to tell you that it has detected an error or warning condition. The SNS notification parameters are associated with each pipeline. It can also use the List Jobs By Status function to find all of the jobs with a given status (e.g., "Completed") or the Read Job function to retrieve detailed information about a particular job.

Like all other AWS services, Elastic Transcoder integrates with AWS Identity and Access Management (IAM), which allows you to control access to the service and to other AWS resources that Elastic Transcoder requires, including Amazon S3 buckets and Amazon SNS topics. By default, IAM users have no access to Elastic Transcoder or to the resources that it uses. If you want IAM users to be able to work with Elastic Transcoder, you must explicitly grant them permissions.



Amazon Elastic Transcoder requires every request made to its control API be authenticated so only authenticated processes or users can create, modify, or delete their own Amazon Transcoder pipelines and presets. Requests are signed with an HMAC-SHA256 signature calculated from the request and a key derived from the user's secret key. Additionally, the Amazon Elastic Transcoder API is only accessible via SSL-encrypted endpoints.

Durability is provided by Amazon S3, where media files are redundantly stored on multiple devices across multiple facilities in an Amazon S3 region. For added protection against users accidentally deleting media files, you can use the Versioning feature in Amazon S3 to preserve, retrieve, and restore every version of every object stored in an Amazon S3 bucket. You can further protect versions using Amazon S3 Versioning's MFA Delete feature. Once enabled for an Amazon S3 bucket, each version deletion request must include the six-digit code and serial number from your multi-factor authentication device.

Amazon AppStream Security

The Amazon AppStream service provides a framework for running streaming applications, particularly applications that require lightweight clients running on mobile devices. It enables you to store and run your application on powerful, parallel-processing GPUs in the cloud and then stream input and output to any client device. This can be a pre-existing application that you modify to work with Amazon AppStream or a new application that you design specifically to work with the service.

The Amazon AppStream SDK simplifies the development of interactive streaming applications and client applications. The SDK provides APIs that connect your customers' devices directly to your application, capture and encode audio and video, stream content across the Internet in near real-time, decode content on client devices, and return user input to the application. Because your application's processing occurs in the cloud, it can scale to handle extremely large computational loads.

Amazon AppStream deploys streaming applications on Amazon EC2. When you add a streaming application through the AWS Management Console, the service creates the AMI required to host your application and makes your application available to streaming clients. The service scales your application as needed within the capacity limits you have set to meet demand. Clients using the Amazon AppStream SDK automatically connect to your streamed application.

In most cases, you'll want to ensure that the user running the client is authorized to use your application before letting him obtain a session ID. We recommend that you use some sort of entitlement service, which is a service that authenticates clients and authorizes their connection to your application. In this case, the entitlement service will also call into the Amazon AppStream REST API to create a new streaming session for the client. After the entitlement service creates a new session, it returns the session identifier to the authorized client as a single-use entitlement URL. The client then uses the entitlement URL to connect to the application. Your entitlement service can be hosted on an Amazon EC2 instance or on [AWS Elastic Beanstalk](#).

Amazon AppStream utilizes an AWS CloudFormation template that automates the process of deploying a GPU EC2 instance that has the AppStream Windows Application and Windows Client SDK libraries installed; is configured for SSH, RDC, or VPN access; and has an elastic IP address assigned to it. By using this template to deploy your standalone streaming server, all you need to do is upload your application to the server and run the command to launch it. You can then use the Amazon AppStream Service Simulator tool to test your application in standalone mode before deploying it into production.

Amazon AppStream also utilizes the STX Protocol to manage the streaming of your application from AWS to local devices. The Amazon AppStream STX Protocol is a proprietary protocol used to stream high-quality application video



over varying network conditions; it monitors network conditions and automatically adapts the video stream to provide a low-latency and high-resolution experience to your customers. It minimizes latency while syncing audio and video as well as capturing input from your customers to be sent back to the application running in AWS.

Analytics Services

Amazon Web Services provides cloud-based analytics services to help you process and analyze any volume of data, whether your need is for managed Hadoop clusters, real-time streaming data, petabyte scale data warehousing, or orchestration.

Amazon Elastic MapReduce (Amazon EMR) Security

Amazon Elastic MapReduce (Amazon EMR) is a managed web service you can use to run Hadoop clusters that process vast amounts of data by distributing the work and data among several servers. It utilizes an enhanced version of the Apache Hadoop framework running on the web-scale infrastructure of Amazon EC2 and Amazon S3. You simply upload your input data and a data processing application into Amazon S3. Amazon EMR then launches the number of Amazon EC2 instances you specify. The service begins the job flow execution while pulling the input data from Amazon S3 into the launched Amazon EC2 instances. Once the job flow is finished, Amazon EMR transfers the output data to Amazon S3, where you can then retrieve it or use it as input in another job flow.

When launching job flows on your behalf, Amazon EMR sets up two Amazon EC2 security groups: one for the master nodes and another for the slaves. The master security group has a port open for communication with the service. It also has the SSH port open to allow you to SSH into the instances, using the key specified at startup. The slaves start in a separate security group, which only allows interaction with the master instance. By default both security groups are set up to not allow access from external sources, including Amazon EC2 instances belonging to other customers. Since these are security groups within your account, you can reconfigure them using the standard EC2 tools or dashboard. To protect customer input and output datasets, Amazon EMR transfers data to and from Amazon S3 using SSL.

Amazon EMR provides several ways to control access to the resources of your cluster. You can use AWS IAM to create user accounts and roles and configure permissions that control which AWS features those users and roles can access. When you launch a cluster, you can associate an Amazon EC2 key pair with the cluster, which you can then use when you connect to the cluster using SSH. You can also set permissions that allow users other than the default Hadoop user to submit jobs to your cluster.

By default, if an IAM user launches a cluster, that cluster is hidden from other IAM users on the AWS account. This filtering occurs on all Amazon EMR interfaces—the console, CLI, API, and SDKs—and helps prevent IAM users from accessing and inadvertently changing clusters created by other IAM users. It is useful for clusters that are intended to be viewed by only a single IAM user and the main AWS account. You also have the option to make a cluster visible and accessible to all IAM users under a single AWS account.

For an additional layer of protection, you can launch the EC2 instances of your EMR cluster into an Amazon VPC, which is like launching it into a private subnet. This allows you to control access to the entire subnetwork. You can also launch the cluster into a VPC and enable the cluster to access resources on your internal network using a VPN connection. You can encrypt the input data before you upload it to Amazon S3 using any common data encryption tool. If you do encrypt the data before it's uploaded, you then need to add a decryption step to the beginning of your job flow when Amazon Elastic MapReduce fetches the data from Amazon S3.



Amazon Kinesis Security

Amazon Kinesis is a managed service designed to handle real-time streaming of big data. It can accept any amount of data, from any number of sources, scaling up and down as needed. You can use Kinesis in situations that call for large-scale, real-time data ingestion and processing, such as server logs, social media or market data feeds, and web clickstream data.

Applications read and write data records to Amazon Kinesis in *streams*. You can create any number of Kinesis streams to capture, store, and transport data. Amazon Kinesis automatically manages the infrastructure, storage, networking, and configuration needed to collect and process your data at the level of throughput your streaming applications need. You don't have to worry about provisioning, deployment, or ongoing-maintenance of hardware, software, or other services to enable real-time capture and storage of large-scale data. Amazon Kinesis also synchronously replicates data across three facilities in an AWS Region, providing high availability and data durability.

In Amazon Kinesis, data records contain a sequence number, a partition key, and a data blob, which is an un-interpreted, immutable sequence of bytes. The Amazon Kinesis service does not inspect, interpret, or change the data in the blob in any way. Data records are accessible for only 24 hours from the time they are added to an Amazon Kinesis stream, and then they are automatically discarded.

Your application is a consumer of an Amazon Kinesis stream, which typically runs on a fleet of Amazon EC2 instances. A Kinesis application uses the Amazon Kinesis Client Library to read from the Amazon Kinesis stream. The Kinesis Client Library takes care of a variety of details for you including failover, recovery, and load balancing, allowing your application to focus on processing the data as it becomes available. After processing the record, your consumer code can pass it along to another Kinesis stream; write it to an [Amazon S3](#) bucket, a [Redshift](#) data warehouse, or a [DynamoDB](#) table; or simply discard it. A connector library is available to help you integrate Kinesis with other AWS services (such as DynamoDB, Redshift, and Amazon S3) as well as third-party products like Apache Storm.

You can control logical access to Kinesis resources and management functions by creating users under your AWS Account using AWS IAM, and controlling which Kinesis operations these users have permission to perform. To facilitate running your producer or consumer applications on an Amazon EC2 instance, you can configure that instance with an IAM role. That way, AWS credentials that reflect the permissions associated with the IAM role are made available to applications on the instance, which means you don't have to use your long-term AWS security credentials. Roles have the added benefit of providing temporary credentials that expire within a short timeframe, which adds an additional measure of protection. See the [Using IAM](#) guide for more information about IAM roles.

The Amazon Kinesis API is only accessible via an SSL-encrypted endpoint (kinesis.us-east-1.amazonaws.com) to help ensure secure transmission of your data to AWS. You must connect to that endpoint to access Kinesis, but you can then use the API to direct AWS Kinesis to create a stream in any AWS Region

AWS Data Pipeline Security

The AWS Data Pipeline service helps you process and move data between different data sources at specified intervals using data-driven workflows and built-in dependency checking. When you create a pipeline, you define data sources, preconditions, destinations, processing steps, and an operational schedule. Once you define and activate a pipeline, it will run automatically according to the schedule you specified.

With AWS Data Pipeline, you don't have to worry about checking resource availability, managing inter-task dependencies, retrying transient failures/timeouts in individual tasks, or creating a failure notification system. AWS Data



Pipeline takes care of launching the AWS services and resources your pipeline needs to process your data (e.g., Amazon EC2 or EMR) and transferring the results to storage (e.g., Amazon S3, RDS, DynamoDB, or EMR).

When you use the console, AWS Data Pipeline creates the necessary IAM roles and policies, including a trusted entities list for you. IAM roles determine what your pipeline can access and the actions it can perform. Additionally, when your pipeline creates a resource, such as an EC2 instance, IAM roles determine the EC2 instance's permitted resources and actions. When you create a pipeline, you specify one IAM role that governs your pipeline and another IAM role to govern your pipeline's resources (referred to as a "resource role"), which can be the same role for both. As part of the security best practice of least privilege, we recommend that you consider the minimum permissions necessary for your pipeline to perform work and define the IAM roles accordingly.

Like most AWS services, AWS Data Pipeline also provides the option of secure (HTTPS) endpoints for access via SSL.

Deployment and Management Services

Amazon Web Services provides a variety of tools to help with the deployment and management of your applications. This includes services that allow you to create individual user accounts with credentials for access to AWS services. It also includes services for creating and updating stacks of AWS resources, deploying applications on those resources, and monitoring the health of those AWS resources. Other tools help you manage cryptographic keys using hardware security modules (HSMs) and log AWS API activity for security and compliance purposes.

AWS Identity and Access Management (AWS IAM)

AWS IAM allows you to create multiple users and manage the permissions for each of these users within your AWS Account. A user is an identity (within an AWS Account) with unique security credentials that can be used to access AWS Services. AWS IAM eliminates the need to share passwords or keys, and makes it easy to enable or disable a user's access as appropriate.

AWS IAM enables you to implement security best practices, such as least privilege, by granting unique credentials to every user within your AWS Account and only granting permission to access the AWS services and resources required for the users to perform their jobs. AWS IAM is secure by default; new users have no access to AWS until permissions are explicitly granted.

AWS IAM is also integrated with the AWS Marketplace, so that you can control who in your organization can subscribe to the software and services offered in the Marketplace. Since subscribing to certain software in the Marketplace launches an EC2 instance to run the software, this is an important access control feature. Using AWS IAM to control access to the AWS Marketplace also enables AWS Account owners to have fine-grained control over usage and software costs.

AWS IAM enables you to minimize the use of your AWS Account credentials. Once you create AWS IAM user accounts, all interactions with AWS Services and resources should occur with AWS IAM user security credentials. More information about AWS IAM is available on the AWS website: <http://aws.amazon.com/iam/>

Roles

An IAM *role* uses temporary security credentials to allow you to delegate access to users or services that normally don't have access to your AWS resources. A role is a set of permissions to access specific AWS resources, but these permissions are not tied to a specific IAM user or group. An authorized entity (e.g., mobile user, EC2 instance) assumes a role and receives temporary security credentials for authenticating to the resources defined in the role. Temporary



security credentials provide enhanced security due to their short life-span (the default expiration is 12 hours) and the fact that they cannot be reused after they expire. This can be particularly useful in providing limited, controlled access in certain situations:

- **Federated (non-AWS) User Access.** Federated users are users (or applications) who do not have AWS Accounts. With roles, you can give them access to your AWS resources for a limited amount of time. This is useful if you have non-AWS users that you can authenticate with an external service, such as Microsoft Active Directory, LDAP, or Kerberos. The temporary AWS credentials used with the roles provide identity federation between AWS and your non-AWS users in your corporate identity and authorization system.

If your organization supports SAML 2.0 (Security Assertion Markup Language 2.0), you can create trust between your organization as an identity provider (IdP) and other organizations as service providers. In AWS, you can configure AWS as the service provider and use SAML to provide your users with federated single-sign on (SSO) to the AWS Management Console or to get federated access to call AWS APIs.

Roles are also useful if you create a mobile or web-based application that accesses AWS resources. AWS resources require security credentials for programmatic requests; however, you shouldn't embed long-term security credentials in your application because they are accessible to the application's users and can be difficult to rotate. Instead, you can let users sign in to your application using Login with Amazon, Facebook, or Google, and then use their authentication information to assume a role and get temporary security credentials.

- **Cross-Account Access.** For organizations who use multiple AWS Accounts to manage their resources, you can set up roles to provide users who have permissions in one account to access resources under another account. For organizations who have personnel who only rarely need access to resources under another account, using roles helps ensure that credentials are provided temporarily, only as needed.
- **Applications Running on EC2 Instances that Need to Access AWS Resources.** If an application runs on an Amazon EC2 instance and needs to make requests for AWS resources such as Amazon S3 buckets or an DynamoDB table, it must have security credentials. Using roles instead of creating individual IAM accounts for each application on each instance can save significant time for customers who manage a large number of instances or an elastically scaling fleet using AWS Auto Scaling.

The temporary credentials include a security token, an Access Key ID, and a Secret Access Key. To give a user access to certain resources, you distribute the temporary security credentials to the user you are granting temporary access to. When the user makes calls to your resources, the user passes in the token and Access Key ID, and signs the request with the Secret Access Key. The token will not work with different access keys. How the user passes in the token depends on the API and version of the AWS product the user is making calls to. More information about temporary security credentials is available on the AWS website: <http://docs.amazonwebservices.com/STS>

The use of temporary credentials means additional protection for you because you don't have to manage or distribute long-term credentials to temporary users. In addition, the temporary credentials get automatically loaded to the target instance so you don't have to embed them somewhere unsafe like your code. Temporary credentials are automatically rotated or changed multiple times a day without any action on your part, and are stored securely by default.

More information about using IAM roles to auto-provision keys on EC2 instances is available in the *Using IAM* guide on the AWS website: <http://docs.amazonwebservices.com/IAM>

Amazon CloudWatch Security

Amazon CloudWatch is a web service that provides monitoring for AWS cloud resources, starting with Amazon EC2. It provides customers with visibility into resource utilization, operational performance, and overall demand patterns—



including metrics such as CPU utilization, disk reads and writes, and network traffic. You can set up CloudWatch alarms to notify you if certain thresholds are crossed, or to take other automated actions such as adding or removing EC2 instances if Auto-Scaling is enabled.

CloudWatch captures and summarizes utilization metrics natively for AWS resources, but you can also have other logs sent to CloudWatch to monitor. You can route your guest OS, application, and custom log files for the software installed on your EC2 instances to CloudWatch, where they will be stored in durable fashion for as long as you'd like. You can configure CloudWatch to monitor the incoming log entries for any desired symbols or messages and to surface the results as CloudWatch metrics. You could, for example, monitor your web server's log files for 404 errors to detect bad inbound links or invalid user messages to detect unauthorized login attempts to your guest OS.

Like all AWS Services, Amazon CloudWatch requires that every request made to its control API be authenticated so only authenticated users can access and manage CloudWatch. Requests are signed with an HMAC-SHA1 signature calculated from the request and the user's private key. Additionally, the Amazon CloudWatch control API is only accessible via SSL-encrypted endpoints.

You can further control access to Amazon CloudWatch by creating users under your AWS Account using AWS IAM, and controlling what CloudWatch operations these users have permission to call.

AWS CloudHSM Security

The AWS CloudHSM service provides customers with dedicated access to a hardware security module (HSM) appliance designed to provide secure cryptographic key storage and operations within an intrusion-resistant, tamper-evident device. You can generate, store, and manage the cryptographic keys used for data encryption so that they are accessible only by you. AWS CloudHSM appliances are designed to securely store and process cryptographic key material for a wide variety of uses such as database encryption, Digital Rights Management (DRM), Public Key Infrastructure (PKI), authentication and authorization, document signing, and transaction processing. They support some of the strongest cryptographic algorithms available, including AES, RSA, and ECC, and many others.

The AWS CloudHSM service is designed to be used with Amazon EC2 and VPC, providing the appliance with its own private IP within a private subnet. You can connect to CloudHSM appliances from your EC2 servers through SSL/TLS, which uses two-way digital certificate authentication and 256-bit SSL encryption to provide a secure communication channel.

Selecting CloudHSM service in the same region as your EC2 instance decreases network latency, which can improve your application performance. You can configure a client on your EC2 instance that allows your applications to use the APIs provided by the HSM, including PKCS#11, MS CAPI and Java JCA/JCE (Java Cryptography Architecture/Java Cryptography Extensions).

Before you begin using an HSM, you must set up at least one partition on the appliance. A cryptographic partition is a logical and physical security boundary that restricts access to your keys, so only you control your keys and the operations performed by the HSM. AWS has administrative credentials to the appliance, but these credentials can only be used to manage the appliance, not the HSM partitions on the appliance. AWS uses these credentials to monitor and maintain the health and availability of the appliance. AWS cannot extract your keys nor can AWS cause the appliance to perform any cryptographic operation using your keys.

The HSM appliance has both physical and logical tamper detection and response mechanisms that erase the cryptographic key material and generate event logs if tampering is detected. The HSM is designed to detect tampering if



the physical barrier of the HSM appliance is breached. In addition, after three unsuccessful attempts to access an HSM partition with HSM Admin credentials, the HSM appliance erases its HSM partitions.

When your CloudHSM subscription ends and you have confirmed that the contents of the HSM are no longer needed, you must delete each partition and its contents as well as any logs. As part of the decommissioning process, AWS zeroizes the appliance, permanently erasing all key material.

AWS CloudTrail Security

AWS CloudTrail provides a log of all requests for AWS resources within your account. For each event recorded, you can see what service was accessed, what action was performed, any parameters for the action, and who made the request. Not only can you see which one of your users or services performed an action on an AWS service, but you can see whether it was as the AWS root account user or an IAM user, or whether it was with temporary security credentials for a role or federated user.

CloudTrail basically captures information about every API call to an AWS resource, whether that call was made from the AWS Management Console, CLI, or an SDK. If the API request returned an error, CloudTrail provides the description of the error, including messages for authorization failures. It even captures AWS Management Console sign-in events, creating a log record every time an AWS account owner, a federated user, or an IAM user simply signs into the console.

Once you have enabled CloudTrail, event logs are delivered every 5 minutes to the Amazon S3 bucket of your choice. The log files are organized by AWS Account ID, region, service name, date, and time. You can configure CloudTrail so that it aggregates log files from multiple regions into a single Amazon S3 bucket. From there, you can then upload them to your favorite log management and analysis solutions to perform security analysis and detect user behavior patterns.

By default, log files are stored indefinitely. The log files are automatically encrypted using Amazon [S3's Server Side Encryption](#) and will remain in the bucket until you choose to delete or archive them. You can use Amazon S3 lifecycle configuration rules to automatically delete old log files or archive them to Amazon Glacier for additional longevity at significant savings.

Like every other AWS service, you can limit access to CloudTrail to only certain users. You can use IAM to control which AWS users can create, configure, or delete AWS CloudTrail trails as well as which users can start and stop logging. You can control access to the log files by applying IAM or Amazon S3 bucket policies. You can also add an additional layer of security by enabling [MFA Delete](#) on your Amazon S3 bucket.

Mobile Services

AWS mobile services make it easier for you to build, ship, run, monitor, optimize, and scale cloud-powered applications for mobile devices. These services also help you authenticate users to your mobile application, synchronize data, and collect and analyze application usage.

Amazon Cognito

Amazon Cognito provides identity and sync services for mobile and web-based applications. It simplifies the task of authenticating users and storing, managing, and syncing their data across multiple devices, platforms, and applications. It provides temporary, limited-privilege credentials for both authenticated and unauthenticated users without having to manage any backend infrastructure.



Cognito works with well-known identity providers like Google, Facebook, and Amazon to authenticate end users of your mobile and web applications. You can take advantage of the identification and authorization features provided by these services instead of having to build and maintain your own. Your application authenticates with one of these identity providers using the provider's SDK. Once the end user is authenticated with the provider, an OAuth or OpenID Connect token returned from the provider is passed by your application to Cognito, which returns a new Cognito ID for the user and a set of temporary, limited-privilege AWS credentials.

To begin using Amazon Cognito, you create an *identity pool* through the Amazon Cognito console. The identity pool is a store of user identity information that is specific to your AWS account. During the creation of the identity pool, you will be asked to create a new [IAM role](#) or pick an existing one for your end users. An IAM role is a set of permissions to access specific AWS resources, but these permissions are not tied to a specific IAM user or group. An authorized entity (e.g., mobile user, EC2 instance) assumes a role and receives temporary security credentials for authenticating to the AWS resources defined in the role. Temporary security credentials provide enhanced security due to their short life-span (the default expiration is 12 hours) and the fact that they cannot be reused after they expire. The role you select has an impact on which AWS services your end users will be able to access with the temporary credentials. By default, Amazon Cognito creates a new role with limited permissions – end users only have access to the Cognito Sync service and Amazon Mobile Analytics. If your application needs access to other AWS resources such as Amazon S3 or DynamoDB, you can modify your roles directly from the IAM management console.

With Amazon Cognito, there's no need to create individual AWS accounts or even IAM accounts for every one of your web/mobile app's end users who will need to access your AWS resources. In conjunction with IAM roles, mobile users can securely access AWS resources and application features, and even save data to the AWS cloud without having to create an account or log in. However, if they choose to do this later, Cognito will merge data and identification information.

Because Amazon Cognito stores data locally as well as in the service, your end users can continue to interact with their data even when they are offline. Their offline data may be stale, but anything they put into the dataset, they can immediately retrieve whether they are online or not. The client SDK manages a local SQLite store so that the application can work even when it is not connected. The SQLite store functions as a cache and is the target of all read and write operations. Cognito's sync facility compares the local version of the data to the cloud version, and pushes up or pulls down deltas as needed. Note that in order to sync data across devices, your identity pool must support authenticated identities. Unauthenticated identities are tied to the device, so unless an end user authenticates, no data can be synced across multiple devices.

With Cognito, your application communicates directly with a supported public identity provider (Amazon, Facebook, or Google) to authenticate users. Amazon Cognito does not receive or store user credentials—only the OAuth or OpenID Connect token received from the identity provider. Once Cognito receives the token, it returns a new Cognito ID for the user and a set of temporary, limited-privilege AWS credentials.

Each Cognito identity has access only to its own data in the sync store, and this data is encrypted when stored. In addition, all identity data is transmitted over HTTPS. The unique Amazon Cognito identifier on the device is stored in the appropriate secure location—on iOS for example, the Cognito identifier is stored in the iOS keychain. User data is cached in a local SQLite database within the application's sandbox; if you require additional security, you can encrypt this identity data in the local cache by implementing encryption in your application.



Amazon Mobile Analytics

Amazon Mobile Analytics is a service for collecting, visualizing, and understanding mobile application usage data. It enables you to track customer behaviors, aggregate metrics, and identify meaningful patterns in your mobile applications. Amazon Mobile Analytics automatically calculates and updates usage metrics as the data is received from client devices running your app and displays the data in the console.

You can integrate Amazon Mobile Analytics with your application without requiring users of your app to be authenticated with an identity provider (like Google, Facebook, or Amazon). For these unauthenticated users, Mobile Analytics works with Amazon Cognito to provide temporary, limited-privilege credentials. To do this, you first create an identity pool in Cognito. The identity pool will use IAM *roles*, which is a set of permissions not tied to a specific IAM user or group but which allows an entity to access specific AWS resources. The entity assumes a role and receives temporary security credentials for authenticating to the AWS resources defined in the role. By default, Amazon Cognito creates a new role with limited permissions – end users only have access to the Cognito Sync service and Amazon Mobile Analytics. If your application needs access to other AWS resources such as Amazon S3 or DynamoDB, you can modify your roles directly from the IAM management console.

You can integrate the AWS Mobile SDK for Android or iOS into your application or use the Amazon Mobile Analytics REST API to send events from any connected device or service and visualize data in the reports. The Amazon Mobile Analytics API is only accessible via an SSL-encrypted endpoint (<https://mobileanalytics.us-east-1.amazonaws.com>).

Applications

AWS applications are managed services that enable you to provide your users with secure, centralized storage and work areas in the cloud.

Amazon WorkSpaces

Amazon WorkSpaces is a managed desktop service that allows you to quickly provision cloud-based desktops for your users. Simply choose a Windows 7 bundle that best meets the needs of your users and the number of WorkSpaces that you would like to launch. Once the WorkSpaces are ready, users receive an email informing them where they can download the relevant client and log into their WorkSpace. They can then access their cloud-based desktops from a variety of endpoint devices, including PCs, laptops, and mobile devices. However, your organization's data is never sent to or stored on the end-user device because Amazon WorkSpaces uses PC-over-IP ([PCoIP](#)), which provides an interactive video stream without transmitting actual data. The PCoIP protocol compresses, encrypts, and encodes the users' desktop computing experience and transmits 'pixels only' across any standard IP network to end-user devices.

In order to access their WorkSpace, users must sign in using a set of unique credentials or their regular Active Directory credentials. When you integrate Amazon WorkSpaces with your corporate Active Directory, each WorkSpace joins your Active Directory domain and can be managed just like any other desktop in your organization. This means that you can use Active Directory Group Policies to manage your users' WorkSpaces to specify configuration options that control the desktop. If you choose not to use Active Directory or other type of on-premises directory to manage your user WorkSpaces, you can create a private cloud directory within Amazon WorkSpaces that you can use for administration.

To provide an additional layer of security, you can also require the use of multi-factor authentication upon signin in the form of a hardware or software token. Amazon WorkSpaces supports MFA using an on-premise Remote Authentication Dial In User Service (RADIUS) server or any security provider that supports RADIUS authentication. It currently supports the PAP, CHAP, MS-CHAP1, and MS-CHAP2 protocols, along with RADIUS proxies.



Each Workspace resides on its own EC2 instance within a VPC. You can create WorkSpaces in a VPC you already own or have the WorkSpaces service create one for you automatically using the WorkSpaces Quick Start option. When you use the Quick Start option, WorkSpaces not only creates the VPC, but it performs several other provisioning and configuration tasks for you, such as creating an Internet Gateway for the VPC, setting up a directory within the VPC that is used to store user and Workspace information, creating a directory administrator account, creating the specified user accounts and adding them to the directory, and creating the Workspace instances. Or the VPC can be connected to an on-premises network using a secure VPN connection to allow access to an existing on-premises Active Directory and other intranet resources. You can add a security group that you create in your Amazon VPC to all the WorkSpaces that belong to your Directory. This allows you to control network access from Amazon WorkSpaces in your VPC to other resources in your Amazon VPC and on-premises network.

Persistent storage for WorkSpaces is provided by Amazon EBS and is automatically backed up twice a day to Amazon S3. If WorkSpaces Sync is enabled on a Workspace, the folder a user chooses to sync will be continuously backed up and stored in Amazon S3. You can also use WorkSpaces Sync on a Mac or PC to sync documents to or from your Workspace so that you can always have access to your data regardless of the desktop computer you are using.

Because it's a managed service, AWS takes care of several security and maintenance tasks like daily backups and patching. Updates are delivered automatically to your WorkSpaces during a weekly maintenance window. You can control how patching is configured for a user's Workspace. By default, Windows Update is turned on, but you have the ability to customize these settings, or use an alternative patch management approach if you desire. For the underlying OS, Windows Update is enabled by default on WorkSpaces, and configured to install updates on a weekly basis. You can use an alternative patching approach or to configure Windows Update to perform updates at a time of your choosing.

You can use IAM to control who on your team can perform administrative functions like creating or deleting WorkSpaces or setting up user directories. You can also set up a Workspace for directory administration, install your favorite Active Directory administration tools, and create organizational units and Group Policies in order to more easily apply Active Directory changes for all your WorkSpaces users.

Amazon WorkDocs

Amazon WorkDocs is a managed enterprise storage and sharing service with feedback capabilities for user collaboration. Users can store any type of file in a WorkDocs folder and allow others to view and download them. Commenting and annotation capabilities work on certain file types such as MS Word, and without requiring the application that was used to originally create the file. WorkDocs notifies contributors about review activities and deadlines via email and performs versioning of files that you have synced using the WorkDocs Sync application.

User information is stored in an Active Directory-compatible network directory. You can either create a new directory in the cloud, or connect Amazon WorkDocs to your on-premises directory. When you create a cloud directory using WorkDocs' quick start setup, it also creates a directory administrator account with the administrator email as the username. An email is sent to your administrator with instructions to complete registration. The administrator then uses this account to manage your directory.

When you create a cloud directory using WorkDocs' quick start setup, it also creates and configures a VPC for use with the directory. If you need more control over the directory configuration, you can choose the standard setup, which allows you to specify your own directory domain name, as well as one of your existing VPCs to use with the directory. If you want to use one of your existing VPCs, the VPC must have an Internet gateway and at least two subnets. Each of the subnets must be in a different Availability Zone.



Using the Amazon WorkDocs Management Console, administrators can view audit logs to track file and user activity by time, IP address, and device, and choose whether to allow users to share files with others outside their organization. Users can then control who can access individual files and disable downloads of files they share.

All data in transit is encrypted using industry-standard SSL. The WorkDocs web and mobile applications and desktop sync clients transmit files directly to Amazon WorkDocs using SSL. WorkDocs users can also utilize Multi-Factor Authentication, or MFA, if their organization has deployed a Radius server. MFA uses the following factors: username, password, and methods supported by the Radius server. The protocols supported are PAP, CHAP, MS-CHAPv1, and MS-CHAPv2.

You choose the AWS Region where each WorkDocs site's files are stored. Amazon WorkDocs is currently available in the US-East (Virginia), US-West (Oregon), and EU (Ireland) AWS Regions. All files, comments, and annotations stored in WorkDocs are automatically encrypted with AES-256 encryption.

Appendix – Glossary of Terms

Access Key ID: A string that AWS distributes in order to uniquely identify each AWS user; it is an alphanumeric token associated with your Secret Access Key.

Access control list (ACL): A list of permissions or rules for accessing an object or network resource. In Amazon EC2, security groups act as ACLs at the instance level, controlling which users have permission to access specific instances. In Amazon S3, you can use ACLs to give read or write access on buckets or objects to groups of users. In Amazon VPC, ACLs act like network firewalls and control access at the subnet level.

AMI: An Amazon Machine Image (AMI) is an encrypted machine image stored in Amazon S3. It contains all the information necessary to boot instances of a customer's software.

API: Application Programming Interface (API) is an interface in computer science that defines the ways by which an application program may request services from libraries and/or operating systems.

Archive: An archive in Amazon Glacier is a file that you want to store and is a base unit of storage in Amazon Glacier. It can be any data such as a photo, video, or document. Each archive has a unique ID and an optional description.

Authentication: Authentication is the process of determining whether someone or something is, in fact, who or what it is declared to be. Not only do users need to be authenticated, but every program that wants to call the functionality exposed by an AWS API must be authenticated. AWS requires that you authenticate every request by digitally signing it using a cryptographic hash function.

Auto-Scaling: An AWS service that allows customers to automatically scale their Amazon EC2 capacity up or down according to conditions they define.

Availability Zone: Amazon EC2 locations are composed of regions and availability zones. Availability zones are distinct locations that are engineered to be insulated from failures in other availability zones and provide inexpensive, low latency network connectivity to other availability zones in the same region.

Bastion host: A computer specifically configured to withstand attack, usually placed on the external/public side of a demilitarized zone (DMZ) or outside the firewall. You can set up an Amazon EC2 instance as an SSH bastion by setting up a public subnet as part of an Amazon VPC.

Bucket: A container for objects stored in Amazon S3. Every object is contained within a bucket. For example, if the object named photos/puppy.jpg is stored in the johnsmith bucket, then it is addressable using the URL <http://johnsmith.s3.amazonaws.com/photos/puppy.jpg>.

Certificate: A credential that some AWS products use to authenticate AWS Accounts and users. Also known as an X.509 certificate. The certificate is paired with a private key.

CIDR Block: Classless Inter-Domain Routing Block of IP addresses.

Client-side encryption: Encrypting data on the client side before uploading it to Amazon S3.

CloudFormation: An AWS provisioning tool that lets customers record the baseline configuration of the AWS resources needed to run their applications so that they can provision and update them in an orderly and predictable fashion.



Cognito: An AWS service that simplifies the task of authenticating users and storing, managing, and syncing their data across multiple devices, platforms, and applications. It works with multiple existing identity providers and also supports unauthenticated guest users.

Credentials: Items that a user or process must have in order to confirm to AWS services during the authentication process that they are authorized to access the service. AWS credentials include passwords, secret access keys as well as X.509 certificates and multi-factor tokens.

Dedicated instance: Amazon EC2 instances that are physically isolated at the host hardware level (i.e., they will run on single-tenant hardware).

Digital signature: A digital signature is a cryptographic method for demonstrating the authenticity of a digital message or document. A valid digital signature gives a recipient reason to believe that the message was created by an authorized sender, and that it was not altered in transit. Digital signatures are used by customers for signing requests to AWS APIs as part of the authentication process.

Direct Connect Service: Amazon service that allows you to provision a direct link between your internal network and an AWS region using a high-throughput, dedicated connection. With this dedicated connection in place, you can then create logical connections directly to the AWS cloud (for example, to Amazon EC2 and Amazon S3) and Amazon VPC, bypassing Internet service providers in the network path.

DynamoDB Service: A managed NoSQL database service from AWS that provides fast and predictable performance with seamless scalability.

EBS: Amazon Elastic Block Store (EBS) provides block-level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance.

ElastiCache: An AWS web service that allows you to set up, manage, and scale distributed in-memory cache environments in the cloud. The service improves the performance of web applications by allowing you to retrieve information from a fast, managed, in-memory caching system, instead of relying entirely on slower disk-based databases.

Elastic Beanstalk: An AWS deployment and management tool that automates the functions of capacity provisioning, load balancing, and auto scaling for customers' applications.

Elastic IP Address: A static, public IP address that you can assign to any instance in an Amazon VPC, thereby making the instance public. Elastic IP addresses also enable you to mask instance failures by rapidly remapping your public IP addresses to any instance in the VPC.

Elastic Load Balancing: An AWS service that is used to manage traffic on a fleet of Amazon EC2 instances, distributing traffic to instances across all availability zones within a region. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits such as taking over the encryption/decryption work from EC2 instances and managing it centrally on the load balancer.

Elastic MapReduce (EMR) Service: An AWS service that utilizes a hosted Hadoop framework running on the web-scale infrastructure of Amazon EC2 and Amazon S3. Elastic MapReduce enables customers to easily and cost-effectively process extremely large quantities of data ("big data").



Elastic Network Interface: Within an Amazon VPC, an Elastic Network Interface is an optional second network interface that you can attach to an EC2 instance. An Elastic Network Interface can be useful for creating a management network or using network or security appliances in the Amazon VPC. It can be easily detached from an instance and reattached to another instance.

Endpoint: A URL that is the entry point for an AWS service. To reduce data latency in your applications, most AWS services allow you to select a regional endpoint to make your requests. Some web services allow you to use a general endpoint that doesn't specify a region; these generic endpoints resolve to the service's us-east-1 endpoint. You can connect to an AWS endpoint via HTTP or secure HTTP (HTTPS) using SSL.

Federated users: User, systems, or applications that are not currently authorized to access your AWS services, but that you want to give temporary access to. This access is provided using the AWS Security Token Service (STS) APIs.

Firewall: A hardware or software component that controls incoming and/or outgoing network traffic according to a specific set of rules. Using firewall rules in Amazon EC2, you specify the protocols, ports, and source IP address ranges that are allowed to reach your instances. These rules specify which incoming network traffic should be delivered to your instance (e.g., accept web traffic on port 80). Amazon VPC supports a complete firewall solution enabling filtering on both ingress and egress traffic from an instance. The default group enables inbound communication from other members of the same group and outbound communication to any destination. Traffic can be restricted by any IP protocol, by service port, as well as source/destination IP address (individual IP or Classless Inter-Domain Routing (CIDR) block).

Guest OS: In a virtual machine environment, multiple operating systems can run on a single piece of hardware. Each one of these instances is considered a guest on the host hardware and utilizes its own OS.

Hash: A cryptographic hash function is used to calculate a digital signature for signing requests to AWS APIs. A cryptographic hash is a one-way function that returns a unique hash value based on the input. The input to the hash function includes the text of your request and your secret access key. The hash function returns a hash value that you include in the request as your signature.

HMAC-SHA1/HMAC-SHA256: In cryptography, a keyed-Hash Message Authentication Code (HMAC or KMAC), is a type of message authentication code (MAC) calculated using a specific algorithm involving a cryptographic hash function in combination with a secret key. As with any MAC, it may be used to simultaneously verify both the data integrity and the authenticity of a message. Any iterative cryptographic hash function, such as SHA-1 or SHA-256, may be used in the calculation of an HMAC; the resulting MAC algorithm is termed HMAC-SHA1 or HMAC-SHA256 accordingly. The cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function, on the size and quality of the key and the size of the hash output length in bits.

Hardware security module (HSM): An HSM is an appliance that provides secure cryptographic key storage and operations within a tamper-resistant hardware device. HSMs are designed to securely store cryptographic key material and use the key material without exposing it outside the cryptographic boundary of the appliance. The AWS CloudHSM service provides customers with dedicated, single-tenant access to an HSM appliance.

Hypervisor: A hypervisor, also called Virtual Machine Monitor (VMM), is computer software/hardware platform virtualization software that allows multiple operating systems to run on a host computer concurrently.



Identity and Access Management (IAM): AWS IAM enables you to create multiple users and manage the permissions for each of these users within your AWS Account.

Identity pool: A store of user identity information in Amazon Cognito that is specific to your AWS Account. Identity pools use IAM roles, which are permissions that are not tied to a specific IAM user or group and that use temporary security credentials for authenticating to the AWS resources defined in the role.

Identity Provider: An online service responsible for issuing identification information for users who would like to interact with the service or with other cooperating services. Examples of identity providers include Facebook, Google, and Amazon.

Import/Export Service: An AWS service for transferring large amounts of data to Amazon S3 or EBS storage by physically shipping a portable storage device to a secure AWS facility.

Instance: An instance is a virtualized server, also known as a virtual machine (VM), with its own hardware resources and guest OS. In EC2, an instance represents one running copy of an Amazon Machine Image (AMI).

IP address: An Internet Protocol (IP) address is a numerical label that is assigned to devices participating in a computer network utilizing the Internet Protocol for communication between its nodes.

IP spoofing: Creation of IP packets with a forged source IP address, called spoofing, with the purpose of concealing the identity of the sender or impersonating another computing system.

Key: In cryptography, a key is a parameter that determines the output of a cryptographic algorithm (called a hashing algorithm). A key pair is a set of security credentials you use to prove your identity electronically and consists of a public key and a private key.

Key rotation: The process of periodically changing the cryptographic keys used for encrypting data or digitally signing requests. Just like changing passwords, rotating keys minimizes the risk of unauthorized access if an attacker somehow obtains your key or determines the value of it. AWS supports multiple concurrent access keys and certificates, which allows customers to rotate keys and certificates into and out of operation on a regular basis without any downtime to their application.

Mobile Analytics: An AWS service for collecting, visualizing, and understanding mobile application usage data. It enables you to track customer behaviors, aggregate metrics, and identify meaningful patterns in your mobile applications.

Multi-factor authentication (MFA): The use of two or more authentication factors. Authentication factors include something you know (like a password) or something you have (like a token that generates a random number). AWS IAM allows the use of a six-digit single-use code in addition to the user name and password credentials. Customers get this single-use code from an authentication device that they keep in their physical possession (either a physical token device or a virtual token from their smart phone).

Network ACLs: Stateless traffic filters that apply to all traffic inbound or outbound from a subnet within an Amazon VPC. Network ACLs can contain ordered rules to allow or deny traffic based upon IP protocol, by service port, as well as source/destination IP address.

Object: The fundamental entities stored in Amazon S3. Objects consist of object data and metadata. The data portion is opaque to Amazon S3. The metadata is a set of name-value pairs that describe the object. These include some default metadata such as the date last modified and standard HTTP metadata such as Content-Type. The developer can also specify custom metadata at the time the Object is stored.

Paravirtualization: In computing, paravirtualization is a virtualization technique that presents a software interface to virtual machines that is similar but not identical to that of the underlying hardware.

Peering: A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IP addresses. Instances in either VPC can communicate with each other as if they are within the same network.

Port scanning: A port scan is a series of messages sent by someone attempting to break into a computer to learn which computer network services, each associated with a "well-known" port number, the computer provides.

Region: A named set of AWS resources in the same geographical area. Each region contains at least two availability zones.

Replication: The continuous copying of data from a database in order to maintain a second version of the database, usually for disaster recovery purposes. Customers can use multiple AZs for their Amazon RDS database replication needs, or use Read Replicas if using MySQL.

Relational Database Service (RDS): An AWS service that allows you to create a relational database (DB) instance and flexibly scale the associated compute resources and storage capacity to meet application demand. Amazon RDS is available for MySQL, Oracle, or Microsoft SQL Server database engines.

Role: An entity in AWS IAM that has a set of permissions that can be assumed by another entity. Use roles to enable applications running on your Amazon EC2 instances to securely access your AWS resources. You grant a specific set of permissions to a role, use the role to launch an Amazon EC2 instance, and let EC2 automatically handle AWS credential management for your applications that run on Amazon EC2.

Route 53: An authoritative DNS system that provides an update mechanism that developers can use to manage their public DNS names, answering DNS queries and translating domain names into IP address so computers can communicate with each other.

Secret Access Key: A key that AWS assigns to you when you sign up for an AWS Account. To make API calls or to work with the command line interface, each AWS user needs the Secret Access Key and Access Key ID. The user signs each request with the Secret Access Key and includes the Access Key ID in the request. To help ensure the security of your AWS Account, the Secret Access Key is accessible only during key and user creation. You must save the key (for example, in a text file that you store securely) if you want to be able to access it again.

Security group: A security group gives you control over the protocols, ports, and source IP address ranges that are allowed to reach your Amazon EC2 instances; in other words, it defines the firewall rules for your instance. These rules specify which incoming network traffic should be delivered to your instance (e.g., accept web traffic on port 80).

Security Token Service (STS): The AWS STS APIs return temporary security credentials consisting of a security token, an Access Key ID, and a Secret Access Key. You can use STS to issue security credentials to users who need temporary



access to your resources. These users can be existing IAM users, non-AWS users (federated identities), systems, or applications that need to access your AWS resources.

Server-side encryption (SSE): An option for Amazon S3 storage for automatically encrypting data at rest. With Amazon S3 SSE, customers can encrypt data on upload simply by adding an additional request header when writing the object. Decryption happens automatically when data is retrieved.

Service: Software or computing ability provided across a network (e.g., Amazon EC2, Amazon S3).

Shard: In Amazon Kinesis, a shard is a uniquely identified group of data records in an Amazon Kinesis stream. A Kinesis stream is composed of multiple shards, each of which provides a fixed unit of capacity.

Signature: Refers to a digital signature, which is a mathematical way to confirm the authenticity of a digital message. AWS uses signatures calculated with a cryptographic algorithm and your private key to authenticate the requests you send to our web services.

Simple Data Base (Simple DB): A non-relational data store that allows AWS customers to store and query data items via web services requests. Amazon SimpleDB creates and manages multiple geographically distributed replicas of the customer's data automatically to enable high availability and data durability.

Simple Email Service (SES): An AWS service that provides a scalable bulk and transactional email-sending service for businesses and developers. In order to maximize deliverability and dependability for senders, Amazon SES takes proactive steps to prevent questionable content from being sent, so that ISPs view the service as a trusted email origin.

Simple Mail Transfer Protocol (SMTP): An Internet standard for transmitting email across IP networks, SMTP is used by the Amazon Simple Email Service. Customers who used Amazon SES can use an SMTP interface to send email, but must connect to an SMTP endpoint via TLS.

Simple Notification Service (SNS): An AWS service that makes it easy to set up, operate, and send notifications from the cloud. Amazon SNS provides developers with the ability to publish messages from an application and immediately deliver them to subscribers or other applications.

Simple Queue Service (SQS): A scalable message queuing service from AWS that enables asynchronous message-based communication between distributed components of an application. The components can be computers or Amazon EC2 instances or a combination of both.

Simple Storage Service (Amazon S3): An AWS service that provides secure storage for object files. Access to objects can be controlled at the file or bucket level and can further restricted based on other conditions such as request IP source, request time, etc. Files can also be encrypted automatically using AES-256 encryption.

Simple Workflow Service (SWF): An AWS service that allows customers to build applications that coordinate work across distributed components. Using Amazon SWF, developers can structure the various processing steps in an application as "tasks" that drive work in distributed applications. Amazon SWF coordinates these tasks, managing task execution dependencies, scheduling, and concurrency based on a developer's application logic.

Single sign-on: The capability to log in once but access multiple applications and systems. A secure single sign-on capability can be provided to your federated users (AWS and non-AWS users) by creating a URL that passes the temporary security credentials to the AWS Management Console.



Snapshot: A customer-initiated backup of an EBS volume that is stored in Amazon S3, or a customer-initiated backup of an RDS database that is stored in Amazon RDS. A snapshot can be used as the starting point for a new EBS volume or Amazon RDS database or to protect the data for long-term durability and recovery.

Secure Sockets Layer (SSL): A cryptographic protocol that provides security over the Internet at the Application Layer. Both the TLS 1.0 and SSL 3.0 protocol specifications use cryptographic mechanisms to implement the security services that establish and maintain a secure TCP/IP connection. The secure connection prevents eavesdropping, tampering, or message forgery. You can connect to an AWS endpoint via HTTP or secure HTTP (HTTPS) using SSL.

Stateful firewall: In computing, a stateful firewall (any firewall that performs stateful packet inspection (SPI) or stateful inspection) is a firewall that keeps track of the state of network connections (such as TCP streams, UDP communication) traveling across it.

Storage Gateway: An AWS service that securely connects a customer's on-premises software appliance with Amazon S3 storage by using a VM that the customer deploys on a host in their data center running VMware ESXi Hypervisor. Data is asynchronously transferred from the customer's on-premises storage hardware to AWS over SSL, and then stored encrypted in Amazon S3 using AES-256.

Temporary security credentials: AWS credentials that provide temporary access to AWS services. Temporary security credentials can be used to provide identity federation between AWS services and non-AWS users in your own identity and authorization system. Temporary security credentials consist of security token, an Access Key ID, and a Secret Access Key.

Transcoder: A system that transcodes (converts) a media file (audio or video) from one format, size, or quality to another. Amazon Elastic Transcoder makes it easy for customers to transcode video files in a scalable and cost-effective fashion.

Transport Layer Security (TLS): A cryptographic protocol that provides security over the Internet at the Application Layer. Customers who used Amazon's Simple Email Service must connect to an SMTP endpoint via TLS.

Tree hash: A tree hash is generated by computing a hash for each megabyte-sized segment of the data, and then combining the hashes in tree fashion to represent ever-growing adjacent segments of the data. Glacier checks the hash against the data to help ensure that it has not been altered en route.

Vault: In Amazon Glacier, a vault is a container for storing archives. When you create a vault, you specify a name and select an AWS region where you want to create the vault. Each vault resource has a unique address.

Versioning: Every object in Amazon S3 has a key and a version ID. Objects with the same key, but different version IDs can be stored in the same bucket. Versioning is enabled at the bucket layer using PUT Bucket versioning.

Virtual Instance: Once an AMI has been launched, the resulting running system is referred to as an instance. All instances based on the same AMI start out identical and any information on them is lost when the instances are terminated or fail.

Virtual MFA: The capability for a user to get the six-digit, single-use MFA code from their smart phone rather than from a token/fob. MFA is the use of an additional factor (the single-use code) in conjunction with a user name and password for authentication.



Virtual Private Cloud (VPC): An AWS service that enables customers to provision an isolated section of the AWS cloud, including selecting their own IP address range, defining subnets, and configuring routing tables and network gateways.

Virtual Private Network (VPN): The capability to create a private, secure network between two locations over a public network such as the Internet. AWS customers can add an IPsec VPN connection between their Amazon VPC and their data center, effectively extending their data center to the cloud while also providing direct access to the Internet for public subnet instances in their Amazon VPC. In this configuration, customers add a VPN appliance on their corporate data center side.

WorkSpaces: An AWS managed desktop service that enables you to provision cloud-based desktops for your users and allows them to sign in using a set of unique credentials or their regular Active Directory credentials.

X.509: In cryptography, X.509 is a standard for a Public Key Infrastructure (PKI) for single sign-on and Privilege Management Infrastructure (PMI). X.509 specifies standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm. Some AWS products use X.509 certificates instead of a Secret Access Key for access to certain interfaces. For example, Amazon EC2 uses a Secret Access Key for access to its Query interface, but it uses a signing certificate for access to its SOAP interface and command line tool interface.

WorkDocs: An AWS managed enterprise storage and sharing service with feedback capabilities for user collaboration.

Changes since last version (Nov 2014):

- Updated compliance programs
- Updated shared security responsibility model
- Updated AWS Account security features
- Reorganized services into categories
- Updated several services with new features: CloudWatch, CloudTrail, CloudFront, EBS, ElastiCache, Redshift, Route 53, S3, Trusted Advisor, and WorkSpaces
- Added Cognito Security
- Added Mobile Analytics Security
- Added WorkDocs Security

Changes since last version (Nov 2013):

- Updated regions
- Updated several services with new features: CloudFront, DirectConnect, DynamoDB, EBS, ELB, EMR, Glacier, IAM, OpsWorks, RDS, Redshift, Route 53, Storage Gateway, and VPC
- Added AppStream Security
- Added CloudTrail Security
- Added Kinesis Security
- Added WorkSpaces Security

Changes since last version (May/June 2013):

- Updated IAM to incorporate roles and API access
- Updated MFA for API access for customer-specified privileged actions
- Updated RDS to add event notification, multi-AZ, and SSL to SQL Server 2012
- Updated VPC to add multiple IP addresses, static routing VPN, and VPC By Default
- Updated several other services with new features: CloudFront, CloudWatch, EBS, ElastiCache, Elastic Beanstalk, Route 53, S3, Storage Gateway
- Added Glacier Security
- Added Redshift Security
- Added Data Pipeline Security
- Added Transcoder Security
- Added Trusted Advisor Security
- Added OpsWorks Security
- Added CloudHSM Security

Changes since last version (May 2011):

- Reorganization to better identify infrastructure versus service-specific security
- Changed Control Environment Summary heading to AWS Compliance Program
- Changed Information and Communication heading to Management and Communication
- Changed Employee Lifecycle heading to Logical Access
- Changed Configuration Management heading to Change Management
- Merged Environmental Safeguards section with Physical Security section
- Incorporated information in Backups section into S3, SimpleDB, and EBS sections



- Update to certifications to reflect SAS70 name change to SSAE 16 and addition of FedRAMP
- Update to Network Security section to add Secure Network Architecture and Network Monitoring and Protection
- Update to IAM to incorporate roles/key provisioning, virtual MFA, temporary security credentials, and single sign on
- Update to regions to include new regions and GovCloud description
- Updated EBS, S3, SimpleDB, RDS, and EMR to clarify service and security descriptions
- Update to VPC to add configuration options, VPN, and Elastic Network Interfaces
- Addition of Amazon Direct Connect Security section
- Addition of Amazon Elastic Load Balancing Security
- Addition of AWS Storage Gateway Security
- Addition of AWS Import/Export Security
- Addition of Auto Scaling Security
- Addition of Amazon DynamoDB Security
- Addition of Amazon ElastiCache Security
- Addition of Amazon Simple Workflow Service (Amazon SWS) Security
- Addition of Amazon Simple Email Service (Amazon SES) Security
- Addition of Amazon Route 53 Security
- Addition of Amazon CloudSearch Security
- Addition of AWS Elastic Beanstalk Security
- Addition of AWS CloudFormation Security
- Updated glossary

Changes since last version (Aug 2010):

- Addition of AWS Identity and Access Management (AWS IAM)
- Addition of Amazon Simple Notification Service (SNS) Security
- Addition of Amazon CloudWatch Security
- Addition of Auto Scaling Security
- Update to Amazon Virtual Private Cloud (Amazon VPC)
- Update to Control Environment
- Removal of Risk Management because it has been expanded in a separate whitepaper

Changes since last version (Nov 2009):

- Major revision

Changes since last version (June 2009):

- Change to Certifications and Accreditations section to reflect SAS70
- Addition of Amazon Virtual Private Cloud (Amazon VPC)
- Addition of Security Credentials section to highlight AWS Multi-Factor Authentication and Key Rotation
- Addition of Amazon Relational Database Service (Amazon RDS) Security

Changes since last version (Sep 2008):

- Addition of Security Design Principles
- Update of Physical Security information and inclusion of background checks
- Backup section updated for clarity with respect to Amazon EBS



- Update of Amazon EC2 Security section to include:
- Certificate-based SSHv2
- Multi-tier security group detail and diagram
- Hypervisor description and Instance Isolation diagram
- Fault Separation
- Addition of Configuration Management
- Amazon S3 section updated for detail and clarity
- Addition of Storage Device Decommissioning
- Addition of Amazon SQS Security
- Addition of Amazon CloudFront Security
- Addition of Amazon Elastic MapReduce Security

Notices

© 2010-2015 Amazon.com, Inc., or its affiliates. This document is provided for informational purposes only. It represents AWS's current product offerings as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.



Data Processing Amendment to Google Apps Agreement

The Customer agreeing to these terms ("**Customer**") and Google Inc., Google Ireland Limited, Google Commerce Limited or Google Asia Pacific Pte. Ltd. (as applicable, "**Google**") have entered into a Google Apps for Work Agreement, Google Apps Enterprise Agreement, Google Apps for Business Agreement, Google Apps for Work via Reseller Agreement, Google Apps Enterprise via Reseller Agreement, Google Apps for Business via Reseller Agreement, Google Apps for Education Agreement or Google Apps for Education via Reseller Agreement, as applicable (as amended to date, the "**Google Apps Agreement**"). This amendment (the "**Data Processing Amendment**") is entered into by Customer and Google as of the Amendment Effective Date and amends the Google Apps Agreement.

The "**Amendment Effective Date**" is: (a) if this Data Processing Amendment is incorporated into the Google Apps Agreement by reference, the effective date of the Google Apps Agreement, as defined in that agreement; or (b) if this Data Processing Amendment is not incorporated into the Google Apps Agreement by reference, the date Customer accepts this Data Processing Amendment by clicking to accept these terms.

If this Data Processing Amendment is not incorporated into the Google Apps Agreement by reference and you are accepting on behalf of Customer, you represent and warrant that: (i) you have full legal authority to bind your employer, or the applicable entity, to these terms; (ii) you have read and understand these terms; and (iii) you agree, on behalf of the party you represent, to this Data Processing Amendment. If you do not have the legal authority to bind Customer, please do not click the "I Accept" button.

1. **Introduction.**

This Data Processing Amendment reflects the parties' agreement with respect to terms governing the processing of Customer Data under the Google Apps Agreement.

2. **Definitions.**

2.1. Capitalized terms used but not defined in this Data Processing Amendment have the meanings given in the Google Apps Agreement. In this Data Processing Amendment, unless expressly stated otherwise:

"**Additional Products**" means products, services and applications that are not part of the Services but that may be accessible, via the Admin Console or otherwise, for use with the Services.

"**Advertising**" means online advertisements displayed by Google to End Users, excluding any advertisements Customer expressly chooses to have Google or any Google Affiliate display in connection with the Services under a separate agreement (for example, Google AdSense advertisements implemented by Customer on a website created by Customer using the "Google Sites" functionality within the Services).

"**Affiliate**" means any entity controlling, controlled by, or under common control with a party, where "control" is defined as (a) the ownership of at least fifty percent (50%) of the equity or beneficial interests of the entity; (b) the right to vote for or appoint a majority of the board of directors or other governing body of the entity; or (c) the power to exercise a controlling influence over the management or policies of the entity.

“Agreement” means the Google Apps Agreement, as amended by this Data Processing Amendment and as may be further amended from time to time in accordance with the Google Apps Agreement.

“Customer Data” means data (which may include personal data and the categories of data referred to in Appendix 1) submitted, stored, sent or received via the Services by Customer, its Affiliates or End Users.

“Data Incident” means (a) any unlawful access to Customer Data stored in the Services or systems, equipment or facilities of Google or its Sub processors, or (b) unauthorized access to such Services, systems, equipment or facilities that results in loss, disclosure or alteration of Customer Data.

“Data Privacy Officer” means Google’s Data Privacy Officer for Apps.

“Data Protection legislation” means, as applicable: (a) any national provisions adopted pursuant to the Directive that are applicable to Customer and or any Customer Affiliates as the controller(s) of the Customer Data; and or (b) the federal Data Protection Act of June 1, 2002 (Switzerland).

“Directive” means Directive 95/46/EC of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the free Movement of Such Data.

“EEA” means the European Economic Area.

“Google Group” means those Google Affiliates involved in provision of the Services to Customer.

“Instructions” means Customer’s written instructions to Google consisting of the Agreement, including instructions to Google to provide the Services and technical support for the Services as set out in the Agreement; instructions given by Customer, its Affiliates and End Users via the Admin Console and otherwise in its and their use of the Services and related technical support services; and any subsequent written instructions given by Customer to Google and acknowledged by Google.

“Model Contract Clauses” or **“MCCs”** means the standard contractual clauses (processors) for the purposes of Article 2 (2) of Directive 95/46/EC for the transfer of personal data to processors established in third countries which do not ensure an adequate level of data protection.

“Safe Harbor or Certification” means a current certification to the U.S. Department of Commerce Safe Harbor framework requirements as set out at the following URL: http://export.gov/safeharbor/eu/eg_main_01_5.asp, or any replacement framework or URL from time to time.

“Services” means, for purposes of this Data Processing Amendment, the Google Apps for Work Services which are described at www.google.com/apps/intl/en/terms/user/features.html (as such services and URL link may be updated or modified by Google from time to time in accordance with the Google Apps Agreement).

“**u processors**” means (a) all Google Group entities that have logical access to and process Customer Data (each, a “**Google Group u processor**”); and (b) all third parties (other than Google Group entities) that are engaged to provide services to Customer and that have logical access to and process Customer Data (each, a “**ird Part u processor**”).

“**erm**” means the term of the Google Apps Agreement, as defined in that agreement.

“**ird Part Auditor**” means a qualified and independent third party auditor, whose then current identity Google will disclose to Customer.

2.2. The terms “personal data”, “processing”, “data subject”, “controller” and “processor” have the meanings given to them in the Directive. The terms “data importer” and “data exporter” have the meanings given to them in the Model Contract Clauses.

. erm.

This Data Processing Amendment will take effect on the Amendment Effective Date and, notwithstanding expiry or termination of the Google Apps Agreement, will remain in effect until, and automatically terminate upon, deletion by Google of all data as described in Section (Data Deletion) of this Data Processing Amendment.

. Data Protection egislation.

The parties agree and acknowledge that the Data Protection Legislation may apply to the processing of Customer Data.

5. Processing of Customer Data.

5.1. **Controller and Processor.** If the Data Protection Legislation applies to the processing of Customer Data, then as between the parties, the parties acknowledge and agree that: (a) Customer is the controller of Customer Data under the Agreement; (b) Google is a processor of such data; (c) Customer will comply with its obligations as a controller under the Data Protection Legislation; and (d) Google will comply with its obligations as a processor under the Agreement. If under the Data Protection Legislation a Customer Affiliate is considered the controller (either alone or jointly with the Customer) with respect to certain Customer Data, Customer represents and warrants to Google that Customer is authorized (i) to give the Instructions to Google and otherwise act on behalf of such Customer Affiliate in relation to such Customer Data as described in this Data Processing Amendment, and (ii) to bind the Customer Affiliate to the terms of this Data Processing Amendment.

5.2. **cope of Processing.** Google will only process Customer Data in accordance with the Instructions, and will not process Customer Data for any other purpose.

5. **. Processing estrictions.** otwithstanding any other term of the Agreement, Google will not process Customer Data for Advertising purposes or serve Advertising in the Services.

5. **. Additional Products.** Customer acknowledges that if it installs, uses, or enables Additional Products, the Services may allow such Additional Products to access Customer Data as required for the interoperation of those Additional Products with the Services. This Data Processing Amendment does not apply to the processing of data transmitted to or from such Additional Products. Customer can enable or disable Additional Products. Customer is not required to use Additional Products in order to use the Services.

. Data ecurit ecurit Compliance Audits.

.1. **Security Measures**. Google will take and implement appropriate technical and organizational measures to protect Customer Data against accidental or unlawful destruction or accidental loss or alteration or unauthorized disclosure or access or other unauthorized processing, as detailed in Appendix 2 (“**Security Measures**”). Google may update or modify the Security Measures from time to time provided that such updates and modifications do not result in the degradation of the overall security of the Services. Customer agrees that it is solely responsible for its use of the Services, including securing its account authentication credentials, and that Google has no obligation to protect Customer Data that Customer elects to store or transfer outside of Google’s and its Subprocessors’ systems (e.g., offline or on premise storage).

.2. **Security Compliance** **Google staff**. Google will take appropriate steps to ensure compliance with the Security Measures by its employees, contractors and Subprocessors to the extent applicable to their scope of performance.

. **Data Incidents**. If Google becomes aware of a Data Incident, Google will promptly notify Customer of the Data Incident, and take reasonable steps to minimize harm and secure Customer Data. Notification(s) of any Data Incident(s) will be delivered to the Notification Email Address provided by Customer in connection with the Agreement or, at Google’s discretion, by direct communication (e.g., by phone call or an in person meeting). Customer acknowledges that it is solely responsible for ensuring the contact information given for purposes of the Notification Email Address is current and valid, and for fulfilling any third party notification obligations. Customer agrees that “Data Incidents” do not include: (i) unsuccessful access attempts or similar events that do not compromise the security or privacy of Customer Data, including pings, port scans, denial of service attacks and other network attacks on firewalls or networked systems; or (ii) accidental loss or disclosure of Customer Data caused by Customer’s use of the Services or Customer’s loss of account authentication credentials. Google’s obligation to report or respond to a Data Incident under this Section will not be construed as an acknowledgement by Google of any fault or liability with respect to the Data Incident.

. **Compliance with Security and Privacy standards** **Cloud reports**. During the Term, Google will maintain the following:

(a) its ISO/IEC 27001:2011 Certification or a comparable certification (“**ISO/IEC 27001:2011 Certification**”) for the Services;

(b) conformity of the Services with ISO/IEC 27001:2011 or a comparable standard (“**ISO/IEC 27001:2011 Conformity**”), as independently verified;

(c) its confidential Service Organization Control (SOC) 2 Report (or a comparable report) on Google’s systems examining logical security controls, physical security controls, and system availability as related to the Services (the “**SOC 2 Report**”), as produced by the Third Party Auditor and updated at least once every eighteen (18) months; and

(d) its Service Organization Control (SOC) Report (or a comparable report) as related to the Services (the “**SOC Report**”), as produced by the Third Party Auditor and updated at least once every eighteen (18) months.

.5. **Auditing Security Compliance**

.5.1. Review of Security Documentation. Google will make the following available for review by Customer:

- (a) the certificate issued in relation to Google's ISO 27001 Certification;
- (b) the then current SOC 2 Report;
- (c) a summary or redacted version of the then current confidential SOC 2 Report; and
- (d) following a request by Customer in accordance with Section 5.2 below, the then current confidential SOC 2 Report.

.5.2. Customer Audits. If Customer (or an authorized Customer Affiliate) has entered into Model Contract Clauses as described in Section 10.2 of this Data Processing Amendment, Customer or such Customer Affiliate may exercise the audit rights granted under clauses 5(f) and 12(2) of such Model Contract Clauses:

- (a) by instructing Google to execute the audit as described in Sections 5.1 and 5.1 above; and or
- (b) following a request by Customer in accordance with Section 5.2 below, by executing an audit as described in such Model Contract Clauses.

.5.3. Additional Business Terms for Review and Audits. Google and Customer (or an authorized Customer Affiliate if applicable) will discuss and agree in advance on:

- (a) the reasonable date(s) of and security and confidentiality controls applicable to any Customer review under Section 5.1(d); and
- (b) the identity of a suitably qualified and independent third party auditor for any audit under Section 5.2(b), and the reasonable start date, scope and duration of and security and confidentiality controls applicable to any such audit.

Google reserves the right to charge a fee (based on Google's reasonable costs) for any review under Section 5.1(d) and or audit under Section 5.2(b). For clarity, Google is not responsible for any costs incurred or fees charged by any third party auditor appointed by Customer (or an authorized Customer Affiliate) in connection with an audit under Section 5.2(b). Nothing in this Section 5.3 varies or modifies any rights or obligations of Customer (or any authorized Customer Affiliate) or Google Inc. under any Model Contract Clauses entered into as described in Section 10.2 (Transfers of Data out of the EEA) of this Data Processing Amendment.

.5.4. Requests for Review and Audits. Any requests under Section 5.1 or 5.2 must be sent to the Data Privacy Officer as described in Section 1.1 (Data Privacy Officer) of this Data Processing Amendment.

6. Data Deletion.

6.1. Deletion of Customer and End Users. During the Term, Google will provide Customer or End Users with the ability to delete Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. Once Customer or End User deletes Customer Data and such Customer Data cannot be recovered by the Customer or End User, such

as from the “trash” (“**Customer-Deleted Data**”), Google will delete such data from its systems as soon as soon as reasonably practicable within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

7.2. **Deletion on Standard Termination**. On expiry or termination of the Google Apps Agreement (or, if applicable, on expiry of any post-termination period during which Google may agree to continue providing the Services), Google will, subject to Section 7.3 (Deletion on Termination for Non-Payment or No Purchase) below, delete all Customer-Deleted Data from its systems as soon as reasonably practicable within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

7.3. **Deletion on Termination for Non-Payment or No Purchase**. On termination of the Google Apps Agreement due to Customer breaching its payment obligations or opting not to purchase the Services at the end of a free trial of the Services, Google will delete all Customer Data from its systems within a maximum period of 180 days, unless applicable legislation or legal process prevents it from doing so.

8. **Access to Data.**

8.1. **Access; Export of Data**. During the Term, Google will provide Customer with access to and the ability to correct, block and export Customer Data in a manner consistent with the functionality of the Services and in accordance with the terms of the Agreement. To the extent Customer, in its use and administration of the Services during the Term, does not have the ability to correct or block Customer Data as required by applicable law, or to migrate Customer Data to another system or service provider, Google will comply with any reasonable requests by Customer to assist in facilitating such actions to the extent Google is legally permitted to do so and has reasonable access to the Customer Data.

8.2. **End User Requests**. During the Term, if Google receives any request from an End User for records relating to that End User’s personal data included in the Customer Data, Google will advise such End User to submit its request to Customer. Customer will be responsible for responding to any such request using the functionality of the Services.

9. **Data Privacy Officer.**

The Data Privacy Officer can be contacted by Customer Administrators at: https://support.google.com/a/contact/gfw_dpo (or via such other means as may be provided by Google). Administrators must be signed in to their Admin Account to use this address.

10. **Data Transfers.**

10.1. **Data Storage and Processing Facilities**. Google may store and process Customer Data in the United States or any other country in which Google or any of its Subprocessors maintains facilities, subject to Section 10.2 (Transfers of Data Out of the EEA) below.

10.2. **Transfers of Data Out of the EEA**. If the storage and processing of Customer Data (as set out in Section 10.1 above) involves transfers of Customer personal data out of the EEA and Data Protection Legislation applies to those transfers, Google will:

10.2.1 ensure that Google Inc. maintains its Safe Harbor Certification, and that the transfers are made in accordance with such Safe Harbor Certification; and/or

10.2.2 ensure that Google Inc. as the data importer of such Customer personal data enters into Model Contract Clauses with Customer (or an authorized Customer Affiliate) as the data exporter of such data, if Customer so requests, and that the transfers are made in accordance with any such Model Contract Clauses; and/or

10.2.3 adopt an alternative solution that achieves compliance with the terms of the Directive for transfers of personal data to a third country, and ensure that the transfers are made in accordance with any such compliance solution.

10.3. **Safe Harbor or Certification and Processing Practices.** While Google Inc. maintains its Safe Harbor Certification pursuant to Section 10.2.1, Google will ensure that: (a) the scope of such Safe Harbor Certification includes Customer Data; and (b) the Google Group's processing practices in respect of Customer Data remain consistent with those described in such Safe Harbor Certification.

10. **Data Center Information.** Google will make available to Customer information about the countries in which data centers used to store Customer Data are located.

11. **Subprocessors.**

11.1. **Subprocessors.** Google may engage Subprocessors to provide parts of the Services and related technical support services, subject to the restrictions in this Data Processing Amendment.

11.2. **Subprocessing Restrictions.** Google will ensure that Subprocessors only access and use Customer Data in accordance with the terms of the Agreement and that Subprocessors are bound by written agreements that require them to provide at least the level of data protection required by the following, as applicable pursuant to Section 10.2 (Transfers of Data Out of the EEA): (a) any Safe Harbor Certification maintained by Google Inc.; (b) any Model Contract Clauses entered into by Google Inc. and Customer (or an authorized Customer Affiliate); and/or (c) any alternative compliance solution adopted by Google.

11.3. **Consent to Subprocessing.** Customer consents to Google subcontracting the processing of Customer Data to Subprocessors in accordance with the Agreement. If the Model Contract Clauses have been entered into as described above, Customer (or, if applicable, an authorized Customer Affiliate) consents to Google Inc. subcontracting the processing of Customer Data in accordance with the terms of the Model Contract Clauses.

11. **Additional Information.** Information about Third Party Subprocessors is available at the following URL: www.google.com/intl/en/work/apps/terms/subprocessors.html, as such URL may be updated by Google from time to time. The information available at the URL is accurate at the time of publication. At the written request of the Customer, Google will provide additional information regarding Subprocessors and their locations. Any such requests must be sent to the Data Privacy Officer for Google Apps as described in Section 9 (Data Privacy Officer) of this Data Processing Amendment.

11. **Termination.** Google will, at least 14 days before appointing any new Third Party Subprocessor, inform Customer of the appointment (including the name and location of such subprocessor and the activities it will perform) either by sending an email to the Notification Email Address or via the Admin Console. If Customer objects to Google's use of any new Third Party Subprocessor, Customer may, as its sole and exclusive remedy, terminate the Google Apps Agreement by giving written notice to Google within 30 days of being informed by Google of the appointment of such subprocessor.

12. Liability Cap.

If Google Inc. and Customer (or an authorized Customer Affiliate) enter into Model Contract Clauses as described above, then, subject to the remaining terms of the Agreement relating to liability (including any specific exclusions from any limitation of liability), the total combined liability of Google and its Affiliates, on the one hand, and Customer and its Affiliates, on the other hand, under or in connection with the Agreement and all those MCCs combined will be limited to the maximum monetary or payment-based liability amount set out in the Agreement.

13. Third Party Beneficiary.

Notwithstanding anything to the contrary in the Agreement, where Google Inc. is not a party to the Agreement, Google Inc. will be a third party beneficiary of Section 10 (Auditing Security Compliance), Section 11.3 (Consent to Subprocessing) and Section 12 (Liability Cap) of this Data Processing Amendment.

14. Effect of Amendment.

To the extent of any conflict or inconsistency between the terms of this Data Processing Amendment and the remainder of the Agreement, the terms of this Data Processing Amendment will govern. Subject to the amendments in this Data Processing Amendment, the Agreement remains in full force and effect.

Appendix Categories of Data and Data Subjects**Categories of Data**

Personal data submitted, stored, sent or received by Customer or End Users via the Services may include the following categories of data: user IDs, email, documents, presentations, images, calendar entries, tasks and other electronic data

Data Subjects

Personal data submitted, stored, sent or received via the Services may concern the following categories of data subjects: End Users including Customer's employees and contractors; the personnel of Customer's customers, suppliers and subcontractors; and any other person who transmits data via the Services, including individuals collaborating and communicating with End Users.

Appendix Security Measures

As of the Amendment Effective Date, Google will take and implement the Security Measures set out in this Appendix to the Data Processing Amendment. Google may update or modify such Security Measures from time to time provided that such updates and modifications do not result in the degradation of the overall security of the Services.

1. Data Center Network Security.**(a) Data Centers.**

Infrastructure. Google maintains geographically distributed data centers. Google stores all production data in physically secure data centers.

Redundancy. Infrastructure systems have been designed to eliminate single points of failure and minimize the impact of anticipated environmental risks. Dual circuits, switches, networks or other necessary devices help provide this redundancy. The Services are designed to allow Google to perform certain types of preventative and corrective maintenance without interruption. All environmental equipment and facilities have documented preventative maintenance procedures that detail the process for and frequency of performance in accordance with the

manufacturer's or internal specifications. Preventative and corrective maintenance of the data center equipment is scheduled through a standard change process according to documented procedures.

Power. The data center electrical power systems are designed to be redundant and maintainable without impact to continuous operations, 24 hours a day, and 7 days a week. In most cases, a primary as well as an alternate power source, each with equal capacity, is provided for critical infrastructure components in the data center. Backup power is provided by various mechanisms such as uninterruptible power supplies (UPS) batteries, which supply consistently reliable power protection during utility brownouts, blackouts, over voltage, under voltage, and out-of-tolerance frequency conditions. If utility power is interrupted, backup power is designed to provide transitory power to the data center, at full capacity, for up to 10 minutes until the diesel generator systems take over. The diesel generators are capable of automatically starting up within seconds to provide enough emergency electrical power to run the data center at full capacity typically for a period of days.

Server Operating Systems. Google servers use a Linux based implementation customized for the application environment. Data is stored using proprietary algorithms to augment data security and redundancy. Google employs a code review process to increase the security of the code used to provide the Services and enhance the security products in production environments.

Businesses Continuity. Google replicates data over multiple systems to help to protect against accidental destruction or loss. Google has designed and regularly plans and tests its business continuity planning/disaster recovery programs.

(b) **Networks Transmission.**

Data Transmission. Data centers are typically connected via high-speed private links to provide secure and fast data transfer between data centers. This is designed to prevent data from being read, copied, altered or removed without authorization during electronic transfer or transport or while being recorded onto data storage media. Google transfers data via Internet standard protocols.

External Attack Surface. Google employs multiple layers of network devices and intrusion detection to protect its external attack surface. Google considers potential attack vectors and incorporates appropriate purpose built technologies into external facing systems.

Intrusion Detection. Intrusion detection is intended to provide insight into ongoing attack activities and provide adequate information to respond to incidents. Google's intrusion detection involves:

1. Tightly controlling the size and make-up of Google's attack surface through preventative measures;
2. Employing intelligent detection controls at data entry points; and
3. Employing technologies that automatically remedy certain dangerous situations.

Incident Response. Google monitors a variety of communication channels for security incidents, and Google's security personnel will react promptly to known incidents.

Encryption Technologies. Google makes HTTPS encryption (also referred to as SSL or TLS connection) available.

2. Access and Site Controls.

(a) Site Controls.

On-site Data Center Security Operation. Google's data centers maintain an on-site security operation responsible for all physical data center security functions 24 hours a day, 7 days a week. The on-site security operation personnel monitor Closed Circuit Television (CCTV) cameras and all alarm systems. On-site Security operation personnel perform internal and external patrols of the data center regularly.

Data Center Access Procedures. Google maintains formal access procedures for allowing physical access to the data centers. The data centers are housed in facilities that require electronic card key access, with alarms that are linked to the on-site security operation. All entrants to the data center are required to identify themselves as well as show proof of identity to on-site security operations. Only authorized employees, contractors and visitors are allowed entry to the data centers. Only authorized employees and contractors are permitted to request electronic card key access to these facilities. Data center electronic card key access requests must be made through e-mail, and require the approval of the requestor's manager and the data center director. All other entrants requiring temporary data center access must: (i) obtain approval in advance from the data center managers for the specific data center and internal areas they wish to visit; (ii) sign in at on-site security operations; and (iii) reference an approved data center access record identifying the individual as approved.

On-site Data Center Security Devices. Google's data centers employ an electronic card key and biometric access control system that is linked to a system alarm. The access control system monitors and records each individual's electronic card key and when they access perimeter doors, shipping and receiving, and other critical areas. Unauthorized activity and failed access attempts are logged by the access control system and investigated, as appropriate. Authorized access throughout the business operations and data centers is restricted based on roles and the individual's job responsibilities. The fire doors at the data centers are alarmed. CCTV cameras are in operation both inside and outside the data centers. The positioning of the cameras has been designed to cover strategic areas including, among others, the perimeter, doors to the data center building, and shipping/receiving. On-site security operations personnel manage the CCTV monitoring, recording and control equipment. Secure cables throughout the data centers connect the CCTV equipment. Cameras record on site via digital video recorders 24 hours a day, 7 days a week. The surveillance records are retained for up to 30 days based on activity.

(b) Access Control.

Infrastructure Security Personnel. Google has, and maintains, a security policy for its personnel, and requires security training as part of the training package for its personnel. Google's infrastructure security personnel are responsible for the ongoing monitoring of Google's security infrastructure, the review of the Services, and responding to security incidents.

Access Control and Privilege Management. Customer's Administrators and End Users must authenticate themselves via a central authentication system or via a single sign on system in order to use the Services. Each application checks credentials in order to allow the display of data to an authorized End User or authorized Administrator.

Internal Data Access Processes and Policies – Access Policy. Google's internal data access processes and policies are designed to prevent unauthorized persons and/or systems from gaining access to systems used to process personal data. Google aims to design its systems to: (i) only allow authorized persons to access data they are authorized to access; and (ii) ensure that personal data cannot be read, copied, altered or removed without authorization during processing, use and after recording. The systems are designed to detect any inappropriate access. Google employs a centralized access management system to control personnel access to production servers, and only provides access to a limited number of authorized personnel. LDAP, Kerberos and a proprietary system utilizing RSA keys are designed to provide Google with secure and flexible access mechanisms. These mechanisms are designed to grant only approved access rights to site hosts, logs, data and configuration information. Google requires the use of unique user IDs, strong passwords, two factor authentication and carefully monitored access lists to minimize the potential for unauthorized account use. The granting or modification of access rights is based on: the authorized personnel's job responsibilities; job duty requirements necessary to perform authorized tasks; and a need to know basis. The granting or modification of access rights must also be in accordance with Google's internal data access policies and training. Approvals are managed by workflow tools that maintain audit records of all changes. Access to systems is logged to create an audit trail for accountability. Where passwords are employed for authentication (e.g., login to workstations), password policies that follow at least industry standard practices are implemented. These standards include password expiry, restrictions on password reuse and sufficient password strength. For access to extremely sensitive information (e.g., credit card data), Google uses hardware tokens.

3. **Data.**

(a) **Data Storage, Isolation & Authentication.**

Google stores data in a multi-tenant environment on Google-owned servers. Data, the Services database and file system architecture are replicated between multiple geographically dispersed data centers. Google logically isolates data on a per End User basis at the application layer. Google logically isolates each Customer's data, and logically separates each End User's data from the data of other End Users, and data for an authenticated End User will not be displayed to another End User (unless the former End User or an Administrator allows the data to be shared). A central authentication system is used across all Services to increase uniform security of data.

The Customer will be given control over specific data sharing policies. Those policies, in accordance with the functionality of the Services, will enable Customer to determine the product sharing settings applicable to End Users for specific purposes. Customer may choose to make use of certain logging capability that Google may make available via the Services, products and APIs. Customer agrees that its use of the APIs is subject to the API Terms of Use. Google agrees that changes to the APIs will not result in the degradation of the overall security of the Services.

(b) **Decommissioned Disks and Disk Erase Policy.**

Certain disks containing data may experience performance issues, errors or hardware failure that lead them to be decommissioned ("Decommissioned Disk"). Every Decommissioned Disk is subject to a series of data destruction processes (the "Disk Erase Policy") before leaving Google's premises either for reuse or destruction. Decommissioned Disks are erased in a multi-step process and verified complete by at least two independent validators. The erase results are logged by the

Decommissioned Disk's serial number for tracking. Finally, the erased Decommissioned Disk is released to inventory for reuse and redeployment. If, due to hardware failure, the Decommissioned Disk cannot be erased, it is securely stored until it can be destroyed. Each facility is audited regularly to monitor compliance with the Disk Erase Policy

Personnel Security.

Google personnel are required to conduct themselves in a manner consistent with the company's guidelines regarding confidentiality, business ethics, appropriate usage, and professional standards. Google conducts reasonably appropriate backgrounds checks to the extent legally permissible and in accordance with applicable local labor law and statutory regulations.

Personnel are required to execute a confidentiality agreement and must acknowledge receipt of, and compliance with, Google's confidentiality and privacy policies. Personnel are provided with security training. Personnel handling Customer Data are required to complete additional requirements appropriate to their role (eg., certifications). Google's personnel will not process Customer Data without authorization.

Subprocessor Security.

Prior to onboarding Subprocessors, Google conducts an audit of the security and privacy practices of Subprocessors to ensure Subprocessors provide a level of security and privacy appropriate to their access to data and the scope of the services they are engaged to provide. Once Google has assessed the risks presented by the Subprocessor, then subject always to the requirements set out in Section 4.2 (Subprocessing Restrictions) of this Data Processing Amendment, the Subprocessor is required to enter into appropriate security, confidentiality and privacy contract terms.

Google Apps Data Processing Amendment, Version 2.0

Getting Started with Microsoft Azure Virtual Machines

Introduction

You can use a Microsoft Azure Virtual Machine when you need a scalable, cloud-based server running a Windows or Linux operating system and any application of your choosing. By taking advantage of Microsoft Azure Infrastructure as a Service (IaaS), you can run a virtual machine on an ongoing basis, or you can stop and restart it later with no loss to your data or server settings.

You can quickly provision a new virtual machine from one of the images available from Microsoft Azure. On the other hand, if you are already running applications in a VMware or Hyper-V virtualized environment, you can easily migrate your virtual machine to Microsoft Azure. Once you have created your virtual machine and added it to Microsoft Azure, you can work with it much like an on-premises server by attaching more disks for data storage or by installing and running applications on it.

Virtual machines rely on Microsoft Azure Storage for high availability. When your virtual machine is provisioned, it is replicated to three separate locations within the data center to which you assign it. You have the option to enable geo-replication to have copies of your virtual machine available in a remote data center region.

Considering Scenarios for a Virtual Machine

Microsoft Azure gives you not only the flexibility to support many application platforms, but also the flexibility to scale up and scale down to suit your requirements. Furthermore, you can quickly provision a new virtual machine in a few minutes. A Microsoft Azure virtual machine is simply a fresh machine preloaded with an operating system of your choice—you can add any needed application easily. Consider the following suggested scenarios as merely a few of the possibilities.

Existing virtual environment. You can move an existing server workload into the cloud simply by copying an existing on-premises VHD to Microsoft Azure. That way, you can leverage cloud benefits such as elasticity, redundancy, and automatic platform updates, while paying only for storage and the time the virtual machine is in use.

Cloud workstation. You can create a virtual machine for use in the cloud that contains applications you commonly use and want to access from any location. Rather than set up a desktop workstation or carry a laptop to run the applications you use frequently, you can install your applications on a virtual machine that is accessible in the cloud. With this type of environment, you can travel easily from location to location with a tablet and then connect to a virtual machine using RDP or SSH to access your routinely used applications such as Matlab, R, IPython, or SQL Server.

Linux cluster. You can create a virtual machine based on a Linux operating system and then change the configuration to run multiple virtual machines as a cluster. By using a cluster, you provide a fast, high capacity, and redundant environment for any custom Big Data and Big Compute applications you choose to install on the virtual machine.

Development or test environment. Another option is to set up virtual machines for development or test purposes that you use solely on an as-needed basis. For example, you could install Visual Studio on a

virtual machine for use during the development phase of a project without the need to allocate on-premises resources to development. Likewise, rather than incur the cost of setting up and maintaining a load test environment that you use only periodically, you can set up virtual machines to run only when you need an independent and isolated environment.

Custom machine images. Perhaps you need a team of people to work with a common set of complex tools and packages. You can take advantage of community-contributed virtual machine images available at VM Depot. Another option is to create and contribute a virtual machine that you preconfigure with selected applications or packages for an online community to reuse in the cloud. See “Using and Contributing Virtual Machines to VM Depot” at <http://research.microsoft.com/en-us/projects/azure/technical-papers.aspx> to learn more.



Creating a Virtual Machine

There are two primary ways to create a virtual machine that runs in Microsoft Azure—mounting an existing disk or using an image.



- A disk is a virtual hard disk (VHD) that already contains an operating system at minimum and optionally includes one or more applications. Any configuration changes you make are persisted on the disk and available the next time you start the VHD.



- An image is a template that contains a base operating system (Windows or Linux) and perhaps one or more applications, but no user accounts. A disk for the operating system is added to Microsoft Azure Storage and assigned to the new virtual machine during the provisioning process. You can use an image to quickly reproduce the same virtual machine multiple times.

Using the Virtual Machine Image Gallery

A simple way to get started using a virtual machine is to create virtual machine from the Image Gallery that you can access from the Microsoft Azure Management Portal. To view the available images, click the New button (which displays as a plus symbol on the bottom left side of the portal), select Virtual Machine in the Compute group, and then click From Gallery, as shown in Figure 1.




Figure 1: Creating a New Virtual Machine

The Create a Virtual Machine dialog box gives you the following three options:

- **Platform Images.** Here you can find a collection of images, as shown in Figure 2, containing only an operating system, such as Windows Server 2012 R2 or OpenLogic CentOS 6.3, or images that contain both an operating system and an application, such as SharePoint Server 2013 or Visual Studio Ultimate 2013. For a list of supported operating system roles and features as well as supported software, see <http://support.microsoft.com/kb/2721672>.
- **My Images.** This section of the Image Gallery displays a list of images that you create and store in Microsoft Azure. To create an image, you can capture an existing Microsoft Azure Virtual Machine, use SUSE Studio to build and migrate an image, or upload your own virtual machine, as explained later in this document.
- **My Disks.** You can create a virtual machine simply by selecting a disk on this list.

Note: In addition to the images available in the Virtual Machine Image Gallery, there are images created and contributed by members of the online community available at [VM Depot](http://vmdepot.net). Refer to the document "Using and Contributing Virtual Machines to VM Depot" at <http://research.microsoft.com/en-us/projects/azure/technical-papers.aspx> for more information.




Figure 2: Selecting a Virtual Machine Image from the Gallery

After you select the image you want, you assign a name to the virtual machine, specify a size, and create the administrator's login name and password for a Windows operating system or specify the authentication method for a Linux operating system. Some platform images have minimum size recommendations noted in the image description. The number of CPU cores and the amount of available memory defines the size of a virtual machine. At one end of the spectrum, a virtual machine can share a CPU core, consume only 768 MB of memory, and support one data disk. The maximum size of a data disk is 1 TB. At the other end of the spectrum, a virtual machine can be allocated eight CPU cores, up to 56 GB of memory, and up to 16 data disks at the time of this writing, but these maximum values are likely to increase as Microsoft Azure adds more powerful hardware offerings in the future. Pricing of the virtual machine depends on the size of the virtual machine you select.

Your next step is to configure the following additional settings for your virtual machine:

- **Cloud Service.** A cloud service is simply a container for your virtual machines. You can host a single virtual machine in a cloud service or assign multiple virtual machines to the same cloud service. When configuring the cloud service, you can request a new cloud service or select one that you created earlier.
- **Cloud Service DNS Name.** You must provide a unique name that you later use to connect to the virtual machine.
- **Subscription.** You must assign the virtual machine to an existing subscription.
- **Region/Affinity Group/Virtual Network.** You must select a data center region to host your virtual machine. Alternatively, if you have created an affinity group or virtual network, you can select either of those options instead.
- **Storage Account.** The disk for your virtual machine that gets created from the image you select must be assigned to a storage account. You can create a new one for your virtual machine or select a storage account that you created earlier.
- **Availability Set.** An optional step is to create an availability set or assign the virtual machine to an existing availability set to ensure high availability. That way, if your physical server is taken

down for maintenance, every service included in the availability set will be accessible through other physical hosts. When multiple virtual machines belong to the same availability set, the Microsoft Azure Fabric Controller places virtual machines across multiple racks, each of which can be considered a fault domain and separate physical points of failure, as shown in Figure 3. In addition, each virtual machine is placed in a separate upgrade domain so that virtual machines in the same availability set are not taken offline at the same time during host updates.




Figure 3: Virtual Machines Assigned to Separate Availability Sets

Last, you configure endpoints for your virtual machine. The image you select has one or more preconfigured endpoints. For example, if the image runs a Windows operating system, the Remote Desktop and PowerShell endpoints are created with public and private ports to use when you or your applications connect to the virtual machine. You can add more endpoints if necessary, such as FTP, SSH, or SMTP for example, and specify whether to assign the TCP or UDP protocol to the endpoint. Additionally, you can change the public and private port numbers from the defaults to alternate values. You can also take advantage of advanced networking features in Microsoft Azure Virtual Network to connect your virtual machine to an on-premises datacenter or a single client machine using an IPsec connection.

When you complete the configurations steps, the virtual machine is provisioned in Microsoft Azure. You can monitor this process, which typically takes only a few minutes, in the Microsoft Azure Management Portal. When provisioning is complete, the virtual machine starts. When its status changes to Running, your new virtual machine is available for you to connect to it as described in the “Using a Virtual Machine” section of this document.

If you are not ready to start working with your virtual machine, select it in the list on the Virtual Machines page in Microsoft Azure Management Portal, and then click Shut Down. That way, you can save usage costs, although there is an ongoing cost associated with the disk in your Microsoft Azure Storage account.

Note: Specific steps for creating a Windows virtual machine from the Image Gallery are available at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-windows-tutorial/>. If you

want to create a Linux virtual machine instead, refer to the instructions at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-tutorial/>.

Capturing a Microsoft Azure Virtual Machine

If you have an existing Windows-server based Microsoft Azure Virtual Machine for which you want to create an image, you can use the Sysprep utility to put the virtual machine into the necessary state. To do this, open a command prompt window, navigate to the `system32\sysprep` directory in the Windows directory of the virtual machine, and execute `sysprep.exe`. In the Sysprep tool, select the Enter System Out-of-Box Experience (OOBE) action, select the Generalize check box, and select Quit as the Shutdown Option. When the `sysprep succeeded.tag` appears in the same folder as the Sysprep tool, log off the virtual machine and then use the Shut Down button on the command bar to stop the virtual machine. When it has successfully stopped, click the Capture button to generate the image. You then name the image and confirm that you have run the Sysprep utility. When you are ready to use the image, you can find it on the My Images section of the Image Gallery described in the “Creating a Virtual Machine” section of this document.

For more detailed instructions, see <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-capture-image-windows-server/>.

Using SUSE Studio to Create a Virtual Machine Image

[SUSE Studio](#) is a web-based service that you can use to create a virtual machine (called an appliance by SUSE Studio) that you can then transfer to run in Microsoft Azure. SUSE Studio offers a wizard interface to simplify the process of creating your virtual machine. You start by selecting an open source operating system, such as SUSE Linux Enterprise 11 SP3, a 32-bit or 64-bit architecture, and a name for your virtual machine. Then you select packages to install, such as Perl or Python. Next you configure the locale, time zone, firewall, and users for the virtual machine in addition to allocating memory and other configuration tasks. You can also specify files to upload into the virtual machine.

After completing the configuration, you use a Build command to produce the virtual machine in one or more specified formats, including a Microsoft Azure-ready image. An Upload link displays for you to move the prepared image into your Microsoft Azure account after the build process completes. Instructions are available to help you download a settings file that you then upload to your SUSE account. Then you specify the Microsoft Azure storage account into which you want the image transferred, and then click the link to go to the SUSE Studio’s dashboard where you click the Upload button to send your image to Microsoft Azure. After the upload completes, you follow the instructions for creating a virtual machine from the Virtual Machine Image Gallery. Your uploaded image displays on the My Images tab of the Images Gallery wizard.

For more information about limitations of Microsoft Azure images created by using SUSE Studio, see <http://susestudio.com/help/use/windows-azure.html>.

Uploading a Virtual Machine to Microsoft Azure

Whether you want to permanently run one of your on-premises virtual machines in Microsoft Azure or simply want to test a configuration outside of your production environment, you can upload a copy of that virtual machine as a page blob to your Microsoft Azure Storage account, as shown in Figure 4. You can even move the virtual machine back on-premises later if necessary. The VHD can be a sysprepped fixed disk VHD or a customized Linux image based on supported distributions.




Figure 4: Mobility of a VHD between On-Premises and Microsoft Azure

Uploading an Existing Windows Virtual Machine

If you are already using a Hyper-V virtual machine, you can upload the associated VHD file to Microsoft Azure and begin using the virtual machine from the cloud. If you are using VMware virtualization, you can use the System Center Virtual Machine Manager to [convert a VMDK file to a VHD file](#) or use [Microsoft Virtual Machine Converter](#). Much like you do with a Microsoft Azure virtual machine, you generalize the operating system of your on-premises virtual machine by using the Sysprep utility. You launch the utility from the system32\sysprep directory in the Windows directory of the virtual machine. Then you select the Enter System Out-of-Box Experience (OOBE) action, select the Generalize check box, and select Shutdown as the Shutdown Option.

Next you must create a new storage account in the Microsoft Azure Management Portal if you do not have one already available. Click New, click Storage, and then click Quick Create. You need to supply a subdomain name as well as a location or affinity group, and choose whether to use geo-replication. You then use Microsoft Azure PowerShell cmdlets, `Get-AzurePublishSettingsFile`, `Import-AzurePublishSettingsFile`, and `Add-AzureVhd`, to upload the VHD file for your virtual machine to your storage account.

After the VHD is uploaded, you can go to the Virtual Machines page in the Microsoft Azure Management Portal, click the Images link above the list of virtual machines, and click Create to set up a new image based on your VHD. You must provide a name for the image, an optional description, the VHD location (which you can obtain by browsing your cloud storage accounts), and the operating system family (Windows, in this case).

For a complete list of steps to perform, refer to <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-create-upload-vhd-windows-server/>.

Uploading a Local Linux Virtual Machine

There are several steps required to prepare and upload a local Linux virtual machine and different steps required for different Linux distributions. At the conclusion of the preparation process, regardless of Linux version, you must install the [Microsoft Azure Linux Agent](#). In addition, you must provision a Microsoft Azure storage account. In the Microsoft Azure Management Portal, click New, click Storage, and then click Quick Create. Next, type in a subdomain name, select a location or affinity group, and specify whether to

use geo-replication. Then just as you do for a Windows VHD, you use Microsoft Azure PowerShell cmdlets, `Get-AzurePublishSettingsFile`, `Import-AzurePublishSettingsFile`, and `Add-AzureVhd`, to upload the VHD file for your virtual machine to your storage account.

You can review the details for uploading a Linux virtual machine at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-linux-create-upload-vhd/>.

Working with a Virtual Machine

The disk for the virtual machine's operating system is limited to 127 GB. The virtual machine also has a disk allocated to temporary storage for data, but this data does not persist if the virtual machine is relocated to other hardware for any reason. Therefore, do not store data in this location if you cannot easily recreate it. These disks are assigned to the drive letters C and D, respectively. If necessary, you can add a data disk to increase your storage capacity and to hold additional applications that you want to install. You can assign any drive letter, other than C or D, to each data disk that you add.

Connecting to a Windows Virtual Machine

If the virtual machine is not currently running, you must first select it in the virtual machines list in the Microsoft Azure Management Portal and then click **Start** in the command bar at the bottom of the page. After the virtual machine is running, click the **Connect** button to generate the remote desktop connection file, click **Open** to use the file, and then click **Connect** to establish the connection to the virtual machine. Type in the administrator login name and password to authenticate, and then click **Yes** to verify the virtual machine's identity. At this point, you can interact with the virtual machine in the same way you would use a local server.

Note: If your firewall has a rule that does not allow you to connect to the RDP port of your virtual machine, you may change the public port of your virtual machine's endpoint.

Connecting to a Linux Virtual Machine

Start the virtual machine if it is not currently running. Then use an SSH client, such as [PuTTY](#) for a Windows client or [OpenSSH](#) for a Linux client, to connect to your virtual machine. To make this connection, you must know the host name and port for the virtual machine. You can get this information by clicking the virtual machine's name in the Microsoft Azure Management Portal, clicking the **Dashboard** link, and then locating the **SSH Details** section in the lower right corner of the page. After connecting to the virtual machine, you can use it just like any other local server.

Attaching a Data Disk

If you need additional storage for your virtual machine, you can attach an empty disk or a VHD that you have preloaded with data. Select your virtual machine in the list of instances, click **Attach** in the command bar, and then click **Attach Empty Disk** or **Attach Disk** as appropriate. For an existing disk, you select the disk in storage and specify a **Host Cache Preference** option. For a new empty disk, a storage location and a file name are automatically assigned, but you can override the file name if you prefer. You must specify a size for the disk in gigabytes, and then select from one of the following **Host Cache Preference** options:

- **None.** Host caching is disabled. This is the default setting for a data disk.
- **Read Only.** Host caching is enabled only for read operations as a performance optimization.
- **Read Write.** Host caching is enabled for both reads and writes. This is the default setting for an operating system disk.

When the disk is attached, you must next initialize the disk by using the Disk Management utility if you are using a Windows virtual machine, and then create a new simple volume. Once the volume is created, the disk is available for use by your applications. If you are using a Linux virtual machine, you use the `fdisk` command to create a new device and a new partition, create a file system on the new partition, add a directory, and then mount the drive.

You can review the steps to attach a data disk for either operating system at <http://azure.microsoft.com/en-us/documentation/articles/storage-windows-attach-disk/>.

Managing a Virtual Machine

You have three options for managing virtual machines—the management portal, scripting, and the Rest API.



- You can use the graphical interface of the Microsoft Azure Management Portal to perform various tasks related to managing your virtual machine. You can start and stop a virtual machine, connect to it, attach and detach disks, capture the VHD as an image, or delete the VHD. In addition, you can configure a different size for the virtual machine, assign it to an availability set, or add new endpoints.



- If you prefer, you can also download and install Microsoft Azure PowerShell cmdlets to use scripts for management tasks. Instructions are available at <http://azure.microsoft.com/en-us/documentation/articles/install-configure-powershell/>. As an example, you can use the [Add-AzureDataDisk](#) cmdlet to add a new data disk to your virtual machine.
- Another similar option is the CLI available for all platforms. You can learn more about the CLI for Windows, Linux, or Mac at <http://azure.microsoft.com/en-us/documentation/articles/xplat-cli/>.



- You can use various language-specific SDKs to manage Microsoft Azure. See the language-specific Developer Centers documentation at <http://azure.microsoft.com/en-us/documentation/>.
- Although the portal, PowerShell cmdlets, CLI, and language-specific SDKs all use common REST APIs behind the scenes, you can also access [Blob Service REST APIs](#) directly if you prefer. With this API, you can manage the containers and blobs in your Storage account, including uploading, downloading, or deleting your VHDs.

AWS PUBLIC SECTOR ACCESS POLICY (STATE)

(Last Updated January 4, 2013)

This AWS Access Policy ("**Access Policy**") governs your access to and use of the Services (as defined below) of Amazon Web Services, Inc. ("**AWS**") provided to you by your systems integrator, reseller, or services provider ("**Provider**"). It sets out the additional rules, conditions and restrictions that apply to you or the entity you represent ("**you**") for use of the Services. In this Access Policy, "**we**", "**us**", or "**our**" means AWS and any of its affiliates. Please see Section 10 for definitions of capitalized terms.

1. Use of the Services.

1.1 Generally. You are provided access to the Services by your Provider. Your use of and access to the Services are governed by the agreement between you and Provider. This Access Policy supplements the terms of such agreement and may be updated by us from time to time. AWS Service Level Agreements do not apply to your use of the Services. Your continued access to and use of the Services is conditioned on your compliance with all laws, rules, regulations, policies and instructions applicable to your use of the Services, including the Policies.

1.2 Account Keys. Provider may provide you with AWS account keys which will allow you to directly access the Services via Provider's account(s). We are not responsible for any activities that occur under these account keys, regardless of whether the activities are undertaken by you, Provider or a third party (including your employees, contractors or agents) and we are also not responsible for unauthorized access to the account.

1.3 Third Party Materials. Through the use of Provider's AWS account(s), you may have access to Third Party Materials, such as software applications provided by third parties, which are made available directly to you by other companies or individuals under separate terms and conditions, including separate fees and charges. Your use of any Third Party Materials is at your sole risk.

2. Your Responsibilities.

2.1 Your Materials. You are solely responsible for the development, content, operation, maintenance, and use of Your Materials with the Services. For example, you are solely responsible for:

- (a) the technical operation of Your Materials, including ensuring that calls you make to any Service are compatible with then-current application program interfaces for that Service;
- (b) compliance of Your Materials with the Acceptable Use Policy, the other Policies, and the law;
- (c) any claims relating to Your Materials;
- (d) properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Materials violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act;
- (e) any action that you permit, assist or facilitate any person or entity to take related to this Access Policy, Your Materials or use of the Services; and
- (f) End Users' use of Your Materials and the Services and ensuring that End Users comply with your obligations under this Access Policy and that the terms of your agreement with each End User are consistent with this Access Policy.

2.2 Other Security and Backup. You or Provider are solely responsible for properly configuring and using the Services and taking steps to maintain appropriate security, protection and backup of Your Materials, including using encryption technology to protect Your Materials from unauthorized access and routinely archiving Your Materials.

2.3 End User Violations. If you become aware of any violation of your obligations under this Access Policy by an End User, you will immediately terminate such End User's access to Your Materials and the Services.

3. Service Interruption.

3.1 General. We may suspend the AWS account(s) through which you access the Services immediately if we determine your or an End User's use of the Services (i) violates the terms of this Access Policy (including the Acceptable Use Policy or Service Terms); (ii) poses a security risk to the Services or any other AWS customer, (iii) may harm our systems or the systems or Materials of any other AWS customer; or (iv) may subject us to liability as a result of any of the foregoing. We will provide notice of any suspension as soon as practicable to Provider, who is solely responsible for providing any notices to you under your agreement with them.

3.2 Scope of Interruption. To the extent practicable, we will (i) suspend your right to access or use only those instances, data, or portions of the Services that caused the suspension, and (ii) limit the suspension to those accounts that caused the suspension. If commercially feasible, access to the Services will be restored once the conditions or circumstances giving rise to the suspension have been removed or corrected. Nothing in this Section 3 will operate to limit your rights or remedies otherwise available to you against Provider under your agreement with them or applicable law.

4. Proprietary Rights.

4.1 Services. As between you and us, we or our licensors own and reserve all right, title, and interest in and to the Services. You have the right to use the Services solely as a licensee of Provider in accordance with this Access Policy and the agreement between you and Provider. We have no obligation to provide the Service to you under this Access Policy, so you must look exclusively to Provider and your agreement with Provider regarding such obligation. Except as expressly provided in this Section 4, you obtain no rights to the Services, the AWS Materials or any Third Party Materials.

4.2 Materials. As a part of the Services, you may have access to AWS Materials and Third Party Materials, which may be subject to additional terms and conditions (including the Terms of Use and Apache Software License). By using those materials, you are subject to such additional terms. You are solely responsible for securing any necessary approvals for the download and use of such materials.

4.3 Restrictions. Neither you nor any End User may use the Services in any manner or for any purpose other than as expressly permitted by this Access Policy and the agreement between you and Provider. Neither you nor any End User may, or may attempt to, (a) modify, alter, tamper with, repair, or otherwise create derivative works of any software included in the Services (except to the extent software included in the Services are provided to you under a separate license that expressly permits the creation of derivative works), (b) reverse engineer, disassemble, or decompile the software included in the Services or apply any other process or procedure to derive the source code of any software included in the Services, or (c) access or use the Services in a way intended to avoid incurring fees or exceeding usage limits or quotas. All rights and access granted to you with respect to the Services are conditioned on your continued compliance with this Access Policy, and you will immediately discontinue your use of the Services if you cannot comply with this Access Policy. You will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, vendors, business partners, or licensors, any patent infringement or other intellectual property infringement claim regarding any Services or AWS Materials that you have used.

4.4 Suggestions. If you provide any Suggestions to us when using the Services, you hereby grant to AWS and its affiliates a perpetual, irrevocable, non-exclusive, worldwide, royalty-free right and license to reproduce, distribute, make derivative works based upon, publicly display, publicly perform, make, have made, use, sell, offer for sale, and import the Suggestions, including the right to sublicense such rights through multiple tiers, alone or in combination.

4.5 Government Rights. If you are using the Services on behalf of the government and these terms fail to meet the government's needs or are inconsistent in any respect with federal or state law, you will immediately discontinue your use of the Services (including any AWS Materials).

5. Representations and Warranties. You represent and warrant that (a) you and your End Users' use of the Services (including any use by your employees and personnel) will not violate this Access Policy; (b) you or your

licensors own all right, title, and interest in and to Your Materials; (c) Your Materials (including the use, development, design, production, advertising, or marketing of your Materials) or the combination of your Materials with other applications, content or processes, do not and will not violate any applicable laws or infringe or misappropriate any third-party rights; and (d) your use of the Services will not cause harm to any End User.

6. Indemnification. Except to the extent prohibited by applicable state law, you will defend, indemnify, and hold harmless us, our licensors and each of our respective employees, officers, directors, and representatives from and against any claims, damages, losses, liabilities, costs, and expenses (including reasonable attorneys' fees) arising out of or relating to any third party claim concerning: (a) your or any End Users' use of the Services (including use by your employees and personnel); (b) Your Materials or the combination of Your Materials with other applications, content or processes, including any claim involving alleged infringement or misappropriation of third-party rights or the use, development, design, production, advertising or marketing of Your Materials; or (c) a dispute between you and any End User. If your ability to comply with the foregoing provision is limited to any extent by the absence of appropriations or government authorization, you will make good faith efforts to obtain sufficient appropriations or authorization for any liabilities arising under this Section 5.

7. Disclaimers. WE PROVIDE THE SERVICES ON AN "AS IS" BASIS TO PROVIDER. WE AND OUR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND TO YOU, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE REGARDING THE SERVICES OR ANY THIRD PARTY MATERIALS, INCLUDING ANY WARRANTY THAT THE SERVICES OR THIRD PARTY MATERIALS WILL BE UNINTERRUPTED, ERROR FREE OR FREE OF HARMFUL COMPONENTS, OR THAT ANY MATERIALS, INCLUDING YOUR MATERIALS OR THE THIRD PARTY MATERIALS, WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED. EXCEPT TO THE EXTENT PROHIBITED BY LAW, WE AND OUR LICENSORS DISCLAIM ALL WARRANTIES, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR USAGE OF TRADE.

8. Limitations of Liability. YOU MUST LOOK SOLELY TO PROVIDER AND YOUR AGREEMENT WITH THEM REGARDING ANY CLAIMS OR DAMAGES RELATED TO THE SERVICES. WE AND OUR AFFILIATES OR LICENSORS WILL NOT BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES (INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA), EVEN IF A PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. FURTHER, NEITHER WE NOR ANY OF OUR LICENSORS WILL BE RESPONSIBLE FOR ANY COMPENSATION, REIMBURSEMENT, OR DAMAGES ARISING IN CONNECTION WITH: (A) YOUR INABILITY TO USE THE SERVICES, INCLUDING AS A RESULT OF ANY (I) SUSPENSION OF YOUR USE OF OR ACCESS TO THE SERVICES, (II) OUR DISCONTINUATION OF ANY OR ALL OF THE SERVICES, OR, (III) ANY UNANTICIPATED OR UNSCHEDULED DOWNTIME OF ALL OR A PORTION OF THE SERVICES FOR ANY REASON; OR (B) ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION, DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR MATERIALS OR OTHER DATA THAT YOU OR ANY END USER SUBMITS OR USES IN CONNECTION WITH THE SERVICES (INCLUDING AS A RESULT OF YOUR OR ANY END USERS' ERRORS, ACTS OR OMISSIONS).

9. Miscellaneous.

9.1 Governing Law; Venue. Except to the extent prohibited by applicable state law, the laws of the State of Washington, without reference to conflict of law rules, govern this Access Policy and any dispute of any sort that might arise between you and us. You irrevocably consent to exclusive jurisdiction and venue of the federal courts located in King County, Washington with respect to any dispute arising in connection with the Services or this Access Policy. We may seek injunctive or other relief in any state, federal, or national court of competent jurisdiction for any actual or alleged infringement of our or any third party's intellectual property or other proprietary rights. The United Nations Convention for the International Sale of Goods does not apply.

9.2 Entire Agreement. This Access Policy is the entire agreement between you and us regarding the Services, including any materials. It supersedes all prior or contemporaneous representations, understandings, agreements, or communications between you and us, whether written or verbal, regarding the subject matter of this Access

Policy. If the terms of this document are inconsistent with the terms contained in your agreement with Provider, the terms contained in this document will control. We will not be bound by, and specifically object to, any term, condition or other provision which is different from or in addition to the provisions of this Access Policy (whether or not it would materially alter it) and which is submitted by you in any order, receipt, acceptance, confirmation, correspondence or other document.

9.3 Survival. The following provisions will survive any termination of your use of the Services: Sections 2.1, 4, 5, 6, 7, 8, 9 and 10.

10. Definitions.

“Acceptable Use Policy” means the policy currently available at <http://aws.amazon.com/aup>, as it may be updated by us from time to time.

“AWS Materials” means Materials we make available in connection with the Services or on the AWS Site to allow access to and use of the Services, including WSDLs; Documentation; sample code; software libraries; command line tools; and other related technology. AWS Materials does not include the Services.

“AWS Service Level Agreement” means all service level agreements that we offer with respect to the Services and post on the AWS Site, as they may be updated by us from time to time.

“AWS Site” means <http://aws.amazon.com> and any successor or related site designated by us.

“Documentation” means the developer guides, getting started guides, user guides, quick reference guides, and other technical and operations manuals, instructions and specifications for the Services currently located at <http://aws.amazon.com/documentation>, as such documentation may be updated by us from time to time.

“End User” means any individual or entity that directly or indirectly through another user: (a) accesses or uses Your Materials; or (b) otherwise accesses or uses the Services through you.

“Materials” means software (including machine images), data, text, audio, video, images or other content.

“Policies” means the Acceptable Use Policy, the Terms of Use, the Service Terms, all restrictions described in the AWS Materials and on the AWS Site, and any other policy or terms referenced in or incorporated into this Access Policy.

“Services” means, collectively or individually (as applicable), the web services made commercially available by us to Provider for use under this Access Policy, including (as applicable) those web services described in the Service Terms.

“Service Terms” means the rights and restrictions for particular Services located at <http://aws.amazon.com/serviceterms>, as they may be updated by us from time to time.

“Suggestions” means all suggested improvements to the Services or AWS Materials that you provide to us.

“Terms of Use” means the terms of use located at <http://aws.amazon.com/terms/>, as they may be updated by us from time to time.

“Third Party Materials” means Materials made available to you by any third party on the AWS Site or in conjunction with the Services.

“Your Materials” means Materials you or any End User (a) run on the Services, (b) cause to interface with the Services, or (c) upload to the Services or otherwise transfer, process, use or store in connection with the Services.

AWS PUBLIC SECTOR ACCESS POLICY (STATE)

(Last Updated January 4, 2013)

This AWS Access Policy ("**Access Policy**") governs your access to and use of the Services (as defined below) of Amazon Web Services, Inc. ("**AWS**") provided to you by your systems integrator, reseller, or services provider ("**Provider**"). It sets out the additional rules, conditions and restrictions that apply to you or the entity you represent ("**you**") for use of the Services. In this Access Policy, "**we**", "**us**", or "**our**" means AWS and any of its affiliates. Please see Section 10 for definitions of capitalized terms.

1. Use of the Services.

1.1 Generally. You are provided access to the Services by your Provider. Your use of and access to the Services are governed by the agreement between you and Provider. This Access Policy supplements the terms of such agreement and may be updated by us from time to time. AWS Service Level Agreements do not apply to your use of the Services. Your continued access to and use of the Services is conditioned on your compliance with all laws, rules, regulations, policies and instructions applicable to your use of the Services, including the Policies.

1.2 Account Keys. Provider may provide you with AWS account keys which will allow you to directly access the Services via Provider's account(s). We are not responsible for any activities that occur under these account keys, regardless of whether the activities are undertaken by you, Provider or a third party (including your employees, contractors or agents) and we are also not responsible for unauthorized access to the account.

1.3 Third Party Materials. Through the use of Provider's AWS account(s), you may have access to Third Party Materials, such as software applications provided by third parties, which are made available directly to you by other companies or individuals under separate terms and conditions, including separate fees and charges. Your use of any Third Party Materials is at your sole risk.

2. Your Responsibilities.

2.1 Your Materials. You are solely responsible for the development, content, operation, maintenance, and use of Your Materials with the Services. For example, you are solely responsible for:

- (a) the technical operation of Your Materials, including ensuring that calls you make to any Service are compatible with then-current application program interfaces for that Service;
- (b) compliance of Your Materials with the Acceptable Use Policy, the other Policies, and the law;
- (c) any claims relating to Your Materials;
- (d) properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Materials violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act;
- (e) any action that you permit, assist or facilitate any person or entity to take related to this Access Policy, Your Materials or use of the Services; and
- (f) End Users' use of Your Materials and the Services and ensuring that End Users comply with your obligations under this Access Policy and that the terms of your agreement with each End User are consistent with this Access Policy.

2.2 Other Security and Backup. You or Provider are solely responsible for properly configuring and using the Services and taking steps to maintain appropriate security, protection and backup of Your Materials, including using encryption technology to protect Your Materials from unauthorized access and routinely archiving Your Materials.

2.3 End User Violations. If you become aware of any violation of your obligations under this Access Policy by an End User, you will immediately terminate such End User's access to Your Materials and the Services.

3. Service Interruption.

3.1 General. We may suspend the AWS account(s) through which you access the Services immediately if we determine your or an End User's use of the Services (i) violates the terms of this Access Policy (including the Acceptable Use Policy or Service Terms); (ii) poses a security risk to the Services or any other AWS customer, (iii) may harm our systems or the systems or Materials of any other AWS customer; or (iv) may subject us to liability as a result of any of the foregoing. We will provide notice of any suspension as soon as practicable to Provider, who is solely responsible for providing any notices to you under your agreement with them.

3.2 Scope of Interruption. To the extent practicable, we will (i) suspend your right to access or use only those instances, data, or portions of the Services that caused the suspension, and (ii) limit the suspension to those accounts that caused the suspension. If commercially feasible, access to the Services will be restored once the conditions or circumstances giving rise to the suspension have been removed or corrected. Nothing in this Section 3 will operate to limit your rights or remedies otherwise available to you against Provider under your agreement with them or applicable law.

4. Proprietary Rights.

4.1 Services. As between you and us, we or our licensors own and reserve all right, title, and interest in and to the Services. You have the right to use the Services solely as a licensee of Provider in accordance with this Access Policy and the agreement between you and Provider. We have no obligation to provide the Service to you under this Access Policy, so you must look exclusively to Provider and your agreement with Provider regarding such obligation. Except as expressly provided in this Section 4, you obtain no rights to the Services, the AWS Materials or any Third Party Materials.

4.2 Materials. As a part of the Services, you may have access to AWS Materials and Third Party Materials, which may be subject to additional terms and conditions (including the Terms of Use and Apache Software License). By using those materials, you are subject to such additional terms. You are solely responsible for securing any necessary approvals for the download and use of such materials.

4.3 Restrictions. Neither you nor any End User may use the Services in any manner or for any purpose other than as expressly permitted by this Access Policy and the agreement between you and Provider. Neither you nor any End User may, or may attempt to, (a) modify, alter, tamper with, repair, or otherwise create derivative works of any software included in the Services (except to the extent software included in the Services are provided to you under a separate license that expressly permits the creation of derivative works), (b) reverse engineer, disassemble, or decompile the software included in the Services or apply any other process or procedure to derive the source code of any software included in the Services, or (c) access or use the Services in a way intended to avoid incurring fees or exceeding usage limits or quotas. All rights and access granted to you with respect to the Services are conditioned on your continued compliance with this Access Policy, and you will immediately discontinue your use of the Services if you cannot comply with this Access Policy. You will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, vendors, business partners, or licensors, any patent infringement or other intellectual property infringement claim regarding any Services or AWS Materials that you have used.

4.4 Suggestions. If you provide any Suggestions to us when using the Services, you hereby grant to AWS and its affiliates a perpetual, irrevocable, non-exclusive, worldwide, royalty-free right and license to reproduce, distribute, make derivative works based upon, publicly display, publicly perform, make, have made, use, sell, offer for sale, and import the Suggestions, including the right to sublicense such rights through multiple tiers, alone or in combination.

4.5 Government Rights. If you are using the Services on behalf of the government and these terms fail to meet the government's needs or are inconsistent in any respect with federal or state law, you will immediately discontinue your use of the Services (including any AWS Materials).

5. Representations and Warranties. You represent and warrant that (a) you and your End Users' use of the Services (including any use by your employees and personnel) will not violate this Access Policy; (b) you or your

licensors own all right, title, and interest in and to Your Materials; (c) Your Materials (including the use, development, design, production, advertising, or marketing of your Materials) or the combination of your Materials with other applications, content or processes, do not and will not violate any applicable laws or infringe or misappropriate any third-party rights; and (d) your use of the Services will not cause harm to any End User.

6. Indemnification. Except to the extent prohibited by applicable state law, you will defend, indemnify, and hold harmless us, our licensors and each of our respective employees, officers, directors, and representatives from and against any claims, damages, losses, liabilities, costs, and expenses (including reasonable attorneys' fees) arising out of or relating to any third party claim concerning: (a) your or any End Users' use of the Services (including use by your employees and personnel); (b) Your Materials or the combination of Your Materials with other applications, content or processes, including any claim involving alleged infringement or misappropriation of third-party rights or the use, development, design, production, advertising or marketing of Your Materials; or (c) a dispute between you and any End User. If your ability to comply with the foregoing provision is limited to any extent by the absence of appropriations or government authorization, you will make good faith efforts to obtain sufficient appropriations or authorization for any liabilities arising under this Section 5.

7. Disclaimers. WE PROVIDE THE SERVICES ON AN "AS IS" BASIS TO PROVIDER. WE AND OUR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND TO YOU, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE REGARDING THE SERVICES OR ANY THIRD PARTY MATERIALS, INCLUDING ANY WARRANTY THAT THE SERVICES OR THIRD PARTY MATERIALS WILL BE UNINTERRUPTED, ERROR FREE OR FREE OF HARMFUL COMPONENTS, OR THAT ANY MATERIALS, INCLUDING YOUR MATERIALS OR THE THIRD PARTY MATERIALS, WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED. EXCEPT TO THE EXTENT PROHIBITED BY LAW, WE AND OUR LICENSORS DISCLAIM ALL WARRANTIES, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR USAGE OF TRADE.

8. Limitations of Liability. YOU MUST LOOK SOLELY TO PROVIDER AND YOUR AGREEMENT WITH THEM REGARDING ANY CLAIMS OR DAMAGES RELATED TO THE SERVICES. WE AND OUR AFFILIATES OR LICENSORS WILL NOT BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES (INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA), EVEN IF A PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. FURTHER, NEITHER WE NOR ANY OF OUR LICENSORS WILL BE RESPONSIBLE FOR ANY COMPENSATION, REIMBURSEMENT, OR DAMAGES ARISING IN CONNECTION WITH: (A) YOUR INABILITY TO USE THE SERVICES, INCLUDING AS A RESULT OF ANY (I) SUSPENSION OF YOUR USE OF OR ACCESS TO THE SERVICES, (II) OUR DISCONTINUATION OF ANY OR ALL OF THE SERVICES, OR, (III) ANY UNANTICIPATED OR UNSCHEDULED DOWNTIME OF ALL OR A PORTION OF THE SERVICES FOR ANY REASON; OR (B) ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION, DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR MATERIALS OR OTHER DATA THAT YOU OR ANY END USER SUBMITS OR USES IN CONNECTION WITH THE SERVICES (INCLUDING AS A RESULT OF YOUR OR ANY END USERS' ERRORS, ACTS OR OMISSIONS).

9. Miscellaneous.

9.1 Governing Law; Venue. Except to the extent prohibited by applicable state law, the laws of the State of Washington, without reference to conflict of law rules, govern this Access Policy and any dispute of any sort that might arise between you and us. You irrevocably consent to exclusive jurisdiction and venue of the federal courts located in King County, Washington with respect to any dispute arising in connection with the Services or this Access Policy. We may seek injunctive or other relief in any state, federal, or national court of competent jurisdiction for any actual or alleged infringement of our or any third party's intellectual property or other proprietary rights. The United Nations Convention for the International Sale of Goods does not apply.

9.2 Entire Agreement. This Access Policy is the entire agreement between you and us regarding the Services, including any materials. It supersedes all prior or contemporaneous representations, understandings, agreements, or communications between you and us, whether written or verbal, regarding the subject matter of this Access

Policy. If the terms of this document are inconsistent with the terms contained in your agreement with Provider, the terms contained in this document will control. We will not be bound by, and specifically object to, any term, condition or other provision which is different from or in addition to the provisions of this Access Policy (whether or not it would materially alter it) and which is submitted by you in any order, receipt, acceptance, confirmation, correspondence or other document.

9.3 Survival. The following provisions will survive any termination of your use of the Services: Sections 2.1, 4, 5, 6, 7, 8, 9 and 10.

10. Definitions.

“Acceptable Use Policy” means the policy currently available at <http://aws.amazon.com/aup>, as it may be updated by us from time to time.

“AWS Materials” means Materials we make available in connection with the Services or on the AWS Site to allow access to and use of the Services, including WSDLs; Documentation; sample code; software libraries; command line tools; and other related technology. AWS Materials does not include the Services.

“AWS Service Level Agreement” means all service level agreements that we offer with respect to the Services and post on the AWS Site, as they may be updated by us from time to time.

“AWS Site” means <http://aws.amazon.com> and any successor or related site designated by us.

“Documentation” means the developer guides, getting started guides, user guides, quick reference guides, and other technical and operations manuals, instructions and specifications for the Services currently located at <http://aws.amazon.com/documentation>, as such documentation may be updated by us from time to time.

“End User” means any individual or entity that directly or indirectly through another user: (a) accesses or uses Your Materials; or (b) otherwise accesses or uses the Services through you.

“Materials” means software (including machine images), data, text, audio, video, images or other content.

“Policies” means the Acceptable Use Policy, the Terms of Use, the Service Terms, all restrictions described in the AWS Materials and on the AWS Site, and any other policy or terms referenced in or incorporated into this Access Policy.

“Services” means, collectively or individually (as applicable), the web services made commercially available by us to Provider for use under this Access Policy, including (as applicable) those web services described in the Service Terms.

“Service Terms” means the rights and restrictions for particular Services located at <http://aws.amazon.com/serviceterms>, as they may be updated by us from time to time.

“Suggestions” means all suggested improvements to the Services or AWS Materials that you provide to us.


“Terms of Use” means the terms of use located at <http://aws.amazon.com/terms/>, as they may be updated by us from time to time.

“Third Party Materials” means Materials made available to you by any third party on the AWS Site or in conjunction with the Services.

“Your Materials” means Materials you or any End User (a) run on the Services, (b) cause to interface with the Services, or (c) upload to the Services or otherwise transfer, process, use or store in connection with the Services.

Trusted Cloud: Microsoft Azure Security, Privacy, and Compliance


April 2015






Contents

Introduction	4
What customers want from cloud providers.....	5
Microsoft Azure: Built for trust	6
Security: Working to keep customer data safe.....	7
Security design and operations	7
Infrastructure protection.....	9
Network protection.....	10
Data protection	11
Identity and access.....	12
Privacy: Customers own and control their data	12
Customers are in control of their data.....	14
Transparency	15
Compliance: Azure conforms to global standards.....	16
Additional resources	18



Introduction

With the emergence of cloud computing, today's IT organizations are playing an increasingly important role in driving business strategy. While cost reduction is still a top priority, scalability and business agility have stepped to the forefront for IT decision makers. As a result, spending on cloud solutions is expected to grow 30 percent from 2013 to 2018, compared with 5 percent overall growth for enterprise IT. And cloud services are keeping pace¹. Analysts expect to see a ten-fold increase in the number of cloud-based solutions on the market in the next four to five years.



Sources:

- 1 Forbes, "Roundup of Cloud Computing Forecasts and Market Estimates, 2015," 1/24/2015. <http://aka.ms/forbes-cloud-2015>
- 2 Intel/McAfee, "Net Losses: Estimating the Global Cost of Cybercrime," June 2014. <http://aka.ms/mcafee-cybercrime-report>
- 3 UK Dept. for Business, Innovation and Skills, "2014 Information Security Breaches Survey," http://aka.ms/uk-gov_breach-survey
- 4 PWC, "Global State of Information Security Survey: 2015," <http://aka.ms/pwc-cybercrime>
- 5 Gemalto, 2014 Breach Level Index Report
- 6 McKinsey & Company, report for World Economic Forum, Jan. 2014

Still, many CIOs hesitate to fully embrace a cloud-first approach. Their hesitation stems in part from anxiety over a wide range of privacy and security related issues. Large-scale data breaches dominated headlines in 2014 and continue in the news today, raising a critical question for IT leaders everywhere: How can organizations build scalable cloud solutions and increase business agility while taking the necessary steps to secure our data and ensure privacy and compliance across the enterprise?

Without a clear answer, security concerns threaten to stall innovation and stifle business growth. IT and business leaders need a trusted partner to bridge the gap between innovation and security. With the right technologies and processes, even the most complex enterprise can move to the cloud with confidence.

“71% of strategic buyers cite scalability, cost and business agility as the most important drivers for using cloud services.”

Gigaom Research

What customers want from cloud providers

Every business has different needs and every business will reap distinct benefits from cloud solutions. Still, customers of all kinds have the same basic concerns about moving to the cloud. They want to retain control of their data, and they want that data to be kept secure and private, all while maintaining transparency and compliance.

Secure our data. The scale and scope of intrusions are growing. In 2014, cyber criminals compromised more than a billion data records in more than 1500 breaches.⁷ In a 2014 report for the World Economic Forum⁸, McKinsey & Company estimated the risk of cyberattacks “could materially slow the pace of technology and business innovation with as much as \$3 trillion in aggregate impact.” In any security attack, target organizations are only as safe as their weakest link. If any component is not secured, then the entire system is at risk. While acknowledging that the cloud can provide increased data security and administrative control, IT leaders are still concerned that migrating to the cloud will leave them more vulnerable to hackers than their current in-house solutions.

Keep our data private. Cloud services raise unique privacy challenges for businesses. As companies look to the cloud to save on infrastructure costs and improve their flexibility, they also worry about losing control of where their data is stored, who is accessing it, and how it gets used. Since the revelations of widespread surveillance by the US government in 2013, privacy concerns have become more accentuated, and the cloud has come under greater scrutiny as a result.

Give us control. Even as they take advantage of the cloud to deploy more innovative solutions, companies are very concerned about losing control of their data. The recent disclosures of government agencies accessing customer data, through both legal and extralegal means, make some CIOs wary of storing their data in the cloud. Many companies are therefore looking to choose where their data resides in the cloud and to control what entities have visibility into that data.

Promote transparency. While security, privacy, and control are important to business decision makers, they also want the ability to independently verify how their data is being stored, accessed, and secured. Businesses understand that they cannot control what they cannot see. To create this sort of visibility for customers, cloud providers must offer transparency of their security, privacy and compliance practices and actions to give customers the information they need to make their own decisions.

⁷ Gemalto, 2014 Breach Level Index Report

⁸ McKinsey & Company, for World Economic Forum, Jan. 2014

Maintain compliance. As companies and government agencies expand their use of cloud technologies, the complexity and scope of standards and regulations continues to evolve. Companies need to know that their compliance standards will be met, and that compliance will evolve as regulations change over time.

Microsoft Azure: Built for trust

Microsoft Azure provides cloud services for a wide range of enterprise and government customers. The core of Microsoft Azure provides four primary functions on which customers build and manage virtual environments, applications, and associated configurations.

Microsoft Azure
Unified platform for modern business



Compute



Data storage



Network services



App services

Global physical infrastructure servers/ networks/ datacenters

- Stores over 10 trillion objects
- Handles on average 127,000 requests/second
- Peak of 880,000 requests/second



Microsoft, with its unique experience and scale, delivers these services to many of the world's leading enterprises and government agencies. Today, the Microsoft cloud infrastructure supports over 1 billion customers across our enterprise and consumer services in 140 countries and supports 10 languages and 24 currencies. Drawing on this history and scale, Microsoft has implemented software development with enhanced security, operational management, and threat mitigation practices, helping it to deliver services that achieve higher levels of security, privacy, and compliance than most customers could achieve on their own.

Microsoft shares best practices with government and commercial organizations and engages in broad security efforts through the creation of centers of excellence, including the Microsoft Digital Crimes Unit, Microsoft Security Response Center, and Microsoft Malware Protection Center.

Security: Working to keep customer data safe

Azure can help reduce the cost, complexity, and risk associated with security and compliance in the cloud. A survey funded by Microsoft and performed by ComScore⁹ found that while many organizations have initial concerns about moving to the cloud, a majority of cloud adopters reported that they achieved significant security benefits. These security benefits are reported because few individual organizations can replicate the technology and operational processes that Microsoft uses to help safeguard its enterprise cloud services and comply with a wide range of international standards. When companies use Azure, they benefit from Microsoft's unmatched scale and experience running compliant online services around the globe. Microsoft's expertise becomes the customer's expertise.


Initial concern



Realized benefit



Microsoft Cloud Experience:



Security Design and Operations

Secure cloud solutions are the result of comprehensive planning, innovative design, and efficient operations. Microsoft makes security a priority at every step, from code development to incident response.

Design for security from the ground up. Azure code development adheres to Microsoft's Security Development Lifecycle (SDL). The SDL is a software development process that helps developers build more secure software and address security compliance requirements while reducing development cost. The SDL became central to Microsoft's development practices a decade ago and is shared freely with the industry and customers. It embeds security requirements into systems and software through the planning, design, development, and deployment phases.

⁹ <http://aka.ms/twc-cloud-trust-study>

“We don’t have the resources to respond to security threats all day, 365 days a year, the way that Microsoft does.”

Bo Wandschneider,
CIO and Associate Vice Principal
Queen’s University (Canada)


Enhancing operational security. Azure adheres to a rigorous set of security controls that govern operations and support. Microsoft deploys combinations of preventive, defensive, and reactive controls including the following mechanisms to help protect against unauthorized developer and/or administrative activity:

- Tight access controls on sensitive data, including a requirement for two-factor smartcard-based authentication to perform sensitive operations.
- Combinations of controls that enhance independent detection of malicious activity.
- Multiple levels of monitoring, logging, and reporting.

Additionally, Microsoft conducts background verification checks of certain operations personnel and limits access to applications, systems, and network infrastructure in proportion to the level of background verification.

Assume breach. One key operational best practice that Microsoft uses to harden its cloud services is known as the “assume breach” strategy. A dedicated “red team” of software security experts simulates real-world attacks at the network, platform, and application layers, testing Azure’s ability to detect, protect against, and recover from breaches. By constantly challenging the security capabilities of the service, Microsoft can stay ahead of emerging threats.

Incident management and response. Microsoft has a global, 24x7 incident response service that works to mitigate the effects of attacks and malicious activity. The incident response team follows established procedures for incident management, communication, and recovery, and uses discoverable and predictable interfaces with internal and external partners alike. In the event of a security incident, the security team follows these five phases:



- **Identification:** If an event indicates a security issue, the incident is assigned a severity classification and appropriately escalated within Microsoft.
- **Containment:** The immediate priority of the escalation team is to ensure the incident is contained and data is safe.
- **Eradication:** After the situation is contained, the escalation team moves toward eradicating any damage caused by the security incident and identifies the root cause of the security issue.
- **Recovery:** Software or configuration updates are applied to the system and services are returned to full working capacity.
- **Lessons Learned:** Each security incident is analyzed to ensure the appropriate mitigations are applied to protect against future recurrence.

Infrastructure Protection

Azure infrastructure includes hardware, software, networks, administrative and operations staff, and the physical data centers that house it all. Azure addresses security risks across its infrastructure.

Physical security. Azure runs in geographically distributed Microsoft facilities, sharing space and utilities with other Microsoft Online Services. Each facility is designed to run 24x7x365 and employs various measures to help protect operations from power failure, physical intrusion, and network outages. These datacenters comply with industry standards (such as ISO 27001) for physical security and availability. They are managed, monitored, and administered by Microsoft operations personnel.

Monitoring and logging. Centralized monitoring, correlation, and analysis systems manage the large amount of information generated by devices within the Azure environment, providing continuous visibility and timely alerts to the teams that manage the service. Additional monitoring, logging, and reporting capabilities provide visibility to customers.

Perimeter 	Buildings 	Computer room 
• Security staff around the clock • Facility setback requirements • Barriers • Fencing	• Alarms • Security operations center • Seismic bracing • Security cameras	• Two-factor access control: biometric and card readers • Cameras • Days of backup power

Update management. Security update management helps protect systems from known vulnerabilities. Azure uses integrated deployment systems to manage the distribution and installation of security updates for Microsoft software. Azure uses a combination of Microsoft and third-party scanning tools to run OS, web application, and database scans of the Azure environment.

Antivirus and antimalware. Azure software components must go through a virus scan prior to deployment. Code is not moved to production without a clean and successful virus scan. In addition, Microsoft provides native antimalware on all Azure VMs. Microsoft recommends that customers run some form of antimalware or antivirus on all virtual machines (VMs). Customers can install Microsoft Antimalware for Cloud Services and Virtual Machines or another antivirus solution on VMs, and VMs can be routinely reimaged to clean out intrusions that may have gone undetected.

Penetration testing. Microsoft conducts regular penetration testing to improve Azure security controls and processes. Microsoft understands that security assessment is also an important part of our customers' application development and deployment. Therefore, Microsoft has established a policy for customers to carry out authorized penetration testing on their own—and only their own—applications hosted in Azure.

DDoS Protection. Azure has a defense system against Distributed Denial-of-Service (DDoS) attacks on Azure platform services. It uses standard detection and mitigation techniques. Azure's DDoS defense system is designed to withstand attacks generated from outside and inside the platform.

Network Protection

Azure networking provides the infrastructure necessary to securely connect VMs to one another and to connect on-premises data centers with Azure VMs. Because Azure's shared infrastructure hosts hundreds of millions of active VMs, protecting the security and confidentiality of network traffic is critical.

In the traditional datacenter model, a company's IT organization controls networked systems, including physical access to networking equipment. In the cloud service model, the responsibilities for network protection and management are shared between the cloud provider and the customer. Customers do not have physical access, but they implement the logical equivalent within their cloud environment through tools such as Guest operating system (OS) firewalls, Virtual Network Gateway configuration, and Virtual Private Networks.

Network isolation. Azure is a multitenant service, meaning that multiple customers' deployments and VMs are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's data from that of others. This provides the scale and economic benefits of multitenant services while rigorously preventing customers from accessing one another's data.

Virtual networks. A customer can assign multiple deployments within a subscription to a virtual network and allow those deployments to communicate with each other through private IP addresses. Each virtual network is isolated from other virtual networks.

VPN and Express Route. Microsoft enables connections from customer sites and remote workers to Azure Virtual Networks using Site-to-Site and Point-to-Site VPNs. For even better performance, customers can use an optional ExpressRoute, a private fiber link into Azure data centers that keeps their traffic off the Internet.

Encrypting communications. Built-in cryptographic technology enables customers to encrypt communications within and between deployments, between Azure regions, and from Azure to on-premises data centers.

"If you're resisting the cloud because of security concerns, you're running out of excuses."

FORRESTER

Data Protection

Azure allows customers to encrypt data and manage keys, and safeguards customer data for applications, platform, system and storage using three specific methods: encryption, segregation, and destruction.

Data isolation. Azure is a multitenant service, meaning that multiple customers' deployments and virtual machines are stored on the same physical hardware.

Protecting data at rest. Azure offers a wide range of encryption capabilities, giving customers the flexibility to choose the solution that best meets their needs. Azure Key Vault helps customers easily and cost effectively streamline key management and maintain control of keys used by cloud applications and services to encrypt data.

Protecting data in transit. For data in transit, customers can enable encryption for traffic between their own VMs and end users. Azure protects data in transit, such as between two virtual networks. Azure uses industry standard transport protocols such as TLS between devices and Microsoft datacenters, and within datacenters themselves.

Encryption. Customers can encrypt data in storage and in transit to align with best practices for protecting confidentiality and data integrity. For data in transit, Azure uses industry-standard transport protocols between devices and Microsoft datacenters and within datacenters themselves. You can enable encryption for traffic between your own virtual machines and end users.

Data redundancy. Customers may opt for in-country storage for compliance or latency considerations or out-of-country storage for security or disaster recovery purposes. Data may be replicated within a selected geographic area for redundancy.

Data destruction. When customers delete data or leave Azure, Microsoft follows strict standards for overwriting storage resources before reuse. As part of our agreements for cloud services such as Azure Storage, Azure VMs, and Azure Active Directory, we contractually commit to specific processes for the deletion of data.

“From a security point of view, I think Azure is a demonstrably more secure environment than most banks’ datacenters.”


John Schlesinger,
Chief Enterprise Architect, Temenos (Switzerland)

Identity and Access

Microsoft has strict controls that restrict access to Azure by Microsoft employees. Azure also enables customers to control access to their environments, data and applications.

Enterprise cloud directory. Azure Active Directory is a comprehensive identity and access management solution in the cloud. It combines core directory services, advanced identity governance, security, and application access management. Azure Active Directory makes it easy for developers to build policy-based identity management into their applications. Azure Active Directory Premium includes additional features to meet the advanced identity and access needs of enterprise organizations. Azure Active Directory enables a single identity management capability across on-premises, cloud, and mobile solutions.

Active Directory



Multi-Factor Authentication. Microsoft Azure provides Multi-Factor Authentication (MFA). This helps safeguard access to data and applications and enables regulatory compliance while meeting user demand for a simple sign-in process for both on-premises and cloud applications. It delivers strong authentication via a range of easy verification options—phone call, text message, or mobile app notification—allowing users to choose the method they prefer.

Access monitoring and logging. Security reports are used to monitor access patterns and to proactively identify and mitigate potential threats. Microsoft administrative operations, including system access, are logged to provide an audit trail if unauthorized or accidental changes are made. Customers can turn on additional access monitoring functionality in Azure and use third-party monitoring tools to detect additional threats. Customers can request reports from Microsoft that provide information about user access to their environments.

Privacy: Customers own and control their data

Customers will only use cloud providers in which they have great trust. They must trust that the privacy of their information will be protected, and that their data will be used in a way that is consistent with their expectations.

We build privacy protections into Azure through Privacy by Design, a program which describes how we build and operate products and services to protect privacy. Standards and processes that support Privacy by Design principles include the Microsoft Online Services Privacy Statement (which details Microsoft's core privacy requirements and practices) and the Microsoft Secure Development Lifecycle (which includes addressing privacy requirements).

“The question is no longer: ‘How do I move to the cloud?’ Instead, it’s ‘Now that I’m in the cloud, how do I make sure I’ve optimized my investment and risk exposure?’”



We then back those protections with strong contractual commitments to safeguard customer data, including offering EU Model Clauses (which provides terms covering the processing of personal information), and complying with international standards.

Microsoft uses customer data stored in Azure only to provide the service, including purposes compatible with providing the service. Azure does not use customer data for advertising or similar commercial purposes.

Contractual commitments. Microsoft was the first major cloud service provider to make contractual privacy commitments that help assure the privacy protections built into in-scope Azure services are strong. Among the many commitments that Microsoft supports are:

- **EU Model Clauses.** EU data protection law regulates the transfer of EU customer personal data to countries outside the European Economic Area (EEA). Microsoft offers customers the EU Standard Contractual Clauses that provide specific contractual guarantees around transfers of personal data for in-scope services. Europe’s privacy regulators have determined that the contractual privacy protections Azure delivers to its enterprise cloud customers meet current EU standards for international transfers of data. Microsoft is the first cloud provider to receive this recognition.
- **US-EU Safe Harbor Framework and the US-Swiss Safe Harbor Program.** Microsoft abides by these frameworks set forth by the US Department of Commerce regarding the collection, use, and retention of data from the EEA and Switzerland.
- **ISO/IEC 27018.** Microsoft is the first major cloud provider to adopt the first international code of practice for cloud privacy. ISO/IEC 27018 was developed to establish a uniform, international approach to protecting the privacy of personal data stored in the cloud. The British Standards Institution independently verified that Microsoft Azure is aligned with the guideline’s code of practice. ISO 27018 controls include a prohibition on the use of customer data for advertising and marketing purposes without the customer’s express consent

Restricted access by Microsoft personnel. Access to customer data by Microsoft personnel is restricted. Customer data is only accessed when necessary to support the customer’s use of Azure. This may include troubleshooting aimed at preventing, detecting, or repairing problems affecting the operation of Azure and the improvement of features that involve the detection of, and protection against, emerging and evolving threats to the user (such as malware or spam). When granted, access is controlled and logged. Strong authentication, including the use of multi-factor authentication, helps limit access to authorized personnel only. Access is revoked as soon as it is no longer needed.

Notification of lawful requests for information. Microsoft believes that customers should control their data whether stored on their premises or in a cloud service. We will not disclose Azure customer data to law enforcement except as a customer directs or where required by law. When governments make a lawful demand for Azure customer data from Microsoft, we strive to be principled, limited in what we disclose, and committed to transparency.

- Microsoft does not provide any third party with direct or unfettered access to customer data. Microsoft only releases specific data mandated by the relevant legal demand.
- If a government wants customer data—including for national security purposes—it needs to follow the applicable legal process, meaning it must serve us with a warrant or court order for content or subpoena for account information. If compelled to disclose customer data, we will promptly notify the customer and provide a copy of the demand unless legally prohibited from doing so.

- Microsoft will only respond to requests for specific accounts and identifiers. There is no blanket or indiscriminate access to Microsoft's customer data. Every request is explicitly reviewed by Microsoft's legal team, who ensures that the requests are valid, rejects those that are not, and makes sure we only provide the data specified in the order.

In its commitment to transparency, Microsoft regularly publishes a Law Enforcement Requests Report that discloses the scope and number of requests we receive.

Greater transparency and simplicity of data use policies.

Microsoft keeps customers informed about the processes to protect data privacy and security, including practices and policies. Microsoft also provides the summaries of independent audits of services, which helps customers pursue their own compliance.

“Just as computer users back up their laptops in case they break or are lost, Estonia is working out how to back up the country, in case it is attacked by Russia.”

The Economist,
reporting on Estonia's Azure cloud backup

Customers are in control of their data

For many organizations, the benefits of moving to the cloud are clear. Still, fear of losing control causes their decision makers to hesitate. Where will data be stored? Who owns the organization's data? Who will be accessing the data? And what happens if the organization wants to switch providers? These are all valid questions—questions Microsoft has in mind when making a clear commitment to provide customers with control over their data. This commitment is unique among major cloud service providers.

Customers own their data. This belief is fundamental to the Microsoft approach. When a customer utilizes Azure, they retain exclusive ownership of their data. Microsoft takes steps to protect many types of data.

Microsoft defines customer data as “all data, including all text, sound, video or image files, and software that are provided to Microsoft by, or on behalf of, Customer through use of the Online Service.” For example, this includes data that you upload for storage or processing and applications that you run in Azure.

Customers can access their own customer data at any at any time and for any reason without assistance from Microsoft. Microsoft will not use customer data or derive information from it for advertising. We will use customer data only to provide the service or for purposes compatible with providing the service.

-
- **Customer data** is all data, including all text, sound, video or image files, and software that are provided to Microsoft by or on behalf of the customer through use of Azure. For example, it includes data uploaded for storage or processing and applications uploaded by the customer for hosting on Azure.
 - **Administrator data** is the information about administrators (including account contact and subscription administrators) supplied during signup, purchase, or administration of Azure, such as name, phone number, and email address.
 - **Metadata** includes configuration and technical settings and information. For example, it includes the disk configuration settings for an Azure virtual machine or the database design for an SQL Database. Metadata does not include information from which customer data could be derived.
 - **Access control data** is data that is used to manage access to other types of data or functions within Azure. It includes passwords, security certificates, and other authentication-related data.
-

“Our brand rests on the continuity of our IT systems, which are now more available running in Azure.”

Andrew Goodin,
Global Manager of Information Systems
Zespri International (New Zealand)

Control over data location. When customers entrust their data to Microsoft, they are not giving up control. For many customers, knowing and controlling the location of their data can be an important element of data privacy, compliance, and governance. Microsoft Azure offers an ever-expanding network of data centers across the globe. Most Azure services permit customers to specify the particular geography where their customer data will be stored. Data may be replicated within a selected geographic area for redundancy, but will not be replicated outside it for redundancy.

Encryption key management. To ensure control over encrypted data, customers have the option to generate and manage their own encryption keys, and determine who is authorized to use them. They also have the option to revoke Microsoft's copy of their encryption key, although this may limit Microsoft's ability to troubleshoot or repair problems and security threats.

Role based access control. Microsoft provides an approach allowing customers to restrict system access to authorized users based on role assignment, role authorization, and permission authorization. Tools in multiple Microsoft cloud services support authorization based on a user's role, simplifying access control across defined groups of users.

Control over data destruction. When customers delete data or leave a Microsoft cloud service, Microsoft follows strict standards for overwriting storage resources before reuse, as well physical destruction of decommissioned hardware, including contractual commitments to specific processes for the deletion of data and the destruction of storage hardware.

Transparency

For customers to effectively exercise their right to control their data, they must have access and visibility to that data. They must know where it is stored. They must also know, through clearly stated and readily available policies and procedures, how the cloud provider helps secure customer data, who can access it, and under what circumstances.

Where and how data is stored and used. Microsoft gives Azure customers visibility to where their customer data is stored in an ever-expanding network of datacenters around the globe. Customers can balance the need to store backups at multiple locations in case of a disaster with the need to keep their data out of certain geographies. Microsoft provides clear data maps and geographic boundary information for all datacenters.

How data is secured. Customers have access to up-to-date information regarding security policies and procedures. Microsoft promotes transparency by publishing and adhering to the Security Development Lifecycle.

Who requests access to customer data. Microsoft will never disclose Azure customer data to a government or law enforcement agency except as directed by the customer or where required by law. In response to lawful demands for Azure customer data, Microsoft strives to be principled, limited in disclosure, and committed to transparency. Microsoft regularly publishes a Law Enforcement Requests Report that discloses the scope and number of government requests received.

Breach notification. In the event that customer data is compromised, Microsoft will notify its customers. Azure has comprehensive, transparent policies that govern incident response from identification all the way through to lessons learned.

Audit standards certifications. Rigorous third-party audits, such as those conducted by the British Standards Institute, verify Azure's adherence to the strict security controls these standards mandate. As part of Microsoft's commitment to transparency, customers can verify Azure's implementation of many security controls by requesting audit results from the certifying third parties.

“By 2020 clouds will stop being referred to as ‘public’ and ‘private’. It will simply be the way business is done and IT is provisioned.”



Customer guidance. Microsoft publishes a Security Response Center Progress Report and a Security Intelligence Report to provide customers with insights into the threat landscape, and provide prescriptive guidance for managing risk to protect their assets.

Transparency Centers. Microsoft operates Transparency Centers that provide government customers with the ability to review source code, reassure themselves of its integrity, and confirm there are no back doors.

Compliance: Azure conforms to global standards

Microsoft invests heavily in the development of robust and innovative compliance processes. The Microsoft compliance framework for online services maps controls to multiple regulatory standards. This enables Microsoft to design and build services using a common set of controls, streamlining compliance across a range of regulations today and as they evolve in the future.

Microsoft compliance processes also make it easier for customers to achieve compliance across multiple services and meet their changing needs efficiently. Together, security-enhanced technology and effective compliance processes enable Microsoft to maintain and expand a rich set of third-party certifications. These help customers demonstrate compliance readiness to their customers, auditors, and regulators. As part of its commitment to transparency, Microsoft shares third-party verification results with its customers.

Certifications and attestations. Azure meets a broad set of international as well as regional and industry-specific compliance standards, such as ISO 27001, FedRAMP, SOC 1 and SOC 2. Azure's adherence to the strict security controls contained in these standards is verified by rigorous third-party audits that demonstrate Azure services work with and meet world-class industry standards, certifications, attestations, and authorizations.

Comprehensive, independently verified compliance. Azure is designed with a compliance strategy that helps customers address business objectives and industry standards and regulations. The security compliance framework includes test and audit phases, security analytics, risk management best practices, and security benchmark analysis to achieve certificates and attestations. Microsoft Azure offers the following certifications for all in-scope services.

CDSA. The Content Delivery and Security Association (CDSA) provides a Content Protection and Security (CPS) standard for compliance with anti-piracy procedures governing digital media. Azure passed the CDSA audit, enabling secure workflows for content development and distribution.

CJIS. Any US state or local agency that wants to access the FBI's Criminal Justice Information Services (CJIS) database is required to adhere to the CJIS Security Policy. Azure is the only major cloud provider that contractually commits to conformance with the CJIS Security Policy, which commits Microsoft to adhere to the same requirements that law enforcement and public safety entities must meet.

CSA CCM. The Cloud Security Alliance (CSA) is a nonprofit, member-driven organization with a mission to promote the use of best practices for providing security assurance within the cloud. The CSA Cloud Controls Matrix (CCM) provides detailed information about how Azure fulfills the security, privacy, compliance, and risk management requirements defined in the CCM version 1.2, and is published in the CSA's Security Trust and Assurance Registry (STAR).

EU Model Clauses. Microsoft offers customers EU Standard Contractual Clauses that provide contractual guarantees around transfers of personal data outside of the EU. Microsoft is the first company to receive joint approval from the EU's Article 29 Working Party that the contractual privacy protections Azure delivers to its enterprise cloud customers meet current EU standards for international transfers of data. This ensures that Azure customers can use Microsoft services to move data freely through our cloud from Europe to the rest of the world.

FDA 21 CFR Part 11. The US Food and Drug Administration (FDA) Code of Federal Regulations (CFR) Title 21 Part 11 lists requirements for the security of electronic records of companies that sell food and drugs manufactured or consumed in the United States. The compliance reports produced by Azure's independent third party SSAE and ISO auditors identify the procedural and technical controls established at Microsoft and can be used to satisfy the requirements of CFR Title 21 Part 11. Microsoft is able to show how relevant controls within these reports have an impact on compliance with the FDA 21 CFR 11 regulations.

FedRAMP. Azure has been granted a Provisional Authority to Operate (P-ATO) from the Federal Risk and Authorization Management Program (FedRAMP) Joint Authorization Board (JAB) at a Moderate impact level based upon the FIPS 199 classification. FedRAMP is a US government program that provides a standard approach to security assessment, authorization, and monitoring for cloud services used by federal agencies and thereby saves the taxpayer and individual organizations the time and cost of conducting their own independent reviews.

FERPA. The Family Educational Rights and Privacy Act (FERPA) is a US federal law that protects the privacy of student educational records. Microsoft agrees to use and disclosure restrictions imposed by FERPA.

FIPS 140-2. Azure complies with the Federal Information Processing Standard (FIPS) Publication 140-2, a US government standard that defines a minimum set of security requirements for products and systems that implement cryptography.

HIPAA. The Health Insurance Portability and Accountability Act (HIPAA) is a US federal law that regulates patient Protected Health Information (PHI). Azure offers customers a HIPAA Business Associate Agreement (BAA), stipulating adherence to certain security and privacy provisions in HIPAA and the HITECH Act. To assist customers in their individual compliance efforts, Microsoft offers a BAA to Azure customers as a contract addendum.

IRAP. Azure has been assessed against the Australian Government Information Security Registered Assessors Program (IRAP), which provides assurance for public sector customers that Microsoft has appropriate and effective security controls.

ISO/IEC 27018. Microsoft is the first cloud provider to have adopted the ISO/IEC 27018 code of practice, covering the processing of personal information by cloud service providers.

ISO/IEC 27001/27002:2013. Azure complies with this standard, which defines the security controls required of an information security management system.

MLPS. Multi-Level Protection Scheme (MLPS) is based on the Chinese state standard issued by the Ministry of Public Security. Azure operated by 21Vianet adheres to this standard, which provides assurance for both the management and technical security of cloud systems.

MTCS. Azure has achieved Level-1 certification with the Multi-Tier Cloud Security Standard for Singapore (MTCS SS), a cloud security standard covering areas such as data security, confidentiality, business impact, and operational transparency, developed under the Singapore Information Technology Standards Committee.

PCI DSS. Azure is Level 1 compliant with Payment Card Industry (PCI) Data Security Standards (DSS) version 3.0, the global certification standard for organizations that accept most payments cards, as well store, process, or transmit cardholder data.

SOC 1 and SOC 2. Azure has been audited against the Service Organization Control (SOC) reporting framework for both SOC 1 Type 2 and SOC 2 Type 2. Both reports are available to customers to meet a wide range of US and international auditing requirements.

The SOC 1 Type 2 audit report attests to the design and operating effectiveness of Azure controls. The SOC 2 Type 2 audit included a further examination of Azure controls related to security, availability, and confidentiality. Azure is audited annually to ensure that security controls are maintained.

TCS CCCPPF. Azure operated by 21Vianet is among the first cloud providers in China to pass the Trusted Cloud Service certification developed by the China Cloud Computing Promotion and Policy Forum (CCCPF).

UK G-Cloud. The UK Government G-Cloud is a cloud computing certification for services used by government entities in the United Kingdom. Azure has received OFFICIAL accreditation from the UK Government Pan Government Accreditor.

Additional resources

Azure Trust Center

<http://azure.microsoft.com/trustcenter>

Cloud Security Alliance Cloud Controls Matrix

<https://cloudsecurityalliance.org/research/ccm/>

Microsoft Cloud Security Readiness Tool


<http://www.microsoft.com/trustedcloud>

Microsoft Online Services Privacy Statement

<http://aka.ms/oneservices-privacy>

Microsoft Privacy Practices


<http://aka.ms/privacy-practices>





NOTE: Certain recommendations contained herein may result in increased data, network, or compute resource usage, and increase your license or subscription costs.

© 2015 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.



Google for Work Security and Compliance Whitepaper

How Google protects your data.

Google™ for Work

This whitepaper applies to the
following Google Apps products

*Google Apps for Work, Education, Government,
Nonprofit, Drive for Work, and Google Apps Unlimited*

Table of Contents



Introduction 1

Google Has a Strong Security Culture 2

- Employee background checks
- Security training for all employees
- Internal security and privacy events
- Our dedicated security team
- Our dedicated privacy team
- Internal audit and compliance specialists
- Collaboration with the security research community

Operational Security 4

- Vulnerability management
- Malware prevention
- Monitoring
- Incident management

Technology with Security at Its Core 6

- State-of-the-art data centers
 - Powering our data centers
 - Environmental impact
- Custom server hardware and software
- Hardware tracking and disposal
- A global network with unique security benefits
- Securing data in transit
- Low latency and highly available solution
- Service availability

Independent Third-Party Certifications 10

- ISO 27001
- SOC 2/3
- FISMA

Data Usage 11

- Our philosophy
- No advertising in Google Apps

Data Access and Restrictions 12

- Administrative access
- For customer administrators
- Law enforcement data requests
- Third-party suppliers

Regulatory compliance 14

- Data processing amendment
- EU Data Protection Directive
 - U.S.-EU and U.S.-Swiss Safe Harbor Frameworks
 - EU model contract clauses
- U.S. Health Insurance Portability and Accountability Act (HIPAA)
- U.S. Family Educational Rights and Privacy Act (FERPA)
- Children's Online Privacy Protection Act of 1998 (COPPA)

Empowering Users and Administrators to Improve Security and Compliance 16

User authentication/authorization features

- 2-step verification
- Security Key
- Single sign-on (SAML 2.0)
- OAuth 2.0 and OpenID Connect

Data management features

- Information Rights Management (IRM)
- Drive audit log
- Drive content compliance / alerting
- Trusted domains for Drive sharing

Email security features

- Secure transport (TLS) enforcement
- Phishing prevention
- Email content compliance
- Objectionable content
- Restricted email delivery

eDiscovery features

- Email retention policy
- Legal holds
- Search/discovery
- Evidence export
- Support for third-party email platforms

Securing endpoints

- Mobile device management (MDM)
- Policy-based Chrome browser security
- Chrome device management

Data recovery

- Restore a recently deleted user
- Restore a user's Drive or Gmail data

Security reports

Conclusion 23



Introduction

Cloud computing offers many advantages and conveniences for today's organizations. Employees can work together in documents in real time from their phone or tablet from any location, and communicate instantly with teammates via video, voice, instant message, or email. No longer tied to a single machine, they have the freedom to work together from anywhere, using any device they choose. Meanwhile, their employers don't shoulder the cost or burden of maintaining servers and constantly updating software. It's no surprise, then, that so many organizations around the world are storing their information and getting work done in the cloud.



The growth of the cloud has thrust the issue of security and trust into the spotlight. That's because cloud services operate very differently from traditional on-premises technology. Rather than residing on local servers, content is now managed on Google servers that are part of our global data center network. In the past, organizations felt that they had complete control over how infrastructure was run and who operated it. Organizations moving to the cloud will rely on cloud suppliers to manage the infrastructure, operations, and delivery of services. In this new world, companies will still control company data, but via cloud-based tools and dashboards. Rather than only using desktop computers, users can now access work files on their personal mobile devices. Customers must assess whether the security controls and compliance of any cloud solution meet their individual requirements. Customers must therefore understand how these solutions protect and process their data. The goal of this whitepaper is to provide an introduction to Google's technology in the context of security and compliance.

As a cloud pioneer, Google fully understands the security implications of the cloud model. Our cloud services are designed to deliver better security than many traditional on-premises solutions. We make security a priority to protect our own operations, but because Google runs on the same infrastructure that we make available to our customers, your organization can directly benefit from these protections. That's why we focus on security, and protection of data is among our primary design criteria. Security drives our organizational structure, training priorities and hiring processes. It shapes our data centers and the technology they house. It's central to our everyday operations and disaster planning, including how we address threats. It's prioritized in the way we handle customer data. And it's the cornerstone of our account controls, our compliance audits and the certifications we offer our customers.

This paper outlines Google's approach to security and compliance for Google Apps, our cloud-based productivity suite. Used by more than five million organizations worldwide, from large banks and retailers with hundreds of thousands of people to fast-growing startups, Google Apps for Work and Education includes Gmail, Calendar, Groups, Drive, Docs, Sheets, Slides, Hangouts, Sites, Talk, Contacts and Vault. Google Apps is designed to help teams work together in new, more efficient ways, no matter where members are located or what device they happen to be using.

This whitepaper will be divided into two main sections: security and compliance. The security section will include details on organizational and technical controls regarding how Google protects your data. The second section on compliance will cover how your data is processed and details on how organizations can meet regulatory requirements.

Google Has a Strong Security Culture

Google has created a vibrant and inclusive security culture for all employees. The influence of this culture is apparent during the hiring process, employee onboarding, as part of ongoing training and in company-wide events to raise awareness.

Employee background checks

Before they join our staff, Google will verify an individual's education and previous employment, and perform internal and external reference checks. Where local labor law or statutory regulations permit, Google may also conduct criminal, credit, immigration, and security checks. The extent of these background checks is dependent on the desired position.

Security training for all employees

All Google employees undergo security training as part of the orientation process and receive ongoing security training throughout their Google careers. During orientation, new employees agree to our [Code of Conduct](#), which highlights our commitment to keep customer information safe and secure. Depending on their job role, additional training on specific aspects of security may be required. For instance, the information security team instructs new engineers on topics like secure coding practices, product design and automated vulnerability testing tools. Engineers also attend technical presentations on security-related topics and receive a security newsletter that covers new threats, attack patterns, mitigation techniques and more.

Internal security and privacy events

Google hosts regular internal conferences to raise awareness and drive innovation in security and data privacy, which are open to all employees. Security and privacy is an ever-evolving area, and Google recognizes that dedicated employee engagement is a key means of raising awareness. One example is "Privacy Week," during which Google hosts events across global offices to raise awareness of privacy in all facets, from software development, data handling and policy enforcement to living our [privacy principles](#). Google also hosts regular "Tech Talks" focusing on subjects that often include security and privacy.



Google employs more than 500 full-time security and privacy professionals, who are part of our software engineering and operations division. Our team includes some of the world's foremost experts in information, application and network security.

Our dedicated security team

Google employs more than 500 full-time security and privacy professionals, who are part of our software engineering and operations division. Our team includes some of the world's foremost experts in information, application and network security. This team is tasked with maintaining the company's defense systems, developing security review processes, building security infrastructure and implementing Google's security policies. Google's dedicated security team actively scans for security threats using commercial and custom tools, penetration tests, quality assurance (QA) measures and software security reviews.

Within Google, members of the information security team review security plans for all networks, systems and services. They provide project-specific consulting services to Google's product and engineering teams. They monitor for suspicious activity on Google's networks, address information security threats, perform routine security evaluations and audits, and engage outside experts to conduct regular security assessments. We specifically built a full-time team, known as [Project Zero](#), that aims to prevent targeted attacks by reporting bugs to software vendors and filing them in an external database.

The security team also takes part in research and outreach activities to protect the wider community of Internet users, beyond just those who choose Google solutions. Some examples of this research would be the discovery of the [POODLE SSL 3.0 exploit](#) and [cipher suite weaknesses](#). The security team also publishes security research papers, [available to the public](#). The security team also organizes and participates in [open-source projects](#) and academic conferences.

Our dedicated privacy team

The Google Privacy team operates independently from product development and security organizations, but participates in every Google product launch. The team reviews design documentation and code audits to ensure that privacy requirements are followed. The Privacy team has built a set of automated monitoring tools to help ensure that products with Customer Data operate as designed and in accordance with our privacy policy. They help release products that reflect strong privacy standards: transparent collection of user data and providing users and administrators with meaningful privacy configuration options, while continuing to be good stewards of any information stored on our platform. After products launch, the privacy team oversees automated processes that audit data traffic to verify appropriate data usage. In addition, the privacy team conducts research providing thought leadership on privacy best practices for our emerging technologies.

Internal audit and compliance specialists

Google has a dedicated internal audit team that reviews compliance with security laws and regulations around the world. As new auditing standards are created, the internal audit team determines what controls, processes, and systems are needed to meet them. This team facilitates and supports independent audits and assessments by third parties.

Collaboration with the security research community

Google has long enjoyed a close relationship with the security research community, and we greatly value their help identifying vulnerabilities in Google Apps and other Google products. Our [Vulnerability Reward Program](#) encourages researchers to report design and implementation issues that may put customer data at risk, offering rewards in the tens of thousands of dollars. In Chrome, for instance, we warn users against malware and phishing, and offer rewards for finding security bugs. Due to our collaboration with the research community, we've squashed more than 700 Chrome security bugs and have rewarded more than \$1.25 million — more than \$2 million has been awarded across Google's various vulnerability rewards programs. We publicly [thank these individuals](#) and list them as contributors to our products and services.

Operational Security

Far from being an afterthought or the focus of occasional initiatives, security is an integral part of our operations.

Vulnerability management

Google administrates a vulnerability management process that actively scans for security threats using a combination of commercially available and purpose-built in-house tools, intensive automated and manual penetration efforts, quality assurance processes, software security reviews and external audits. The vulnerability management team is responsible for tracking and following up on vulnerabilities. Once a vulnerability requiring remediation has been identified, it is logged, prioritized according to severity, and assigned an owner. The vulnerability management team tracks such issues and follows up frequently until they can verify that the issues have been remediated.

Google also maintains relationships and interfaces with members of the security research community to track reported issues in Google services and open-source tools. More information about reporting security issues can be found at [Google Application Security](#).



Malware prevention

An effective malware attack can lead to account compromise, data theft, and possibly additional access to a network. Google takes these threats to its networks and its customers very seriously and uses a variety of methods to prevent, detect and eradicate malware. Google helps tens of millions of people every day to protect themselves from harm by showing warnings to users of Google Chrome, Mozilla Firefox and Apple Safari when they attempt to navigate to websites that would steal their personal information or install software designed to take over their computers. Malware sites or email attachments install malicious software on users' machines to steal private information, perform identity theft, or attack other computers. When people visit these sites, software that takes over their computer is downloaded without their knowledge. Google's malware strategy begins with infection prevention by using manual and automated scanners to scour Google's search index for websites that may be vehicles for malware or phishing. Approximately one billion people use [Google's Safe Browsing](#) on a regular basis. Google's Safe Browsing technology examines billions of URLs per day looking for unsafe websites. Every day, we discover thousands of new unsafe sites, many of which are legitimate websites that have been compromised. When we detect unsafe sites, we show warnings on Google Search and in web browsers. In addition to our Safe Browsing solution, Google operates [VirusTotal](#), a free online service that analyzes files and URLs enabling the identification of viruses, worms, trojans and other kinds of malicious content detected by antivirus engines and website scanners. VirusTotal's mission is to help in improving the antivirus and security industry and make the Internet a safer place through the development of free tools and services.

Google makes use of multiple antivirus engines in Gmail, Drive, servers and workstations to help identify malware that may be missed by antivirus signatures.

Monitoring

Google's security monitoring program is focused on information gathered from internal network traffic, employee actions on systems and outside knowledge of vulnerabilities. At many points across our global network, internal traffic is inspected for suspicious behavior, such as the presence of traffic that might indicate botnet connections. This analysis is performed using a combination of open-source and commercial tools for traffic capture and parsing. A proprietary correlation system built on top of Google technology also supports this analysis. Network analysis is supplemented by examining system logs to identify unusual behavior, such as attempted access of customer data. Google security engineers place standing search alerts on public data repositories to look for security incidents that might affect the company's infrastructure. They actively review inbound security reports and monitor public mailing lists, blog posts, and wikis. Automated network analysis helps determine when an unknown threat may exist and escalates to Google security staff, and network analysis is supplemented by automated analysis of system logs.

Google helps tens of millions of people every day to protect themselves from harm by showing warnings to users of Google Chrome, Mozilla Firefox and Apple Safari when they attempt to navigate to websites that would steal their personal information or install software designed to take over their computers.



Incident management

We have a rigorous incident management process for security events that may affect the confidentiality, integrity, or availability of systems or data. If an incident occurs, the security team logs and prioritizes it according to its severity. Events that directly impact customers are assigned the highest priority. This process specifies courses of action, procedures for notification, escalation, mitigation, and documentation. Google's security incident management program is structured around the NIST guidance on handling incidents (NIST SP 800–61). Key staff are trained in forensics and handling evidence in preparation for an event, including the use of third-party and proprietary tools. Testing of incident response plans is performed for key areas, such as systems that store sensitive customer information. These tests take into consideration a variety of scenarios, including insider threats and software vulnerabilities. To help ensure the swift resolution of security incidents, the Google security team is available 24/7 to all employees. If an incident involves customer data, Google or its partners will inform the customer and support investigative efforts via our support team.

Technology with Security at Its Core

Google Apps runs on a technology platform that is conceived, designed and built to operate securely. Google is an innovator in hardware, software, network and system management technologies. We custom-designed our servers, proprietary operating system, and geographically distributed data centers. Using the principles of “defense in depth,” we’ve created an IT infrastructure that is more secure and easier to manage than more traditional technologies.

State-of-the-art data centers

Google's focus on security and protection of data is among [our primary design criteria](#). Google data center physical security features a layered security model, including safeguards like custom-designed electronic access cards, alarms, vehicle access barriers, perimeter fencing, metal detectors, and biometrics, and the data center floor features laser beam intrusion detection. Our data centers are monitored 24/7 by high-resolution interior and exterior cameras that can detect and track intruders. Access logs,

activity records, and camera footage are available in case an incident occurs. Data centers are also routinely patrolled by experienced security guards who have undergone rigorous background checks and training. As you get closer to the data center floor, security measures also increase. Access to the data center floor is only possible via a security corridor which implements multifactor access control using security badges and biometrics. Only approved employees with specific roles may enter. Less than one percent of Googlers will ever step foot in one of our data centers.

Powering our data centers

To keep things running 24/7 and ensure uninterrupted services, Google's data centers feature redundant power systems and environmental controls. Every critical component has a primary and alternate power source, each with equal power. Diesel engine backup generators can provide enough emergency electrical power to run each data center at full capacity. Cooling systems maintain a constant operating temperature for servers and other hardware, reducing the risk of service outages. Fire detection and suppression equipment helps prevent damage to hardware. Heat, fire, and smoke detectors trigger audible and visible alarms in the affected zone, at security operations consoles, and at remote monitoring desks.

Environmental impact

Google reduces environmental impact of running our data centers by designing and building our own facilities. We install smart temperature controls, use “free-cooling” techniques like using outside air or reused water for cooling, and redesign how power is distributed to reduce unnecessary energy loss. To gauge improvements, we calculate the performance of each facility using comprehensive efficiency measurements. We're the first major Internet services company to gain external certification of our high environmental, workplace safety and energy management standards throughout our data centers. Specifically, we received voluntary [ISO 14001](#), [OHSAS 18001](#) and [ISO 50001](#) certifications. In a nutshell, these standards are built around a very simple concept: Say what you're going to do, then do what you say—and then keep improving.

Custom server hardware and software

Google's data centers house energy-efficient custom, purpose-built servers and network equipment that we design and manufacture ourselves. Unlike much commercially available hardware, Google servers don't include unnecessary components such as video cards, chipsets, or peripheral connectors, which can introduce vulnerabilities. Our production servers run a custom-designed operating system (OS) based on a stripped-down and hardened version of Linux. Google's servers and their OS are designed for the sole purpose of providing Google services. Server resources are dynamically allocated, allowing for flexibility in growth and the ability to adapt quickly and efficiently, adding or reallocating resources based on customer demand. This homogeneous environment is maintained by proprietary software that continually monitors systems for binary modifications. If a modification is found that differs from the standard Google image, the system is automatically returned to its official state. These automated, self-healing mechanisms are designed to enable Google to monitor and remediate destabilizing events, receive notifications about incidents, and slow down potential compromise on the network.

Hardware tracking and disposal

Google meticulously tracks the location and status of all equipment within our data centers from acquisition to installation to retirement to destruction, via bar codes and asset tags. Metal detectors and video surveillance are implemented to help make sure no equipment leaves the data center floor without authorization. If a component fails to pass a performance test at any point during its lifecycle, it is removed from inventory and retired. Google hard drives leverage technologies like FDE (full disk encryption) and drive locking, to protect data at rest. When a hard drive is retired, authorized individuals verify that the disk is erased by writing zeros to the drive and performing a multiple-step verification process to ensure the drive contains no data. If the drive cannot be erased for any reason, it is stored securely until it can be physically destroyed. Physical destruction of disks is a multistage process beginning with a crusher that deforms the drive, followed by a shredder that breaks the drive into small pieces, which are then recycled at a secure facility. Each data center adheres to a strict disposal policy and any variances are immediately addressed.

A global network with unique security benefits

Google's IP data network consists of our own fiber, public fiber, and undersea cables. This allows us to deliver highly available and low latency services across the globe.

In other cloud services and on-premises solutions, customer data must make several journeys between devices, known as "hops," across the public Internet. The number of hops depends on the distance between the customer's ISP and the solution's data center. Each additional hop introduces a new opportunity for data to be attacked or intercepted. Because it's linked to most ISPs in the world, Google's global network improves the security of data in transit by limiting hops across the public Internet.

Defense in depth describes the multiple layers of defense that protect Google's network from external attacks. Only authorized services and protocols that meet our security requirements are allowed to traverse it; anything else is automatically dropped. Industry-standard firewalls and access control lists (ACLs) are used to enforce network segregation. All traffic is routed through custom GFE (Google Front End) servers to detect and stop malicious requests and Distributed Denial of Service (DDoS) attacks. Additionally, GFE servers are only allowed to communicate with a controlled list of servers internally; this "default deny" configuration prevents GFE servers from accessing unintended resources. Logs are routinely examined to reveal any exploitation of programming errors. Access to networked devices is restricted to authorized personnel.

Google's IP data network consists of our own fiber, public fiber, and undersea cables. This allows us to deliver highly available and low latency services across the globe.

Securing data in transit

Data is most vulnerable to unauthorized access as it travels across the Internet or within networks. For this reason, securing data in transit is a high priority for Google. Data traveling between a customer's device and Google is encrypted using HTTPS/TLS (Transport Layer Security). In fact, Google was the first major cloud provider to enable HTTPS/TLS by default. When sending to or receiving email from a non-Google user, all links of the chain (device, browser, provider of the email service) have to be strong and work together to make encryption work. We believe this is so important that we report on the industry's adoption of TLS on our [safe email site](#).

Google has also upgraded all our RSA certificates to 2048-bit keys, making our encryption in transit for Google Apps and all other Google services even stronger.

Perfect forward secrecy (PFS) minimizes the impact of a compromised key, or a cryptographic breakthrough. It protects network data by using a short-term key that lasts only a couple of days and is only held in memory, rather than a key that's used for years and kept on durable storage. Google was the first major web player to enable perfect forward secrecy by default.

Google encrypts all Google Apps data as it moves between our data centers on our private network.

Low latency and highly available solution

Google designs the components of our platform to be highly redundant. This redundancy applies to our server design, how we store data, network and Internet connectivity, and the software services themselves. This "redundancy of everything" includes the handling of errors by design and creates a solution that is not dependant on a single server, data center, or network connection. Google's data centers are geographically distributed to minimize the effects of regional disruptions such as natural disasters and local outages. In the event of hardware, software, or network failure, data is automatically shifted from one facility to another so that Google Apps customers can continue working in most cases without interruption. Customers with global workforces can collaborate on documents, video conferencing and more without additional configuration or expense. Global teams share a highly performant and low latency experience as they work together on a single global network.

Google's highly redundant infrastructure also helps protect our customers from data loss. For Google Apps, our recovery point objective (RPO) target is zero, and our recovery time objective (RTO) design target is also zero. We aim to achieve these targets through live or synchronous replication:

Google's data centers are geographically distributed to minimize the effects of regional disruptions such as natural disasters and local outages.



actions you take in Google Apps Products are simultaneously replicated in two data centers at once, so that if one data center fails, we transfer your data over to the other one that's also been reflecting your actions. Customer data is divided into digital pieces with random file names. Neither their content nor their file names are stored in readily human-readable format, and stored customer data cannot be traced to a particular customer or application just by inspecting it in storage. Each piece is then replicated in near-real time over multiple disks, multiple servers, and multiple data centers to avoid a single point of failure. To further prepare for the worst, we conduct disaster recovery drills in which we assume that individual data centers—including our corporate headquarters—won't be available for 30 days. We regularly test our readiness for plausible scenarios as well as more imaginative crises, like alien and zombie invasions.

Our highly redundant design has allowed Google to achieve an uptime of 99.984% for Gmail for the last years with **no scheduled downtime**. Simply put, when Google needs to service or upgrade our platform, users do not experience downtime or maintenance windows.

Service availability

Some of Google's services may not be available in some jurisdictions. Often these interruptions are temporary due to network outages, but others are permanent due to government-mandated blocks. Google's Transparency Report also shows [recent and ongoing disruptions of traffic](#) to Google products. We provide this data to help the public analyze and understand the availability of online information.

Independent Third-Party Certifications

Google's customers and regulators expect independent verification of our security, privacy, and compliance controls. In order to provide this, we undergo several independent third-party audits on a regular basis. For each one, an independent auditor examines our data centers, infrastructure, and operations. Regular audits are conducted to certify our compliance with the auditing standards ISO 27001, SOC 2 and SOC 3, as well as with the U.S. Federal Information Security Modernization Act of 2014 (FISMA) for Google Apps for Government. When customers consider Google Apps, these certifications can help them confirm that the product suite meets their security, compliance and data processing needs.

ISO 27001

ISO 27001 is one of the most widely recognized and accepted independent security standards. Google has earned it for the systems, technology, processes, and data centers that run Google Apps. Our compliance with the international standard was certified by Ernst & Young CertifyPoint, an ISO certification body accredited by the Dutch Accreditation Council (a member of the International Accreditation Forum, or IAF). Our ISO 27001 certificate and scoping document are available in [our Trust Center](#).

SOC 2/3

In 2014, the American Institute of Certified Public Accountants (AICPA) Assurance Services Executive Committee (ASEC) released the revised version of the Trust Services Principles and Criteria (TSP). SOC (Service Organization Controls) is an audit framework for non-privacy principles that include security, availability, processing integrity, and confidentiality. Google has both SOC 2 and SOC 3 reports. Our SOC 3 report is available for [download](#) without a nondisclosure agreement. The SOC 3 confirms our compliance with the principles of security, availability, processing integrity and confidentiality.

FISMA

FISMA is a U.S. federal law pertaining to the information security of federal agencies' information systems. The law requires agencies to ensure that their systems and those operated by service providers such as Google meet minimum security requirements specified by the National Institute of Standards and Technology (NIST). Google, as a cloud provider, has one of the longest track records in meeting these requirements, and we maintain a current authorization to operate (ATO) for Google Apps for Government. The Federal Risk and Authorization Management Program (FedRAMP) adds a number of requirements on top of the FISMA requirements that cloud providers must meet. For more information on the FedRAMP program and vendor status, please visit [fedramp.gov](#).

Data Usage

Our philosophy

Google Apps customers own their data, not Google. The data that Google Apps organizations and users put into our systems is theirs, and we do not scan it for advertisements nor sell it to third parties. We offer our customers a detailed [data processing amendment](#) that describes our commitment to protecting customer data. It states that Google will not process data for any purpose other





than to fulfill our contractual obligations. Furthermore, if customers delete their data, we commit to deleting it from our systems within 180 days. Finally, we provide tools that make it easy for customer administrators to take their data with them if they choose to stop using our services, without penalty or additional cost imposed by Google.

No advertising in Google Apps

There is **no** advertising in the [Google Apps Core Services](#), and we have no plans to change this in the future. Google does not collect, scan or use data in Google Apps Core Services for advertising purposes. Customer administrators can restrict access to Non-Core Services from the Google Apps Admin console. Google indexes customer data to provide beneficial services, such as spam filtering, virus detection, spellcheck and the ability to search for emails and files within an individual account.

Data Access and Restrictions

Administrative access

To keep data private and secure, Google logically isolates each customer's Google Apps data from that of other customers and users, even when it's stored on the same physical server. Only a small group of Google employees have access to customer data. For Google employees, access rights and levels are based on their job function and role, using the concepts of least-privilege and need-to-know to match access privileges to defined responsibilities. Google employees are only granted a limited set of default permissions to access company resources, such as employee email and Google's internal employee portal. Requests for additional access follow a formal process that involves a request and an approval from a data or system owner, manager, or other executives, as dictated by Google's security policies. Approvals are managed by workflow tools that maintain audit records of all changes. These tools control both the modification of authorization settings and the approval process to ensure consistent application of the approval policies. An employee's authorization settings are used to control access to all resources, including data and systems for Google Apps products. Support services are only provided to authorized customer administrators whose identities have been verified in several ways. Googler access is monitored and audited by our dedicated security, privacy, and internal audit teams.

For customer administrators

Within customer organizations, administrative roles and privileges for Google Apps are configured and controlled by the customer. This means that individual team members can manage certain services or perform specific administrative functions without gaining access to all settings and data. Integrated audit logs offer a detailed history of administrative actions, helping customers monitor internal access to data and adherence to their own policies.

Law enforcement data requests

The customer, as the data owner, is primarily responsible for responding to law enforcement data requests; however, like other technology and communications companies, Google may receive direct requests from governments and courts around the world about how a person has used the company's services. We take measures to protect customers' privacy and limit excessive requests while also meeting our legal obligations. Respect for the privacy and security of data you store with Google remains our priority as we comply with these legal requests. When we receive such a request, our team reviews the request to make sure it satisfies legal requirements and Google's policies. Generally speaking, for us to comply, the request must be made in writing, signed by an authorized official of the requesting agency and issued under an appropriate law. If we believe a request is overly broad, we'll seek to narrow it, and we push back often and when necessary. For example, in 2006 Google was the only major search company that refused a U.S. government request to hand over two months of user search queries. We objected to the subpoena, and eventually a court denied the government's request. In some cases we receive a request for all information associated with a Google account, and we may ask the requesting agency to limit it to a specific product or service. We believe the public deserves to know the full extent to which governments request user information from Google. That's why we became the first company to start regularly publishing reports about government data requests. Detailed information about data requests and Google's response to them is available in our [Transparency Report](#). It is Google's policy to notify customers about requests for their data unless specifically prohibited by law or court order.

We believe the public deserves to know the full extent to which governments request user information from Google. That's why we became the first company to start regularly publishing reports about government data requests.

Third-party suppliers

Google directly conducts virtually all data processing activities to provide our services. However, Google may engage some [third-party suppliers](#) to provide services related to Google Apps, including customer and technical support. Prior to onboarding third-party suppliers, Google conducts an assessment of the security and privacy practices of third-party suppliers



to ensure they provide a level of security and privacy appropriate to their access to data and the scope of the services they are engaged to provide. Once Google has assessed the risks presented by the third-party supplier, the supplier is required to enter into appropriate security, confidentiality, and privacy contract terms.

Regulatory Compliance

Our customers have varying regulatory [compliance](#) needs. Our clients operate across regulated industries, including finance, pharmaceutical and manufacturing.

Google contractually commits to the following:

- Google will maintain adherence to ISO 27001 and SOC 2/3 audits during the term of the agreement.
- Defined Security Standards. Google will define how data is processed, stored, and protected through specific defined security standards.
- Access to our Data Privacy Officer. Customers may contact Google's Data Privacy Officer for questions or comments.
- Data Portability. Administrators can export customer data in [standard formats](#) at any time during the term of the agreement. Google does not charge a fee for exporting data.

Data processing amendment

Google takes a global approach to our commitments on data processing. Google and many of our customers operate in a global environment. We offer **all** of our users the same high level of protections via our [data processing amendment](#). The commitments in our data processing amendment are designed to facilitate compliance with jurisdictional-specific laws or regulations. Your organization can opt into our data processing amendment by following the instructions in our [Help Center](#).

EU Data Protection Directive

The Article 29 Working Party is an independent European advisory body focused on data protection and privacy. They have provided guidance on how to meet European data privacy requirements when engaging with cloud computing providers. Google provides capabilities and contractual commitments created to meet data protection recommendations provided by the Article 29 Working Party.

U.S.-EU and U.S.-Swiss Safe Harbor Frameworks

More than half of Google's business customers are based outside of the United States, with many of them operating in Europe. These businesses must comply with the European Commission's Data Protection Directive, which regulates the transfer of personal data within the European Union. The [U.S.-EU Safe Harbor Framework](#) provides a method for European companies to transfer personal data outside the European Union in a manner that's consistent with the Directive. Google [certifies](#) that it adheres to its principles and to those of the U.S.-Swiss Safe Harbor Framework.

EU model contract clauses

In 2010, the European Commission approved model contract clauses as a means of compliance with the requirements of the Directive. The effect of this decision is that by incorporating certain provisions into a contract, personal data can flow from those subject to the Directive to providers outside the EU or the European Economic Area. Google has a broad customer base in Europe. By adopting [EU model contract clauses](#), we're offering customers an additional option for compliance with the Directive.

U.S. Health Insurance Portability and Accountability Act (HIPAA)

Google Apps supports our customers' compliance with the U.S. Health Insurance Portability and Accountability Act (HIPAA), which governs the confidentiality and privacy of protected health information (PHI). Customers who are subject to HIPAA and wish to use Google Apps with PHI must sign a [business associate agreement \(BAA\)](#) with Google. The BAA covers Gmail, Google Calendar, Google Drive, Google Sites and Google Apps Vault.

U.S. Family Educational Rights and Privacy Act (FERPA)

More than 30 million students rely on Google Apps for Education. Google Apps for Education services comply with FERPA (Family Educational Rights and Privacy Act) and our commitment to do so is included in our agreements.

Children's Online Privacy Protection Act of 1998 (COPPA)

Protecting children online is important to us. We contractually require Google Apps for Education schools to obtain parental consent that COPPA calls for to use our services, and our services can be used in compliance with COPPA.



Empowering Users and Administrators to Improve Security and Compliance

Google builds security into its structure, technology, operations and approach to customer data. Our robust security infrastructure and systems become the default for each and every Google Apps customer. But beyond these levels, users are actively empowered to enhance and customize their individual security settings to meet their business needs through dashboards and account security wizards. Google Apps also offers administrators full control to configure infrastructure, applications and system integrations in a single dashboard via our Admin console — regardless of the size of the organization. This approach simplifies administration and configuration. Consider deployment of DKIM (a phishing prevention feature) in an on-premise email system. Administrators would need to patch and configure every server separately, and any misconfiguration would cause a service outage. Using our Admin console, DKIM is configured in minutes across thousands or hundreds of thousands of accounts with peace of mind and no outage or maintenance window required. Administrators have many powerful tools at their disposal, such as authentication features like 2-step verification and single sign-on, and email security policies like secure transport (TLS) enforcement, which can be configured by organizations to meet security and system integration requirements. Below are some key features that can help customize Google Apps for your security and compliance needs:

User authentication/authorization features

2-step verification

[2-step verification](#) adds an extra layer of security to Google Apps accounts by requiring users to enter a verification code in addition to their username and password when they sign in. This can greatly reduce the risk of unauthorized access if a user's password is compromised. Verification codes are delivered on a one-time basis to a user's Android, BlackBerry, iPhone, or other mobile phone. Administrators can choose to turn on 2-step verification for their domain at any time.

Security Key

[Security Key](#) is an enhancement for 2-step verification. Google, working with the FIDO Alliance standards organization, developed the Security Key — an actual physical key used to access your Google Account. It sends an encrypted signature rather than a code, and helps ensure that your login cannot be phished. Google for Work admins will be able to easily deploy, monitor and manage the Security Key at scale with new controls in the Admin console with no additional software to install. IT admins will see where and when employees last used their keys with usage tracking and reports. If Security Keys are lost, admins can easily revoke access to those keys and provide backup codes so employees can still sign-in and get work done.

Single sign-on (SAML 2.0)

Google Apps offers customers a [single sign-on \(SSO\) service](#) that lets users access multiple services using the same sign-in page and authentication credentials. It is based on SAML 2.0, an XML standard that allows secure web domains to exchange user authentication and authorization data. For additional security, SSO accepts public keys and certificates generated with either the RSA or DSA algorithm. Customer organizations can use the SSO service to integrate single sign-on for Google Apps into their LDAP or other SSO system.

OAuth 2.0 and OpenID Connect

Google Apps supports [OAuth 2.0 and OpenID Connect](#), an open protocol for authentication and authorization. This allows customers to configure one single sign-on service (SSO) for multiple cloud solutions. Users can log on to third-party applications through Google Apps—and vice versa—without re-entering their credentials or sharing sensitive password information.

Data management features

Information Rights Management (IRM)

With [Information Rights Management](#) ("IRM") you can disable downloading, printing and copying from the advanced sharing menu — perfect for when the file you're sharing is only meant for a few select people. This new option is available for any file stored in Google Drive, including documents, spreadsheets and presentations created in Google Docs.

Google Apps also offers administrators full control to configure infrastructure, applications and system integrations in a single dashboard via our Admin console — regardless of the size of the organization.

Drive audit log

The [Drive audit log](#) lists every time your domain's users view, create, update, delete or share Drive content. This includes content you create in Google Docs, Sheets, Slides and other Google Apps, as well as content created elsewhere that you upload to Drive, such as PDFs and Word files.

Drive content compliance / alerting

Google Apps for Work has [an additional feature](#) that allows Administrators to keep track of when specific actions are taken in Drive and can set up [custom Drive alerts](#). So if you want to know when a file containing the word "confidential" in the title is shared outside the company, now you'll know. And there are more events coming to Drive audit, including download, print and preview alerts.

Trusted domains for Drive sharing

Google Apps for Work and Education administrators will allow for domain [whitelisting](#). End users can share to those trusted domains, but can't share to other external domains. Great for partnerships, subsidiaries or other arrangements where certain domains are trusted and users are allowed to share to them.

Email Security features

Secure transport (TLS) enforcement

Google Apps administrators can require that email to or from specific domains or email addresses be encrypted with [Transport Layer Security \(TLS\)](#). For instance, a customer organization may choose to transmit all messages to its outside legal counsel via a secure connection. If TLS is not available at a specified domain, inbound mail will be rejected and outbound mail will not be transmitted.

Phishing prevention

Spammers can sometimes forge the "From" address on an email message so that it appears to come from a reputable organization's domain. Known as [phishing](#), this practice is often an attempt to collect sensitive data. To help prevent phishing, Google participates in the [DMARC program](#), which lets domain owners tell email providers how to handle unauthenticated messages from their domain. Google Apps customers can implement DMARC by creating a DMARC record within their admin settings and implementing an SPF record and DKIM keys on all outbound mail streams.

Email content compliance

Administrators can choose to scan Google Apps email messages for [predefined sets of words, phrases, text patterns or numerical patterns](#). They can create rules that either reject matching emails before they reach their intended recipients or deliver them with modifications. Customers have used this setting to monitor sensitive or restricted data, such as credit card information, internal project code names, URLs, telephone numbers, employee identification numbers, and social security numbers.

Objectionable content

The [objectionable content](#) setting enables administrators to specify what action to perform for messages based on custom word lists. With objectionable content policies, administrators choose whether messages containing certain words (such as obscenities) are rejected or delivered with modifications; for example, to notify others when the content of a message matches the rules that you set. Administrators can also configure this setting to reject outbound emails that may contain sensitive company information; for example, by setting up an outbound filter for the word *confidential*.

Restricted email delivery

By default, users with Gmail accounts at your domain can send mail to and receive mail from any email address. However, in some cases, administrators may want to [restrict the email addresses](#) your users can exchange mail with. For example, a school might want to allow its students to exchange mail with the faculty and other students, but not with people outside of the school. Use the Restrict delivery setting to allow the sending or receiving of email messages only from addresses or domains that administrators specify. When administrators add a Restrict delivery setting, users cannot communicate with anyone, except those authorized. Users who attempt to send mail to a domain not listed will see a message that specifies a policy prohibiting mail to that address, confirming that the mail is unsent. Users receive only authenticated messages from listed domains. Messages sent from unlisted domains—or messages from listed domains that can't be verified using DKIM or SPF records—are returned to the sender with a message about the policy.

Google Apps administrators can require that email to or from specific domains or email addresses be encrypted with Transport Layer Security (TLS).

eDiscovery features

eDiscovery allows organizations to stay prepared in case of lawsuits and other legal matters. [Google Vault](#) is the eDiscovery solution for Google Apps that lets customers retain, archive, search and export their business Gmail. Administrators can also search and export files stored in Google Drive.

Email retention policy

[Retention rules](#) control how long certain messages in your domain are retained before they are removed from user mailboxes and expunged from all Google systems. Google Apps allows you to set a default retention rule for your entire domain. For more advanced implementations, [Google Vault](#) allows administrators to create custom retention rules to retain specific content. This advanced configuration allows administrators to specify the number of days to retain messages, whether to delete them permanently after their retention periods, whether to retain messages with specific labels, and whether to let users manage email deletion themselves.

Legal holds

[Google Vault](#) allows administrators to place [legal holds](#) on users to preserve all their emails and on-the-record chats indefinitely in order to meet legal or other retention obligations. You can place legal holds on all content in a user's account, or target specific content based on dates and terms. If a user deletes messages that are on hold, the messages are removed from the user's view, but they are not deleted from Google servers until the hold is removed.

Search/discovery

[Google Vault](#) allows administrators to [search Gmail and Drive accounts](#) by user account, organizational unit, date or keyword. Search results include email, on-the-record chats, Google file types and non-Google file types such as PDF, DOCX and JPG.

Evidence export

[Google Vault](#) allows administrators to have the ability to [export](#) specific email, on-the-record chats and files to standard formats for additional processing and review in a manner that supports legal matters while respecting chain of custody guidelines.

Support for third-party email platforms

The [comprehensive mail storage setting](#) ensures that a copy of all sent or received mail in your domain—including mail sent or received by non-Gmail mailboxes—is stored in the associated users' Gmail mailboxes. For organizations that reroute mail to non-Gmail mail servers, this setting also ensures storage of mail in Gmail mailboxes for archiving and eDiscovery purposes.

Securing endpoints

Mobile device management (MDM)

[Mobile device management](#) in Google Apps eliminates the need for on-premises device or third-party management solutions.

Administrators can enforce policies over mobile devices in their organization, encrypt data on devices, and perform actions like remotely wiping or locking lost or stolen devices. This type of control helps ensure the security of business data, even if employees choose to work on their personal phones and tablets. Mobile device management in Google Apps works with Android, iOS, Windows Phone, and smartphones and tablets using Microsoft Exchange ActiveSync, such as BlackBerry 10.

Policy-based Chrome browser security

All of the tools and features in Google Apps are best supported by Google Chrome. **Administrators can apply [security and usage policies](#) across Windows, OSX, Linux, iOS, and Android.**

Chrome's standard security features include Safe Browsing, sandboxing, and managed updates that protect users from malicious sites, viruses, malware, and phishing attacks. There are also measures in place to prevent cross-site scripting, which attackers can use to steal private data. Google Apps administrators can deploy Chrome for Work across their organization and customize it to meet their needs.

Over [280 policies](#) help administrators control how employees use Chrome across devices. For example, administrators can enable automatic updates to get the latest security fixes, block or allow specific apps, and configure support for legacy browsers.

Administrators can enforce policies over mobile devices in their organization, encrypt data on devices, and perform actions like remotely wiping or locking lost or stolen devices.

Chrome device management

The Google Apps Admin Console applies policy to [Chrome devices](#) such as Chromebooks, Chromeboxes, and Chromebox for meetings, which are fast, secure, and cost-effective computers that run Chrome as an operating system. Administrators can easily manage security and other settings for their organization's Chrome devices from a single place. They can configure Chrome features for their users, set up access to VPNs and WiFi networks, pre-install apps and extensions, restrict sign-in to certain users, and more.

Data Recovery

Restore a recently deleted user

An administrator can [restore a deleted user account](#) for up to five days after date of deletion. After five days, the Admin console permanently deletes the user account, and it can't be restored, even if you contact Google technical support. Please note that only customer Administrators can delete accounts.

Restore a user's Drive or Gmail data

An administrator can restore a user's Drive or Gmail data for up to 25 days after the data is removed from the user's trash. Google will delete all Customer-deleted data from its systems as soon as reasonably practicable and within a maximum period of 180 days.

An administrator can restore a user's Drive or Gmail data for up to 25 days after date of deletion.

Security reports

Google Apps administrators have access to [security reports](#) that provide vital information on their organization's exposure to data compromise. They can quickly discover which particular users pose security risks by eschewing 2-step verification, installing external apps, or sharing documents indiscriminately. Administrators can also choose to [receive alerts](#) when suspicious login activity occurs, indicating a possible security threat.

Conclusion

The protection of user data is a primary design consideration for all of Google's infrastructure, applications and personnel operations. Protection of user data is far from being an afterthought or the focus of occasional initiatives, it's an integral part of what we do. We believe that Google can offer a level of protection that very few can match. Because protecting your data is part of our core business, Google can develop security innovations such as 2-step authentication and stronger encryption methods. We are able to make extensive investments in security, resources and expertise at a scale that few can afford. Our scale of operations and collaboration with the security research community enable Google to address vulnerabilities quickly or prevent them entirely. Google's security and operational procedures are verified by independent third-party auditors.

Data protection is more than just security. Google's strong contractual commitments make sure our customers maintain control over the data and how it is processed, including the assurance that your data is not used for advertising or any other purpose than to deliver Google Apps services.

For these reasons and more, over five million organizations across the globe, including 64 percent of the Fortune 500, trust Google with their most valuable asset: their information. Google will continue to invest in our platform to allow our users to benefit from our services in a secure and transparent manner.

Data protection is more than just security. Google's strong contractual commitments make sure our customers maintain control over the data and how it is processed, including the assurance that your data is not used for advertising or any other purpose than to deliver Google Apps services.



Google Apps Service Level Agreement

Google Apps SLA. During the Term of the applicable Google Apps Agreement (the "Agreement"), the Google Apps Covered Services web interface will be operational and available to Customer at least 99.9% of the time in any calendar month (the "Google Apps SLA"). If Google does not meet the Google Apps SLA, and if Customer meets its obligations under this Google Apps SLA, Customer will be eligible to receive the Service Credits described below. This Google Apps SLA states Customer's sole and exclusive remedy for any failure by Google to meet the Google Apps SLA.

Definitions. The following definitions shall apply to the Google Apps SLA.

"Downtime" means, for a domain, if there is more than a five per cent user error rate. Downtime is measured based on server-side error rate.

"Google Apps Covered Services" means the Gmail, Google Calendar, Google Talk, Google Docs and Drive, Google Groups, Google Sites and Google Apps Vault components of the Service. This does not include the Gmail Labs functionality, Google Apps – Postini Services, Gmail Voice or Video Chat components of the Service.

"Monthly Uptime Percentage" means total number of minutes in a calendar month minus the number of minutes of Downtime suffered in a calendar month, divided by the total number of minutes in a calendar month.

"Service" means the Google Apps for Business service (also known as Google Apps Premier Edition), Google Apps for Government service, Google Apps for ISPs service (also known as Google Apps Partner Edition), Google Apps for Education service (also known as Google Apps Education Edition) or Google Apps Vault (as applicable) provided by Google to Customer under the Agreement.

"Service Credit" means the following:

Monthly Uptime Percentage	Days of Service added to the end of the Service term (or monetary credit equal to the value of days of service for monthly postpay billing customers), at no charge to Customer
< 99.9% - >= 99.0%	3
< 99.0% - >= 95.0%	7
< 95.0%	15

Customer Must Request Service Credit. In order to receive any of the Service Credits described above, Customer must notify Google within thirty days from the time Customer becomes eligible to receive a Service Credit. Failure to comply with this requirement will forfeit Customer's right to receive a Service Credit.

Maximum Service Credit. The aggregate maximum number of Service Credits to be issued by Google to Customer for all Downtime that occurs in a single calendar month shall not exceed fifteen days of Service added to the end of Customer's term for the Service (or the value of 15 days of service in the form of a monetary credit to a monthly billing customer's account). Service Credits may not be exchanged for, or converted to, monetary amounts, except for customers who are on Google's monthly billing plan.

Google Apps SLA Exclusions. The Google Apps SLA does not apply to any services that expressly exclude this Google Apps SLA (as stated in the documentation for such services) or any performance issues: (i) caused by factors described in the "Force Majeure" section of the Agreement; or (ii) that resulted from Customer's equipment or third-party equipment, or both (not within the primary control of Google).

Appendix A

Additional Use Rights and Restrictions

These Additional Use Rights and Restrictions shall apply to each Enterprise's use of Products and must be incorporated in each Government Contract.

1. General.

Right to use. Purchasing Entity may access and use In-Scope Services, and install and use a Client (if any) included with its Subscription, only as described in this agreement. All other rights are reserved.

Acceptable use. Purchasing Entity will use In-Scope Services only per the AUP. Purchasing Entity will not use In-Scope Services in any way that infringes a third party's patent, copyright, or trademark or misappropriates its trade secret. Purchasing Entity may not reverse engineer, decompile, work around technical limits in, or disassemble In-Scope Services, except if applicable law permits despite this limit. Purchasing Entity may not rent, lease, lend, resell, transfer, or host In-Scope Services to or for third parties.

Compliance.

Purchasing Entity will comply with all laws and regulations applicable to its use of In-Scope Services.

In providing In-Scope Services, Contractor and Microsoft will comply with all laws and regulations (including applicable security breach notification law) that generally apply to IT service providers. Microsoft will comply with all laws and regulations applicable to its provision of the Online Services, including security breach notification law. However, Microsoft is not responsible for compliance with any laws or regulations applicable to Customer or Customer's industry or government function that are not generally applicable to information technology service providers. Microsoft does not determine whether Customer Data includes information subject to any specific law or regulation. All Security Incidents are subject to the Security Incident Notification terms below.

Purchasing Entity will obtain any consents required: (1) to allow it to access, monitor, use, and disclose user data; and (2) for Contractor and Microsoft to provide the In-Scope Services. If Purchasing Entity is an educational institution, it will obtain any parental consent for end users' use of In-Scope Services as required by applicable law.

Customer Data. Customer Data is used only to provide Purchasing Entity In-Scope Services. This use may include troubleshooting to prevent, find and fix problems with those In-Scope Services' operation. It may also include improving features for finding and protecting against threats to users. Neither Contractor nor Microsoft will derive information from Customer Data for any advertising or other commercial purposes. As between the parties, the Purchasing Entity retains all right, title and interest in and to Customer Data. Neither Contractor nor Microsoft acquires rights in Customer Data, other than the rights Purchasing Entity grants to Microsoft through Contractor to provide the Online Services to Purchasing Entity. This paragraph does not affect Microsoft's rights in software or Online Services Contractor and Microsoft license to Purchasing Entity.

Contractor and Microsoft will enable Purchasing Entity to keep Customer Data in DPT Services separate from Microsoft's consumer services.

Customer Data will not be disclosed unless required by law or allowed by this agreement. Purchasing Entity's contact information may be provided so that a requestor can contact it. If law requires disclosure, Contractor and Microsoft will use commercially reasonable efforts to notify Purchasing Entity, if permitted.

Customer Data may be transferred to, and stored and processed in, any country Contractor or Microsoft maintain facilities, except when a tenant for DPT Services is provisioned in the United States, in which case Customer Content shall be processed and stored at rest in the Continental United States.

Security Incident Notification for DPT Services. Solely for DPT Services, the following terms and conditions apply:

If Microsoft becomes aware of any Security Incident (as defined herein), Microsoft will promptly (1) notify Purchasing Entity of the Security Incident; (2) investigate the Security Incident and provide Purchasing Entity with detailed information about the Security Incident; and (3) take reasonable steps to mitigate the effects and to minimize any damage resulting from the Security Incident.

Notification(s) of Security Incidents will be delivered to one or more of Purchasing Entity's administrators by any means Microsoft selects, including via email. It is Purchasing Entity's sole responsibility to ensure Purchasing Entity's administrators maintain accurate contact information on each applicable Online Services portal. Microsoft's obligation to report or respond to a Security Incident under this section is not an acknowledgement by Microsoft of any fault or liability with respect to the Security Incident.

Purchasing Entity must notify Microsoft promptly about any possible misuse of its accounts or authentication credentials or any Security Incident related to an Online Service.

5-day Security Incident notification for Government Community Cloud Services. As an exception to the foregoing, notification of Security Incident solely pertaining to Government Community Cloud Services will be delivered within 5 days after Microsoft determines that a Security Incident has occurred, provided that Purchasing Entity must comply with the following requirements:

For each Online Service Tenant or Azure subscription, as applicable, as a condition of receiving notifications within 5 days, as set forth in the preceding paragraph, Purchasing Entity must register the following information by sending email to ols-notifications@microsoft.com, and must keep such information current at all times:

- 1) Purchasing Entity's Microsoft Online Direct Routing Domain (MODRD);
- 2) For one or more individual(s) to be contacted, each of whom must be registered as an administrator on the applicable Online Services, each of the following:
 - a. Name;
 - b. Title;
 - c. Email address registered as an administrator on the Online Services;
 - d. Email address not registered as a user on the Online Services;
- 3) Name of Purchasing Entity;
- 4) Enrollment number (provided by Contractor and assigned by Microsoft) to represent Purchasing Entity's Tenant or Azure Subscription, on Contractor's subcontract with Microsoft.

Microsoft expects to change the above process by which the Purchasing Entity for each tenant or subscription will be able to register their MODRD and other information for five-day Security Incident notification pursuant to these terms and conditions. In the event that a Purchasing Entity is notified by Microsoft, in the administrative console or otherwise, of revised instructions necessary to ensure five-day Security Incident notification, the Purchasing Entity must comply with such revised instructions.

Reimbursement of Security Incident Remediation Costs. To the extent that a Security Incident for a DPT Service results from Microsoft's failure to comply with its obligations under the Master Agreement (and, where applicable, Participating Addenda), and subject to the Limitation of Liability section in this Agreement, Microsoft will reimburse Purchasing Entities for reasonable out-of-pocket remediation costs incurred by such Purchasing Entities in connection with that Security Incident. "Reasonable out-of-pocket remediation costs" are costs that (a) are customary, reasonable and expected to be paid by entities similar to Purchasing Entity, based on the nature and scope of the Security Incident, and (b) do not arise from or relate to Purchasing Entity's violation of (i) laws applicable to Purchasing Entity or (ii) Purchasing Entity's obligations to third parties, and (c) in no event include costs arising related to compliance with laws applicable to Purchasing Entity or its industry or government function that are not generally applicable to information technology services providers. Purchasing Entity must document all such expenditures and, upon Microsoft's request, those expenditures must be validated by an independent, internationally-recognized

third party industry expert chosen by both parties. For avoidance of doubt, the costs reimbursed by Microsoft under this paragraph will be characterized as direct damages subject to the limitation on liability set forth in this Section, and not as special damages excluded under the "EXCLUSION OF CERTAIN DAMAGES" in the Limitation of Liability section.

Changes. In-Scope Services may be changed periodically. As such, Purchasing Entity's use of them will be subject to a revised OST, then-current as of the date it renews its subscription. For In-Scope Services that include client software, Purchasing Entity may be required to run a client software upgrade on its devices after a change to maintain full functionality, in accordance with the terms and conditions of the OST.

Use rights. Use rights specific to each In-Scope Service are posted online at the link to the OST, which is hereby incorporated by reference into this agreement

2. Confidentiality and Security.

For DPT Services, Contractor and Microsoft will (a) implement and maintain all appropriate administrative, physical, technical, procedural safeguards and organizational measures, internal controls, and data security routines intended to protect Customer Data against accidental loss or change, unauthorized disclosure or access, unlawful destruction, hacks, introduction of viruses, disabling devices, malware, and other forms of malicious or inadvertent acts that can disrupt Purchasing Entity's access to its Customer Data; and (b) not disclose Customer Data, except as required by law or expressly allowed. Neither party will make any public statement about this agreement's terms without the other's prior written consent.

3. Term, Termination, and Suspension.

Term and termination. This agreement will remain in effect for three years subject to each party's right under the Master Agreement and applicable law to terminate for convenience. It may be extended, based upon the mutual consent of the parties, during the term of the Master Agreement.

Customer Data. Purchasing Entity may extract Customer Data at any time. If Purchasing Entity's Subscription expires or terminates, Contractor and Microsoft will keep your Customer Data in a limited account for at least 90 days so Purchasing Entity may extract it. Contractor and Microsoft may delete Purchasing Entity's Customer Data after that, and specifically with regard to DPT Services will delete Purchasing Entity's Customer Data after no more than 180 days.

Regulatory. If a government rule or regulation applies to Contractor or Microsoft, but not generally to other businesses, and makes it difficult to operate In-Scope Services without change, or Contractor or Microsoft believe this agreement or an In-Scope Service may conflict with the rule or regulation, Microsoft may change the applicable In-Scope Service(s) or Contractor may terminate the agreement. If Microsoft changes In-Scope Services to come into compliance, and Purchasing Entity does not like the change, Purchasing Entity may terminate.

Suspension. Microsoft may suspend use of In-Scope Services: (1) if reasonably needed to prevent unauthorized Customer Data access; (2) if Purchasing Entity does not promptly respond under §5 to intellectual property claims; (3) for non-payment; or (4) if Purchasing Entity violates the AUP. A suspension will be in effect only while the condition or need exists and, if under clause (1) or (2), will apply to the minimum extent necessary. Contractor or Microsoft will notify Purchasing Entity before suspension, unless doing so may increase damages. Contractor will notify Purchasing Entity at least 30 days before suspending for non-payment. If Purchasing Entity does not fully address the reasons for suspension within 60 days after suspension, Contractor may terminate Purchasing Entity's Subscription.

4. Limited warranty; disclaimer.

Contractor warrants that In-Scope Services will meet the SLA terms during the Subscription; Purchasing Entity's only remedy for breach of warranty is stated in the SLA. *Contractor provides no (and disclaims to*

the extent permitted by law any) other warranties, express, implied, or statutory, including warranties of merchantability or fitness for a particular purpose.

5. Indemnification.

a. Subject to the exceptions below and the NASPO Participants' (as defined below) compliance with the notice and defense provisions below, in the event of any defect or deficiency in any Microsoft Products or Services purchased by a Participating Entity or Purchasing Entity, Microsoft agrees to defend NASPO, NASPO ValuePoint, the Lead State, Participating Entities, and Purchasing Entities, along with their officers, agents, and employees (collectively, the "NASPO Participants") against third party claims, damages or causes of action including reasonable attorneys' fees and related costs for any death, injury, or damage to tangible property suffered by such third party and caused by the negligence, or willful misconduct of Microsoft, its employees or subcontractors or volunteers, at any tier, during the performance of this Master Agreement (a "PI Claim"). This clause shall not be construed to bar any legal remedies Microsoft may have with respect to the NASPO Participants' failure to fulfill their obligations pursuant to the Master Agreement or any Participating Addendum.

To qualify for such defense, the Participating Entities/Purchasing Entities shall promptly notify Microsoft of any PI Claim of which the Participating Entities/Purchasing Entities become aware which may give rise to a right of defense pursuant to this Section. Notice of any PI Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Participating Entities'/Purchasing Entities' first learning of such proceeding. If the Participating Entity's/Purchasing Entity's laws require approval of a third party to defend Participating Entity/Purchasing Entity, Participating Entity/Purchasing Entity will seek such approval and if approval is not received, Microsoft is not required to defend that Participating Entity/Purchasing Entity. If a PI Claim is settled, to the extent permitted by law, the Participating Entities/Purchasing Entities shall not publicize the settlement and will cooperate with Microsoft so that Microsoft can make every effort to ensure the settlement agreement contains a non-disclosure provision.

Notwithstanding anything to the contrary contained herein, Participating Entities/Purchasing Entities agree that Microsoft has no obligation for any PI Claim covered by this Section arising out of or resulting from the Participating Entities'/Purchasing Entities' or any of their respective employees', contractors' or agents' acts of negligence, gross negligence or misconduct. THE FOREGOING SHALL CONSTITUTE EACH AND EVERY PARTICIPATING ENTITY'S/PURCHASING ENTITY'S SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL PI CLAIMS.

b. **Indemnification** – Intellectual Property. Subject to the NASPO Participant's compliance with the notice and defense requirements and exceptions set forth below, Microsoft agrees to defend the NASPO Participants against any claims made by an unaffiliated third party that any Product or Service or its proper or reasonably expected or acceptable use infringes that third party's patent, copyright, or trademark or makes unlawful use of its trade secret (an "Intellectual Property Claim").

(1) Microsoft's obligations under this section shall not extend to any claims based on:

- (a) Microsoft's compliance with a Participating Entity's/Purchasing Entity's designs, specifications or instructions; or
- (b) Microsoft's use of technical information or technology provided by the Participating Entity/Purchasing Entity; or
- (c) Non-Microsoft software, modifications a Participating Entity/Purchasing Entity makes to, or any specifications or materials a Participating Entity/Purchasing Entity provides or makes available for, a Product; or
- (d) Participating Entity's/Purchasing Entity's combination of the Product or Service with a Non-Microsoft product, data or business process; or damages based on the use of a Non-Microsoft product, data or business process; or

(e) Participating Entity's/Purchasing Entity's use of either Microsoft's trademarks or the use or redistribution of a Product or Service in violation of this Master Agreement, Participating Addendum, or any other agreement incorporating its terms; or

(f) Participating Entity's/Purchasing Entity's use of a Product or Service after Microsoft notifies Participating Entity/Purchasing Entity to discontinue that use due to a third party claim.

(2) To qualify for such defense, the involved NASPO Participants (the "Indemnified Party") shall promptly notify Microsoft of any Intellectual Property Claim of which the Indemnified Party become aware which may give rise to right of defense pursuant to this Section. Notice of any Intellectual Property Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Indemnified Party's learning of such proceeding. If the Indemnified Party's laws require approval of a third party to defend the Indemnified Party, the Indemnified Party will seek such approval and if approval is not received, Microsoft is not required to defend that Indemnified Party. In the event the Indemnified Party does not authorize sole control to Microsoft over any claims that may arise under this subsection, then the parties agree that Microsoft will be granted authorization to equally participate in any proceeding subject to this subsection. The Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Microsoft's reasonable request and expense, information and assistance necessary for such defense. If Microsoft fails to vigorously pursue the defense or settlement of the Intellectual Property Claim, the Indemnified Party may assume the defense or settlement of it and Microsoft shall be liable for all costs and expenses, including reasonable attorneys' fees and related costs, incurred by the Indemnified Party in the pursuit of the Intellectual Property Claim. Unless otherwise agreed in writing, this section is not subject to any limitations of liability in this Master Agreement or in any other document executed in conjunction with this Master Agreement.

If Microsoft reasonably believes that a Product or Service may infringe or misappropriate a third-party's intellectual property rights, Microsoft will seek to:

- i. procure for Participating Entities/Purchasing Entities the right to continue to use the Product or Service; or
- ii. modify or replace it with a functional equivalent to make it non-infringing and notify Participating Entities/Purchasing Entities to discontinue use of the prior version, which Participating Entities/Purchasing Entities must do immediately.

If the foregoing options are not commercially reasonable for Microsoft, or if required by a valid judicial or government order, Microsoft may cause Contractor to terminate Participating Entities'/Purchasing Entities' license or access rights in the Product or Service. In such a case, Microsoft will provide Participating Entities'/Purchasing Entities with notice and refund any amounts Participating Entities'/Purchasing Entities have paid for those rights to the Product or Service.

THE FOREGOING SHALL CONSTITUTE THE LEAD STATE'S AND EACH AND EVERY PARTICIPATING AND/OR PURCHASING ENTITIES' SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL INTELLECTUAL PROPERTY CLAIMS.

6. *Limitation of Liability.*

The parties agree that, other than In-Scope Services, no other Microsoft Products shall be purchased hereunder. As applicable to In-Scope Services, the following terms and conditions shall apply:

a. General. The total liability of each party for In-Scope Services hereunder (Microsoft, Contractor and each Purchasing Entity, including their Affiliates and contractors), for claims arising under this Agreement, is limited to direct damages up to the amount Purchasing Entity paid for the In-Scope Service during the prior 12 months before the cause of action arose; but in no event will a party's aggregate liability for any In-Scope Service exceed the total amount paid for that In-Scope Service under this Agreement. In the case of In-Scope Services provided free of charge, previews, or code that a Purchasing Entity is authorized to redistribute to third parties without separate payment to Contractor, Microsoft's liability is limited to direct damages up to U.S. \$5,000. These limitations apply regardless of whether the liability is based on breach of contract, tort (including negligence), strict liability, breach of warranties, or any other legal theory.

b. Affiliates and contractors. Contractor and Purchasing Entity each agree not to bring any action against the other's Affiliates or contractors in respect of any matter disclaimed on their behalf in this Agreement. Each party will be responsible for its actions in the event of any breach of this provision.

c. EXCLUSION OF CERTAIN DAMAGES. Neither party nor their Affiliates or contractors will be liable for any indirect, consequential, special or incidental damages, or damages for lost profits, revenues, business interruption, or loss of business information in connection with this agreement, even if advised of the possibility of such damages or if such possibility was reasonably foreseeable.

d. Limits. *The limits and exclusions in this section titled "Limitation of liability" do not apply to either party's (1) obligations under the section titled "Defense of third party claims", (2) liability for damages caused by either party's gross negligence or willful misconduct, or that of its employees or its agents, and awarded by a court of final adjudication (provided that, in jurisdictions that do not recognize a legal distinction between "gross negligence" and "negligence," "gross negligence" as used in this subsection shall mean "recklessness"); and (3) liability for violation of its confidentiality obligations (except obligations related to Customer Data) or the other party's intellectual property rights.*

7. Agreement mechanics

Purchasing Entity must send notice by regular mail, return receipt requested, to the address on the Portal (effective when delivered). Contractor or Microsoft may email notice to your account administrators (effective when sent). Purchasing Entity may not assign this agreement, or any right or duty under it. If part of this agreement is held unenforceable, the rest remains in force. Failure to enforce this agreement is not a waiver. The parties are independent contractors. This agreement does not create an agency, partnership, or joint venture. This agreement is governed by the laws applicable to Purchasing Entity, without regard to conflict of laws. This agreement (including the SLA and OST) and Contractor's price sheet are the parties' entire agreement on this subject and supersedes any concurrent or prior communications. Agreement terms that require performance, or apply to events that may occur, after termination or expiration will survive, including §5. In-Scope Service and any Client associated therewith are subject to U.S. export jurisdiction. Purchasing Entity must comply with the U.S. Export Administration Regulations, the International Traffic in Arms Regulations, and end-user, end-use, and destination restrictions. For more information, see <http://www.microsoft.com/exporting/>. Microsoft delivers In-Scope Services, and the rights granted to us also apply to them.

8. Miscellaneous Microsoft Terms and Conditions

A. Accessibility. Microsoft supports the government's obligation to provide accessible technologies to its citizens with disabilities as required by Section 508 of the Rehabilitation Act of 1973, and its state law counterparts (including applicable California provisions). Microsoft encourages its customers (including Purchasing Entities under the Master Agreement and Participating Addenda) to judiciously compare product accessibility performance. The Voluntary Product Accessibility Templates ("VPATs") for the Microsoft technologies used in providing the online services can be found at Microsoft's VPAT page. Further information regarding Microsoft's commitment to accessibility can be found at www.microsoft.com/enable.

B. DPT Services Information Security Policy. For each DPT Service, Microsoft follows a written data security policy ("Information Security Policy") that complies with the control standards and frameworks shown in the table below.

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Office 365 Services	Yes	Yes	Yes	Yes*	Yes*
Microsoft Dynamics CRM Online Services	Yes	Yes	Yes	Yes	Yes

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Microsoft Azure Core Services	Yes	Yes	Yes	Varies**	Varies**
Microsoft Intune Online Services	Yes	Yes	Yes	Yes	Yes

*Does not include Yammer Enterprise.

**Current scope is detailed in the audit report and summarized in the Microsoft Azure Trust Center.

Microsoft may add industry or government standards at any time. Microsoft will not eliminate a standard or framework in the table above, unless it is no longer used in the industry and it is replaced with a successor (if any). Azure Government Services meet a separate set of control standards and frameworks, as detailed on the Microsoft Azure Trust Center.

Subject to non-disclosure obligations, Microsoft will make each Information Security Policy available to Purchasing Entity, along with other information reasonably requested by Purchasing Entity regarding Microsoft security practices and policies.

Purchasing Entity is solely responsible for reviewing each Information Security Policy and making an independent determination as to whether it meets Purchasing Entity's requirements.

C. **Microsoft Audits of DPT Services.** For each DPT Service, Microsoft will conduct audits of the security of the computers, computing environment and physical data centers that it uses in processing Customer Data (including personal data), as follows:

- Where a standard or framework provides for audits, an audit of such control standard or framework will be initiated at least annually for each DPT Service.
- Each audit will be performed according to the standards and rules of the regulatory or accreditation body for each applicable control standard or framework.
- Each audit will be performed by qualified, independent, third party security auditors at Microsoft's selection and expense.

Each audit will result in the generation of an audit report ("Microsoft Audit Report"), which will be Microsoft's Confidential Information. The Microsoft Audit Report will clearly disclose any material findings by the auditor. Microsoft will promptly remediate issues raised in any Microsoft Audit Report to the satisfaction of the auditor.

If Purchasing Entity requests, Microsoft will provide Purchasing Entity with each Microsoft Audit Report so that Purchasing Entity can verify Microsoft's compliance with the security obligations under the DPT. The Microsoft Audit Report will be subject to non-disclosure and distribution limitations of Microsoft and the auditor.

D. **HIPAA/HITECH Business Associate.** If Purchasing Entity is a "covered entity" or a "business associate" and includes "protected health information" in Customer Data, as those terms are defined in 45 CFR § 160.103, this agreement includes execution of the HIPAA Business Associate Agreement ("BAA"), the full text of which identifies the DPT Services to which it applies and is available at <http://aka.ms/BAA>. Purchasing Entity may opt out of the BAA by instructing Contractor to send the following information to Microsoft in a written notice (under the terms of the Contractor's subcontract with Microsoft):

- the full legal name of Purchasing Entity and any Affiliate that is opting out;
- the volume licensing agreement number to which the opt out applies.

E. **FedRAMP.** During the term of a Purchasing Entity's subscription for Government Community Cloud Services, those services will be operated in accordance with a written data security policy and control framework that is consistent with the requirements of NIST 800-53 Revision 4, or successor standards and guidelines (if any), established to support Federal Risk and Authorization Management Program (FedRAMP) accreditation at a Moderate Impact level. Microsoft intends for Government Community Cloud Services to support FedRAMP Authority to Operate ("ATO"), and Microsoft will use commercially reasonable efforts to obtain an ATO from a Federal agency, and to

maintain such ATO through continuous monitoring processes and by conducting regular FedRAMP audits.

- F. **Background Checks.** Microsoft performs the following background checks on all US personnel who have potential to access Customer Data. Adherence to this policy is one of the control procedures addressed by the Microsoft Audit Report per the section of the Microsoft Online Services Terms titled "Microsoft Audits of Online Services." Such Background Checks will be performed in accordance with the Fair Credit Reporting Act and will consist of Social Security Number trace, seven (7) year felony and misdemeanor criminal records check of federal, state, or local records (as applicable) for job related crimes, Office of Foreign Assets Control List (OFAC) check, Bureau of Industry and Security List (BIS) check and Office of Defense Trade Controls Debarred Persons List (DDTC) check.

Additionally, for States with which Microsoft has entered into a Microsoft CJIS Information Agreement and provided an FBI CJIS Addendum Certification to the applicable CJIS Systems Agency (CSA, e.g. the California Department of Justice), for its CJIS Covered Services (which are most of the Microsoft Government Community Cloud Services), Microsoft submits its personnel that have access to unencrypted Criminal Justice Information and other Customer Content (including both employees and subcontractors, where applicable) for FBI NCIC fingerprint background checks and adjudication by each such CSA.

- G. **Click-Through Terms.** For the Microsoft In-Scope Services sold and licensed to Purchasing Entities hereunder, no click-through licensing terms presented to end users or administrators, as they pertain to the delivery, operation or use of such services, shall be binding. For clarity, to the extent that certain services may present terms of use for a web portal used to administer and configure the In-Scope Services, or to download software in conjunction with the In-Scope Services, such terms of use shall be binding, to the extent that there are no equivalent terms and conditions in the Master Agreement which pertain to the use of such web portals.

9. **Definitions.**

"AUP" means the section of the OST titled "Acceptable Use Policy" "Azure Consumption-Based Services" means any Microsoft Online Service made available for Purchasing Entity's use and consumption under the Microsoft Azure Administrative Portal, following Purchasing Entity's initial subscription order for such services, subject to consumption fees billed to Purchasing Entity (or debited against its prepaid Azure monetary commitment balance, if applicable) as described in the applicable Product Terms.

"Client" means device software, if any, that Contractor or Microsoft provides to Purchasing Entity with In-Scope Services.

"Customer Content" means the subset of Customer Data created by users, which includes the following:

- For Office 365 Services, Customer Content shall at least include Exchange Online mailbox content (e-mail body, calendar entries, and the content of e-mail attachments), SharePoint Online site content and the files stored within that site, and Skype for Business Online archived conversations.
- For Microsoft Dynamics CRM Online Services, Customer Content means the entities of Customer Data managed by the Microsoft Dynamics CRM Online Services.
- For Microsoft Azure Core Services, Customer Content includes all Customer Data except as may be noted in the Microsoft Azure Trust Center (which Microsoft may update from time to time, but will not add exceptions for existing Microsoft Azure Core Services in general release).
- For Microsoft Intune Online Services, Customer Content includes all Customer Data except as may be noted in the "Data Location" section of the Microsoft Intune Trust Center.

"Customer Data" means all data, including all text, sound, or image files that are provided to Microsoft by, or on behalf of, Purchasing Entity through its use of In-Scope Services.

"DPT" means the section of the OST titled "Data Processing Terms."

“DPT Services” means the Microsoft Online Services to which the DPT applies. As of the date this agreement was executed by the parties, DPT Services include “Microsoft Dynamics CRM Online Services,” “Office 365 Services,” “Microsoft Intune Online Services,” and “Microsoft Azure Core Services,” which consist of the following component Microsoft Online Services:

DPT Services	
Microsoft Dynamics CRM Online Services	Microsoft Dynamics CRM Online services made available through volume licensing or the Microsoft online services portal, excluding (1) Microsoft Dynamics CRM for supported devices, which includes but is not limited to Microsoft Dynamics CRM Online services for tablets and/or smartphones and (2) any separately-branded service made available with or connected to Microsoft Dynamics CRM Online, such as Microsoft Social Engagement, Parature, from Microsoft, and Microsoft Dynamics Marketing.
Office 365 Services	The following services, each as a standalone service or as included in an Office 365-branded plan or suite: Exchange Online, Exchange Online Archiving, Exchange Online Protection, Advanced Threat Protection, SharePoint Online, OneDrive for Business, Project Online, Skype for Business Online, Sway, Office Online, Delve Analytics, Customer Lockbox, and Yammer Enterprise. Office 365 Services do not include Office 365 ProPlus, any portion of PSTN Services that operate outside of Microsoft’s control, any client software, or any separately branded service made available with an Office 365-branded plan or suite, such as a Bing or a service branded “for Office 365.”
Microsoft Azure Core Services	Cloud Services (web and worker roles), Virtual Machines (including with SQL Server), Storage (Blobs, Tables, Queues), Virtual Network, Traffic Manager, Batch, Web Sites, BizTalk Services, Media Services, Mobile Services, Service Bus, Notification Hub, Workflow Manager, Express Route, Scheduler, Multi-Factor Authentication, Active Directory, Rights Management Service, SQL Database, and HDInsight.
Microsoft Intune Online Services	The cloud service portion of Microsoft Intune such as the Microsoft Intune Add-on Product or a management service provided by Microsoft Intune such as Mobile Device Management for Office 365.

“Government Community Cloud Services” means the Microsoft Online Services identified as such (a subset of DPT Services) in the DPT and Product Terms.

“In-Scope Services” means the Microsoft Online Services for which subscription licenses are sold to Purchasing Entity by Contractor. In-Scope Services include, but are not limited to DPT Services. Except for Azure Consumption-Based Services (which are billed based upon the quantities of such services consumed by Purchasing Entity), Purchasing Entity will order individual licenses from Contractor for all In-Scope Services Licenses.

“Master Agreement” means the agreement (described therein as “NASPO ValuePoint Master Agreement for Cloud Services”) between Contractor and the Lead State, subject to the terms and conditions of the Participating Addendum between Contractor and Purchasing Entity’s State (if different from Utah), and which resulted from award of contract under Utah Solicitation Number CH16012.

“Microsoft” means Microsoft Corporation, which is subcontractor to Contractor and delivers the In-Scope Services.

“Online Services” means the Microsoft-hosted services for which Purchasing Entity acquires Licenses from Contractor hereunder.

“OST” means the Microsoft “Online Services Terms,” which are additional terms that apply to Purchasing Entity’s use of Online Services published on the Volume Licensing Site and updated from time to time, and is subject to terms and conditions in the OST which govern which version of OST is applicable to a particular subscription order.

“Portal” means the Online Services Portal for each In-Scope Service.

“Product Terms” means the document that provides information about Microsoft Products and Professional Services available through volume licensing. The Product Terms document is available on the Volume Licensing Site and is updated from time to time.

“Security Incident” means any unlawful access, use, theft or destruction to any Customer Data stored on Microsoft’s equipment or in Microsoft’s facilities, or unauthorized access to such equipment or facilities resulting in use, theft, loss, disclosure, alteration or destruction of Customer Data. All references to “Data Breach” in the Master Agreement shall be deemed to mean Security Incident.

“Service Level Agreement” (“SLA”) means the document which specifies the standards to which Microsoft agrees to adhere and by which it measures the level of service for an In-Scope Service. Microsoft’s current and archived prior version SLAs are available at the Volume Licensing Site. With respect to any In-Scope Service ordered under the this Agreement, the most current SLA available at the onset of a subscription License term shall apply to that In-Scope Service for the duration of that subscription License Term, after which (for any subsequent renewal subscription License term) it will be superseded by the version current at the time of renewal.

“Subscription” means an order for a quantity of an In-Scope Service, including but not limited to (a) orders for User Subscription Licenses for one or more In-Scope Services; and (b) an order for the purpose of establishing a subscription for Azure Consumption-Based Services.

“Volume Licensing Site” means <http://www.microsoft.com/licensing/contracts> or a successor site.

APPENDIX 13-GOOGLE-1

Exhibit A13.1 lists Google's cloud services.

Exhibit A13.1. Google Cloud Services.

Category	Product Name	Product Description
SaaS	Google Apps Auth User	<i>Google Apps Auth User includes the ability to define Google Apps user identities for managing Android applications and devices only.</i>
SaaS	Google Apps Message Encryption	<i>With Google Apps for Work, Edu, or Unlimited, customers can add Google Apps Message Encryption to set policies to route their outbound email to an encryption server for external recipients to securely receive and send email to their end users.</i>
SaaS	Google Apps No Gmail	<i>Google Apps No Gmail includes Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps for Work	<i>Google Apps for Work includes mail (30 GB of storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps for EDU, Unlimited	<i>For qualifying K-12 and higher education organizations, Google Apps for EDU, Unlimited includes mail, Google Vault eDiscovery and Archiving (unlimited storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps Deskless	<i>Google Apps Deskless includes the same services above; this offering is intended for nontechnical workers and infrequent users.</i>
SaaS	Google Apps Unlimited	<i>Google Apps Unlimited includes mail, Google Vault eDiscovery and Archiving (unlimited storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps Unlimited Deskless	<i>Google Apps Unlimited customers can blend this product into their contract. This offering, which includes the same services above, is intended for nontechnical workers and infrequent users.</i>
SaaS	Google Apps Vault	<i>Google Apps Vault provides Vault eDiscovery and Archiving (unlimited storage).</i>
SaaS	Google Apps Vault Deskless	<i>Google Apps customers can blend this product into their contract. This offering, which includes the same services above, is intended for nontechnical workers and infrequent users.</i>

APPENDIX-UNISYS-1

Exhibits X and Y list the Unisys solutions and services provided by Unisys.

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud.

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
Unisys Justice Law Enforcement and Border Security Solutions for the Cloud		
Secure Image Management Solution	Unisys developed a highly secure solution to meet the forensic integrity standards and policy requirements of several of the largest police departments in the world. The secure storage of images and other multimedia content occurs in a tamperproof information archive – in effect, a digital evidence vault.	All
U-LEAF	U-LEAF provides a POLE-type data model (Person, Object, Location and Event) for the storage and recording of incidents and entities. The POLE model allows entities to be recorded in the system once. The recorded entities, however, can be linked to other entities and events as many times as necessary, to build the picture of an incident, or a network of associations.	All
Unisys Social Services Solutions for the Cloud		
Unisys 311	The Unisys “311 solution” is a Multi- Channel Citizen Service Delivery and Engagement solution that can be used by a Local Government or any Governmental Agency to deliver and manage non-emergency services, respond to inquiries and engage with their constituents. This offers multi-channel interaction including social media and mobile applications, supported by a knowledgebase and GIS systems and also interfaces with existing systems to manage work orders and service requests. The solution can be configured to meet specific needs and extend its functionality and flexibility through other apps, and integration to mobile apps, social media and other systems.	All
State/Local Government Enterprise Regulatory System (AMANDA)	The Unisys AMANDA Platform (from CSDC) is designed to provide a collection of back office functions such as Licensing, Permitting, Inspections, Land Use, Planning and a number of other functions in various form factors (desktop, mobile and tablet) and can also be deployed either on premise, hosted, or delivered via the Cloud. The solution includes the design, delivery, deployment and ongoing support and maintenance of the solution. The AMANDA platform can be configured specific to client needs in order to extend its functionality and interface with other systems as required.	All
Unisys Enterprise Content Management for the Cloud		
Infomage	Unisys Infomage is the Enterprise Content Management (ECM), Business Process Management (BPM) and Record Management (RM) solution for organizations looking to significantly improve business processes that depend heavily high-volume paper documents, documents generated from Internet transactions, Office documents, and other electronic documents that need to be accessed for automated and manual processes. Infomage can easily capture, manage, store, and access the content required for cases, inquiries, and process-centric work, regardless of data structure or document origination, with a single	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	intuitive user interface. Unisys Infolmage brings together ECM, imaging, workflow, document management, record management and integration technologies to form an integrated end-to-end solution.	
Unisys Horizontal Solutions for the Cloud		
Stealth	The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. We built a better way to deal with advanced threats for our clients by applying novel approaches to new threats. By substituting traditional hardware topology for software-based cryptography, our Stealth Microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users.	Private and Hybrid Cloud Only. For Public and Community Cloud, see AWS offering above. IaaS, PaaS and SaaS
VantagePoint	VantagePoint can also extend beyond the boundaries of IT with uses cases to support relevant business scenarios in, • Security: Strengthens the security posture of organizations through greater visibility into the Stealth enabled infrastructure. • Advanced Data Analytics: Cuts across data silos and sources relevant data to optimize the performance of Advanced Analytics solution. • Cloud and Infrastructure Services: Enables hybrid IT by seamlessly integrating with a variety of data sources and services and presenting them via a digital control panel, irrespective of technology underpinnings. • End User Services: Supports millennial users via personalized and secure access to relevant data and services, across endpoints. • Service Management: Cuts across ITSM platforms to improve service delivery through service automation, orchestration, and aggregation capabilities. • Facilities/Crisis Management: Helps organizations monitor their facilities and keep their Facilities Managers abreast of any security threats and risks. Also acts as a standard communication vehicle for users at times of crisis. (178) With VantagePoint, our clients gain a personalized, secure, and intuitive data and service aggregation platform that cuts across strategic and operational dimensions of business and accelerates digital transformation.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys Cloud Hosting	Unisys Hosted Private Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can provide, integrate and scale the cloud infrastructure to meet client's needs. Our management platform enables a single pane of glass to effectively manage client's workloads. Unisys Hybrid Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can source, integrate and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our management platform enables a single pane of glass to effectively	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	manage a variety of clouds within an organization, whether they be public or private.	
ServiceNow	ServiceNow offers enterprise service management software for human resources, law, facilities management, finance, marketing, and field operations in the cloud. ServiceNow has its STAR Self-Assessment available on the CSA's website. ServiceNow specializes in ITSM applications and provides forms-based workflow application development. ServiceNow has open integration options to variety platforms such as: Salesforce, SharePoint, and BMC Remedy Action Request System.	All
Unisys ClearPath Forward!	Unisys ClearPath Forward is an Intel based fabric computing platform designed to run mission critical applications that require predictable performance and low transport latency. The ClearPath Forward platform provides hardware partitioning technology similar to the Mainframes, but is designed to run Windows and Linux operating environments on commodity server hardware. The fabric interconnect included with the ClearPath Forward platform provides transport speeds at 56Gbps or higher. The ClearPath Forward platform fabric can be used in a hybrid mode with VMware.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys White Label Offerings	Unisys has strategic, long lasting relationships in the industry and is pleased to offer our partner products surrounding the cloud market space. This is an evolving list. Our current partner list includes: SalesForce, Oracle, Verizon, SHI, Decision Lens, Aptio, NetApp, VMware, EMC, NetApp, Birst, Okta, Box, WatchDox. We have created webpage on our public site that is updated with our list of current partners. The page is available here: http://www.unisys.com/ms/wsca-cloud-hosting/	All

Exhibit Y. Unisys Services Offerings.

Unisys Professional Services	Description
Cloud Advisory Services	Unisys Cloud Advisory Services provide strategic and financial guidance on aligning IT with business objectives. This starts with a roadmap that outlines the vision of a hybrid IT based on a combination of existing data center, internal cloud, and external cloud resources to provide agility, flexibility, and control.
Data Center Planning, Design, and Implementation Services	Unisys Data Center Planning, Design, and Implementation Services offer a complete range of services that delivers cohesive, end-to-end optimization of data centers. With a wide range of services for discovery, analysis, optimization, virtualization, consolidation, and migration of data centers that can complement client efforts and fill gaps in skills and capacities and a combination of world-class people, processes, and technology with program and project management expertise, Unisys transforms clients' existing data centers to a business engine that provides agility at a lower cost.
CloudBuild Services	Unisys CloudBuild Services enable organizations to successfully build a cloud that is integrated with the overall business process, transforming their existing infrastructure to an agile IT-as-a-service model. The Unisys "8-Tracks" model, a 360-degree cloud view approach, covers eight critical data center domains. Together with ConOps, which includes industry best practices, Unisys enables a client's cloud infrastructure to meet the security, regulatory, and compliance requirements that enable us to deliver the most secure and reliable cloud in the marketplace.
Hybrid Cloud Strategy	Unisys Hybrid Cloud strategy helps organizations overcome key challenges when planning to implement a cloud environment. Unisys provides governments and businesses with a comprehensive cloud architecture that gives clients cloud on their

Unisys Professional Services	Description
	terms. We can source, integrate, and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our homegrown VantagePoint management platform enables a single pane of glass to effectively manage a variety of clouds in an organization, whether they be public or private.
Unisys Platform Services	
Platform Assessment Services	Unisys Platform Assessment Services enable governments and businesses to foster their platform adoption initiatives. Unisys analyzes organizations' business goals and objectives along with their technology landscape of their enterprise application portfolio and recommend a best-fit Platform Suitability Analysis to help them make the right decision in their move to PaaS.
Architecture Design Services	Unisys Architecture Design Services assist governments and businesses with designing an effective architecture suited for cloud. We provide several concepts and best practices that are essential to build highly scalable applications in the cloud – be it on-premises, public, or hybrid.
PaaS Enablement Services	Unisys PaaS Enablement Services analyze an organization's current set of applications and provide a strategy to use the best practices and engagement models of various platform providers such as Microsoft Azure, IBM Bluemix, SAP HCP, and SFDC AppCloud.
Application Migration Services	Unisys Application Migration Services assist organizations in their cloud migration activities by using tools and frameworks adhering to Unisys best practices that promise minimal downtime without affecting the day-to-day business processes.
Application Portability Services	Unisys Application Portability Services help organizations to develop cross-platform applications that can be scaled across multiple cloud platforms.
Application Development Services	Unisys Application Development Services help organizations to build rich, interactive applications focused on business logic and workflows using visual tools as well as cloud-based tools, architectures, and services that make their applications cloud ready.
IoT Development Services	Unisys IoT Development Services enable organizations to use IoT Services across various cloud platforms that enable apps to communicate and consume data collected by the connected devices, sensors, and gateways.
API Management Services	Unisys API Management Services provide a solution that addresses the aspects of the application programming interface (API) life cycle for on-premises and cloud environments and offer capabilities to create, run, manage, secure, and monetize APIs and microservices that deliver an integrated user experience and enable rapid deployment and simplified administration of APIs.
Horizontal Technology Integration Services	
Horizontal Technology Integration Services	Unisys is a world-recognized leader in integrating technology infrastructures across platforms and vendors. We particularly specialize in integrating solutions surrounding service desks, cloud provisioning, data center management, and the ITIL framework.
Unisys Service Management Services	
Maturity and Platform Technology Assessments	Unisys Service Management Consulting Services leverage subject matter experts with more than 15 years of experience in working and delivering services for ITIL and service management. We leverage that knowledge to provide best practices, lessons learned, and roadmap development for Service Management disciplines and benchmark them with industry and operational best practices.
Service Management Platform Implementation	Unisys brings experience in moving existing ITSM tool information, processes, and requirements to new cloud enterprise Service Management platforms. By leveraging industry best practices and structured methodologies, Unisys can establish transition and transformation plans that minimize risk to clients' ongoing operations and deliver speed to value with the new cloud-based Service Management platform.
Enterprise Cloud Service Management Platform Support	Unisys can provide ongoing day-to-day support in the ongoing management of a client's Service Management platform. This service allows the organization to focus on the business aspects instead of the routine activities necessary to support and maintain a platform. Unisys provides cost-effective solutions that leverage staff in many locations, driving the cost of servicing the environment to effective and efficient levels.

Appendix A

Additional Use Rights and Restrictions

These Additional Use Rights and Restrictions shall apply to each Enterprise's use of Products and must be incorporated in each Government Contract.

1. General.

Right to use. Purchasing Entity may access and use In-Scope Services, and install and use a Client (if any) included with its Subscription, only as described in this agreement. All other rights are reserved.

Acceptable use. Purchasing Entity will use In-Scope Services only per the AUP. Purchasing Entity will not use In-Scope Services in any way that infringes a third party's patent, copyright, or trademark or misappropriates its trade secret. Purchasing Entity may not reverse engineer, decompile, work around technical limits in, or disassemble In-Scope Services, except if applicable law permits despite this limit. Purchasing Entity may not rent, lease, lend, resell, transfer, or host In-Scope Services to or for third parties.

Compliance.

Purchasing Entity will comply with all laws and regulations applicable to its use of In-Scope Services.

In providing In-Scope Services, Contractor and Microsoft will comply with all laws and regulations (including applicable security breach notification law) that generally apply to IT service providers. Microsoft will comply with all laws and regulations applicable to its provision of the Online Services, including security breach notification law. However, Microsoft is not responsible for compliance with any laws or regulations applicable to Customer or Customer's industry or government function that are not generally applicable to information technology service providers. Microsoft does not determine whether Customer Data includes information subject to any specific law or regulation. All Security Incidents are subject to the Security Incident Notification terms below.

Purchasing Entity will obtain any consents required: (1) to allow it to access, monitor, use, and disclose user data; and (2) for Contractor and Microsoft to provide the In-Scope Services. If Purchasing Entity is an educational institution, it will obtain any parental consent for end users' use of In-Scope Services as required by applicable law.

Customer Data. Customer Data is used only to provide Purchasing Entity In-Scope Services. This use may include troubleshooting to prevent, find and fix problems with those In-Scope Services' operation. It may also include improving features for finding and protecting against threats to users. Neither Contractor nor Microsoft will derive information from Customer Data for any advertising or other commercial purposes. As between the parties, the Purchasing Entity retains all right, title and interest in and to Customer Data. Neither Contractor nor Microsoft acquires rights in Customer Data, other than the rights Purchasing Entity grants to Microsoft through Contractor to provide the Online Services to Purchasing Entity. This paragraph does not affect Microsoft's rights in software or Online Services Contractor and Microsoft license to Purchasing Entity.

Contractor and Microsoft will enable Purchasing Entity to keep Customer Data in DPT Services separate from Microsoft's consumer services.

Customer Data will not be disclosed unless required by law or allowed by this agreement. Purchasing Entity's contact information may be provided so that a requestor can contact it. If law requires disclosure, Contractor and Microsoft will use commercially reasonable efforts to notify Purchasing Entity, if permitted.

Customer Data may be transferred to, and stored and processed in, any country Contractor or Microsoft maintain facilities, except when a tenant for DPT Services is provisioned in the United States, in which case Customer Content shall be processed and stored at rest in the Continental United States.

Security Incident Notification for DPT Services. Solely for DPT Services, the following terms and conditions apply:

If Microsoft becomes aware of any Security Incident (as defined herein), Microsoft will promptly (1) notify Purchasing Entity of the Security Incident; (2) investigate the Security Incident and provide Purchasing Entity with detailed information about the Security Incident; and (3) take reasonable steps to mitigate the effects and to minimize any damage resulting from the Security Incident.

Notification(s) of Security Incidents will be delivered to one or more of Purchasing Entity's administrators by any means Microsoft selects, including via email. It is Purchasing Entity's sole responsibility to ensure Purchasing Entity's administrators maintain accurate contact information on each applicable Online Services portal. Microsoft's obligation to report or respond to a Security Incident under this section is not an acknowledgement by Microsoft of any fault or liability with respect to the Security Incident.

Purchasing Entity must notify Microsoft promptly about any possible misuse of its accounts or authentication credentials or any Security Incident related to an Online Service.

5-day Security Incident notification for Government Community Cloud Services. As an exception to the foregoing, notification of Security Incident solely pertaining to Government Community Cloud Services will be delivered within 5 days after Microsoft determines that a Security Incident has occurred, provided that Purchasing Entity must comply with the following requirements:

For each Online Service Tenant or Azure subscription, as applicable, as a condition of receiving notifications within 5 days, as set forth in the preceding paragraph, Purchasing Entity must register the following information by sending email to ols-notifications@microsoft.com, and must keep such information current at all times:

- 1) Purchasing Entity's Microsoft Online Direct Routing Domain (MODRD);
- 2) For one or more individual(s) to be contacted, each of whom must be registered as an administrator on the applicable Online Services, each of the following:
 - a. Name;
 - b. Title;
 - c. Email address registered as an administrator on the Online Services;
 - d. Email address not registered as a user on the Online Services;
- 3) Name of Purchasing Entity;
- 4) Enrollment number (provided by Contractor and assigned by Microsoft) to represent Purchasing Entity's Tenant or Azure Subscription, on Contractor's subcontract with Microsoft.

Microsoft expects to change the above process by which the Purchasing Entity for each tenant or subscription will be able to register their MODRD and other information for five-day Security Incident notification pursuant to these terms and conditions. In the event that a Purchasing Entity is notified by Microsoft, in the administrative console or otherwise, of revised instructions necessary to ensure five-day Security Incident notification, the Purchasing Entity must comply with such revised instructions.

Reimbursement of Security Incident Remediation Costs. To the extent that a Security Incident for a DPT Service results from Microsoft's failure to comply with its obligations under the Master Agreement (and, where applicable, Participating Addenda), and subject to the Limitation of Liability section in this Agreement, Microsoft will reimburse Purchasing Entities for reasonable out-of-pocket remediation costs incurred by such Purchasing Entities in connection with that Security Incident. "Reasonable out-of-pocket remediation costs" are costs that (a) are customary, reasonable and expected to be paid by entities similar to Purchasing Entity, based on the nature and scope of the Security Incident, and (b) do not arise from or relate to Purchasing Entity's violation of (i) laws applicable to Purchasing Entity or (ii) Purchasing Entity's obligations to third parties, and (c) in no event include costs arising related to compliance with laws applicable to Purchasing Entity or its industry or government function that are not generally applicable to information technology services providers. Purchasing Entity must document all such expenditures and, upon Microsoft's request, those expenditures must be validated by an independent, internationally-recognized

third party industry expert chosen by both parties. For avoidance of doubt, the costs reimbursed by Microsoft under this paragraph will be characterized as direct damages subject to the limitation on liability set forth in this Section, and not as special damages excluded under the "EXCLUSION OF CERTAIN DAMAGES" in the Limitation of Liability section.

Changes. In-Scope Services may be changed periodically. As such, Purchasing Entity's use of them will be subject to a revised OST, then-current as of the date it renews its subscription. For In-Scope Services that include client software, Purchasing Entity may be required to run a client software upgrade on its devices after a change to maintain full functionality, in accordance with the terms and conditions of the OST.

Use rights. Use rights specific to each In-Scope Service are posted online at the link to the OST, which is hereby incorporated by reference into this agreement

2. Confidentiality and Security.

For DPT Services, Contractor and Microsoft will (a) implement and maintain all appropriate administrative, physical, technical, procedural safeguards and organizational measures, internal controls, and data security routines intended to protect Customer Data against accidental loss or change, unauthorized disclosure or access, unlawful destruction, hacks, introduction of viruses, disabling devices, malware, and other forms of malicious or inadvertent acts that can disrupt Purchasing Entity's access to its Customer Data; and (b) not disclose Customer Data, except as required by law or expressly allowed. Neither party will make any public statement about this agreement's terms without the other's prior written consent.

3. Term, Termination, and Suspension.

Term and termination. This agreement will remain in effect for three years subject to each party's right under the Master Agreement and applicable law to terminate for convenience. It may be extended, based upon the mutual consent of the parties, during the term of the Master Agreement.

Customer Data. Purchasing Entity may extract Customer Data at any time. If Purchasing Entity's Subscription expires or terminates, Contractor and Microsoft will keep your Customer Data in a limited account for at least 90 days so Purchasing Entity may extract it. Contractor and Microsoft may delete Purchasing Entity's Customer Data after that, and specifically with regard to DPT Services will delete Purchasing Entity's Customer Data after no more than 180 days.

Regulatory. If a government rule or regulation applies to Contractor or Microsoft, but not generally to other businesses, and makes it difficult to operate In-Scope Services without change, or Contractor or Microsoft believe this agreement or an In-Scope Service may conflict with the rule or regulation, Microsoft may change the applicable In-Scope Service(s) or Contractor may terminate the agreement. If Microsoft changes In-Scope Services to come into compliance, and Purchasing Entity does not like the change, Purchasing Entity may terminate.

Suspension. Microsoft may suspend use of In-Scope Services: (1) if reasonably needed to prevent unauthorized Customer Data access; (2) if Purchasing Entity does not promptly respond under §5 to intellectual property claims; (3) for non-payment; or (4) if Purchasing Entity violates the AUP. A suspension will be in effect only while the condition or need exists and, if under clause (1) or (2), will apply to the minimum extent necessary. Contractor or Microsoft will notify Purchasing Entity before suspension, unless doing so may increase damages. Contractor will notify Purchasing Entity at least 30 days before suspending for non-payment. If Purchasing Entity does not fully address the reasons for suspension within 60 days after suspension, Contractor may terminate Purchasing Entity's Subscription.

4. Limited warranty; disclaimer.

Contractor warrants that In-Scope Services will meet the SLA terms during the Subscription; Purchasing Entity's only remedy for breach of warranty is stated in the SLA. *Contractor provides no (and disclaims to*

the extent permitted by law any) other warranties, express, implied, or statutory, including warranties of merchantability or fitness for a particular purpose.

5. Indemnification.

a. Subject to the exceptions below and the NASPO Participants' (as defined below) compliance with the notice and defense provisions below, in the event of any defect or deficiency in any Microsoft Products or Services purchased by a Participating Entity or Purchasing Entity, Microsoft agrees to defend NASPO, NASPO ValuePoint, the Lead State, Participating Entities, and Purchasing Entities, along with their officers, agents, and employees (collectively, the "NASPO Participants") against third party claims, damages or causes of action including reasonable attorneys' fees and related costs for any death, injury, or damage to tangible property suffered by such third party and caused by the negligence, or willful misconduct of Microsoft, its employees or subcontractors or volunteers, at any tier, during the performance of this Master Agreement (a "PI Claim"). This clause shall not be construed to bar any legal remedies Microsoft may have with respect to the NASPO Participants' failure to fulfill their obligations pursuant to the Master Agreement or any Participating Addendum.

To qualify for such defense, the Participating Entities/Purchasing Entities shall promptly notify Microsoft of any PI Claim of which the Participating Entities/Purchasing Entities become aware which may give rise to a right of defense pursuant to this Section. Notice of any PI Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Participating Entities'/Purchasing Entities' first learning of such proceeding. If the Participating Entity's/Purchasing Entity's laws require approval of a third party to defend Participating Entity/Purchasing Entity, Participating Entity/Purchasing Entity will seek such approval and if approval is not received, Microsoft is not required to defend that Participating Entity/Purchasing Entity. If a PI Claim is settled, to the extent permitted by law, the Participating Entities/Purchasing Entities shall not publicize the settlement and will cooperate with Microsoft so that Microsoft can make every effort to ensure the settlement agreement contains a non-disclosure provision.

Notwithstanding anything to the contrary contained herein, Participating Entities/Purchasing Entities agree that Microsoft has no obligation for any PI Claim covered by this Section arising out of or resulting from the Participating Entities'/Purchasing Entities' or any of their respective employees', contractors' or agents' acts of negligence, gross negligence or misconduct. THE FOREGOING SHALL CONSTITUTE EACH AND EVERY PARTICIPATING ENTITY'S/PURCHASING ENTITY'S SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL PI CLAIMS.

b. **Indemnification** – Intellectual Property. Subject to the NASPO Participant's compliance with the notice and defense requirements and exceptions set forth below, Microsoft agrees to defend the NASPO Participants against any claims made by an unaffiliated third party that any Product or Service or its proper or reasonably expected or acceptable use infringes that third party's patent, copyright, or trademark or makes unlawful use of its trade secret (an "Intellectual Property Claim").

(1) Microsoft's obligations under this section shall not extend to any claims based on:

- (a) Microsoft's compliance with a Participating Entity's/Purchasing Entity's designs, specifications or instructions; or
- (b) Microsoft's use of technical information or technology provided by the Participating Entity/Purchasing Entity; or
- (c) Non-Microsoft software, modifications a Participating Entity/Purchasing Entity makes to, or any specifications or materials a Participating Entity/Purchasing Entity provides or makes available for, a Product; or
- (d) Participating Entity's/Purchasing Entity's combination of the Product or Service with a Non-Microsoft product, data or business process; or damages based on the use of a Non-Microsoft product, data or business process; or

(e) Participating Entity's/Purchasing Entity's use of either Microsoft's trademarks or the use or redistribution of a Product or Service in violation of this Master Agreement, Participating Addendum, or any other agreement incorporating its terms; or

(f) Participating Entity's/Purchasing Entity's use of a Product or Service after Microsoft notifies Participating Entity/Purchasing Entity to discontinue that use due to a third party claim.

(2) To qualify for such defense, the involved NASPO Participants (the "Indemnified Party") shall promptly notify Microsoft of any Intellectual Property Claim of which the Indemnified Party become aware which may give rise to right of defense pursuant to this Section. Notice of any Intellectual Property Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Indemnified Party's learning of such proceeding. If the Indemnified Party's laws require approval of a third party to defend the Indemnified Party, the Indemnified Party will seek such approval and if approval is not received, Microsoft is not required to defend that Indemnified Party. In the event the Indemnified Party does not authorize sole control to Microsoft over any claims that may arise under this subsection, then the parties agree that Microsoft will be granted authorization to equally participate in any proceeding subject to this subsection. The Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Microsoft's reasonable request and expense, information and assistance necessary for such defense. If Microsoft fails to vigorously pursue the defense or settlement of the Intellectual Property Claim, the Indemnified Party may assume the defense or settlement of it and Microsoft shall be liable for all costs and expenses, including reasonable attorneys' fees and related costs, incurred by the Indemnified Party in the pursuit of the Intellectual Property Claim. Unless otherwise agreed in writing, this section is not subject to any limitations of liability in this Master Agreement or in any other document executed in conjunction with this Master Agreement.

If Microsoft reasonably believes that a Product or Service may infringe or misappropriate a third-party's intellectual property rights, Microsoft will seek to:

- i. procure for Participating Entities/Purchasing Entities the right to continue to use the Product or Service; or
- ii. modify or replace it with a functional equivalent to make it non-infringing and notify Participating Entities/Purchasing Entities to discontinue use of the prior version, which Participating Entities/Purchasing Entities must do immediately.

If the foregoing options are not commercially reasonable for Microsoft, or if required by a valid judicial or government order, Microsoft may cause Contractor to terminate Participating Entities'/Purchasing Entities' license or access rights in the Product or Service. In such a case, Microsoft will provide Participating Entities'/Purchasing Entities with notice and refund any amounts Participating Entities'/Purchasing Entities have paid for those rights to the Product or Service.

THE FOREGOING SHALL CONSTITUTE THE LEAD STATE'S AND EACH AND EVERY PARTICIPATING AND/OR PURCHASING ENTITIES' SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL INTELLECTUAL PROPERTY CLAIMS.

6. *Limitation of Liability.*

The parties agree that, other than In-Scope Services, no other Microsoft Products shall be purchased hereunder. As applicable to In-Scope Services, the following terms and conditions shall apply:

a. General. The total liability of each party for In-Scope Services hereunder (Microsoft, Contractor and each Purchasing Entity, including their Affiliates and contractors), for claims arising under this Agreement, is limited to direct damages up to the amount Purchasing Entity paid for the In-Scope Service during the prior 12 months before the cause of action arose; but in no event will a party's aggregate liability for any In-Scope Service exceed the total amount paid for that In-Scope Service under this Agreement. In the case of In-Scope Services provided free of charge, previews, or code that a Purchasing Entity is authorized to redistribute to third parties without separate payment to Contractor, Microsoft's liability is limited to direct damages up to U.S. \$5,000. These limitations apply regardless of whether the liability is based on breach of contract, tort (including negligence), strict liability, breach of warranties, or any other legal theory.

b. Affiliates and contractors. Contractor and Purchasing Entity each agree not to bring any action against the other's Affiliates or contractors in respect of any matter disclaimed on their behalf in this Agreement. Each party will be responsible for its actions in the event of any breach of this provision.

c. EXCLUSION OF CERTAIN DAMAGES. Neither party nor their Affiliates or contractors will be liable for any indirect, consequential, special or incidental damages, or damages for lost profits, revenues, business interruption, or loss of business information in connection with this agreement, even if advised of the possibility of such damages or if such possibility was reasonably foreseeable.

d. Limits. *The limits and exclusions in this section titled "Limitation of liability" do not apply to either party's (1) obligations under the section titled "Defense of third party claims", (2) liability for damages caused by either party's gross negligence or willful misconduct, or that of its employees or its agents, and awarded by a court of final adjudication (provided that, in jurisdictions that do not recognize a legal distinction between "gross negligence" and "negligence," "gross negligence" as used in this subsection shall mean "recklessness"); and (3) liability for violation of its confidentiality obligations (except obligations related to Customer Data) or the other party's intellectual property rights.*

7. Agreement mechanics

Purchasing Entity must send notice by regular mail, return receipt requested, to the address on the Portal (effective when delivered). Contractor or Microsoft may email notice to your account administrators (effective when sent). Purchasing Entity may not assign this agreement, or any right or duty under it. If part of this agreement is held unenforceable, the rest remains in force. Failure to enforce this agreement is not a waiver. The parties are independent contractors. This agreement does not create an agency, partnership, or joint venture. This agreement is governed by the laws applicable to Purchasing Entity, without regard to conflict of laws. This agreement (including the SLA and OST) and Contractor's price sheet are the parties' entire agreement on this subject and supersedes any concurrent or prior communications. Agreement terms that require performance, or apply to events that may occur, after termination or expiration will survive, including §5. In-Scope Service and any Client associated therewith are subject to U.S. export jurisdiction. Purchasing Entity must comply with the U.S. Export Administration Regulations, the International Traffic in Arms Regulations, and end-user, end-use, and destination restrictions. For more information, see <http://www.microsoft.com/exporting/>. Microsoft delivers In-Scope Services, and the rights granted to us also apply to them.

8. Miscellaneous Microsoft Terms and Conditions

A. Accessibility. Microsoft supports the government's obligation to provide accessible technologies to its citizens with disabilities as required by Section 508 of the Rehabilitation Act of 1973, and its state law counterparts (including applicable California provisions). Microsoft encourages its customers (including Purchasing Entities under the Master Agreement and Participating Addenda) to judiciously compare product accessibility performance. The Voluntary Product Accessibility Templates ("VPATs") for the Microsoft technologies used in providing the online services can be found at Microsoft's VPAT page. Further information regarding Microsoft's commitment to accessibility can be found at www.microsoft.com/enable.

B. DPT Services Information Security Policy. For each DPT Service, Microsoft follows a written data security policy ("Information Security Policy") that complies with the control standards and frameworks shown in the table below.

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Office 365 Services	Yes	Yes	Yes	Yes*	Yes*
Microsoft Dynamics CRM Online Services	Yes	Yes	Yes	Yes	Yes

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Microsoft Azure Core Services	Yes	Yes	Yes	Varies**	Varies**
Microsoft Intune Online Services	Yes	Yes	Yes	Yes	Yes

*Does not include Yammer Enterprise.

**Current scope is detailed in the audit report and summarized in the Microsoft Azure Trust Center.

Microsoft may add industry or government standards at any time. Microsoft will not eliminate a standard or framework in the table above, unless it is no longer used in the industry and it is replaced with a successor (if any). Azure Government Services meet a separate set of control standards and frameworks, as detailed on the Microsoft Azure Trust Center.

Subject to non-disclosure obligations, Microsoft will make each Information Security Policy available to Purchasing Entity, along with other information reasonably requested by Purchasing Entity regarding Microsoft security practices and policies.

Purchasing Entity is solely responsible for reviewing each Information Security Policy and making an independent determination as to whether it meets Purchasing Entity's requirements.

C. **Microsoft Audits of DPT Services.** For each DPT Service, Microsoft will conduct audits of the security of the computers, computing environment and physical data centers that it uses in processing Customer Data (including personal data), as follows:

- Where a standard or framework provides for audits, an audit of such control standard or framework will be initiated at least annually for each DPT Service.
- Each audit will be performed according to the standards and rules of the regulatory or accreditation body for each applicable control standard or framework.
- Each audit will be performed by qualified, independent, third party security auditors at Microsoft's selection and expense.

Each audit will result in the generation of an audit report ("Microsoft Audit Report"), which will be Microsoft's Confidential Information. The Microsoft Audit Report will clearly disclose any material findings by the auditor. Microsoft will promptly remediate issues raised in any Microsoft Audit Report to the satisfaction of the auditor.

If Purchasing Entity requests, Microsoft will provide Purchasing Entity with each Microsoft Audit Report so that Purchasing Entity can verify Microsoft's compliance with the security obligations under the DPT. The Microsoft Audit Report will be subject to non-disclosure and distribution limitations of Microsoft and the auditor.

D. **HIPAA/HITECH Business Associate.** If Purchasing Entity is a "covered entity" or a "business associate" and includes "protected health information" in Customer Data, as those terms are defined in 45 CFR § 160.103, this agreement includes execution of the HIPAA Business Associate Agreement ("BAA"), the full text of which identifies the DPT Services to which it applies and is available at <http://aka.ms/BAA>. Purchasing Entity may opt out of the BAA by instructing Contractor to send the following information to Microsoft in a written notice (under the terms of the Contractor's subcontract with Microsoft):

- the full legal name of Purchasing Entity and any Affiliate that is opting out;
- the volume licensing agreement number to which the opt out applies.

E. **FedRAMP.** During the term of a Purchasing Entity's subscription for Government Community Cloud Services, those services will be operated in accordance with a written data security policy and control framework that is consistent with the requirements of NIST 800-53 Revision 4, or successor standards and guidelines (if any), established to support Federal Risk and Authorization Management Program (FedRAMP) accreditation at a Moderate Impact level. Microsoft intends for Government Community Cloud Services to support FedRAMP Authority to Operate ("ATO"), and Microsoft will use commercially reasonable efforts to obtain an ATO from a Federal agency, and to

maintain such ATO through continuous monitoring processes and by conducting regular FedRAMP audits.

- F. **Background Checks.** Microsoft performs the following background checks on all US personnel who have potential to access Customer Data. Adherence to this policy is one of the control procedures addressed by the Microsoft Audit Report per the section of the Microsoft Online Services Terms titled "Microsoft Audits of Online Services." Such Background Checks will be performed in accordance with the Fair Credit Reporting Act and will consist of Social Security Number trace, seven (7) year felony and misdemeanor criminal records check of federal, state, or local records (as applicable) for job related crimes, Office of Foreign Assets Control List (OFAC) check, Bureau of Industry and Security List (BIS) check and Office of Defense Trade Controls Debarred Persons List (DDTC) check.

Additionally, for States with which Microsoft has entered into a Microsoft CJIS Information Agreement and provided an FBI CJIS Addendum Certification to the applicable CJIS Systems Agency (CSA, e.g. the California Department of Justice), for its CJIS Covered Services (which are most of the Microsoft Government Community Cloud Services), Microsoft submits its personnel that have access to unencrypted Criminal Justice Information and other Customer Content (including both employees and subcontractors, where applicable) for FBI NCIC fingerprint background checks and adjudication by each such CSA.

- G. **Click-Through Terms.** For the Microsoft In-Scope Services sold and licensed to Purchasing Entities hereunder, no click-through licensing terms presented to end users or administrators, as they pertain to the delivery, operation or use of such services, shall be binding. For clarity, to the extent that certain services may present terms of use for a web portal used to administer and configure the In-Scope Services, or to download software in conjunction with the In-Scope Services, such terms of use shall be binding, to the extent that there are no equivalent terms and conditions in the Master Agreement which pertain to the use of such web portals.

9. **Definitions.**

"AUP" means the section of the OST titled "Acceptable Use Policy" "Azure Consumption-Based Services" means any Microsoft Online Service made available for Purchasing Entity's use and consumption under the Microsoft Azure Administrative Portal, following Purchasing Entity's initial subscription order for such services, subject to consumption fees billed to Purchasing Entity (or debited against its prepaid Azure monetary commitment balance, if applicable) as described in the applicable Product Terms.

"Client" means device software, if any, that Contractor or Microsoft provides to Purchasing Entity with In-Scope Services.

"Customer Content" means the subset of Customer Data created by users, which includes the following:

- For Office 365 Services, Customer Content shall at least include Exchange Online mailbox content (e-mail body, calendar entries, and the content of e-mail attachments), SharePoint Online site content and the files stored within that site, and Skype for Business Online archived conversations.
- For Microsoft Dynamics CRM Online Services, Customer Content means the entities of Customer Data managed by the Microsoft Dynamics CRM Online Services.
- For Microsoft Azure Core Services, Customer Content includes all Customer Data except as may be noted in the Microsoft Azure Trust Center (which Microsoft may update from time to time, but will not add exceptions for existing Microsoft Azure Core Services in general release).
- For Microsoft Intune Online Services, Customer Content includes all Customer Data except as may be noted in the "Data Location" section of the Microsoft Intune Trust Center.

"Customer Data" means all data, including all text, sound, or image files that are provided to Microsoft by, or on behalf of, Purchasing Entity through its use of In-Scope Services.

"DPT" means the section of the OST titled "Data Processing Terms."

“DPT Services” means the Microsoft Online Services to which the DPT applies. As of the date this agreement was executed by the parties, DPT Services include “Microsoft Dynamics CRM Online Services,” “Office 365 Services,” “Microsoft Intune Online Services,” and “Microsoft Azure Core Services,” which consist of the following component Microsoft Online Services:

DPT Services	
Microsoft Dynamics CRM Online Services	Microsoft Dynamics CRM Online services made available through volume licensing or the Microsoft online services portal, excluding (1) Microsoft Dynamics CRM for supported devices, which includes but is not limited to Microsoft Dynamics CRM Online services for tablets and/or smartphones and (2) any separately-branded service made available with or connected to Microsoft Dynamics CRM Online, such as Microsoft Social Engagement, Parature, from Microsoft, and Microsoft Dynamics Marketing.
Office 365 Services	The following services, each as a standalone service or as included in an Office 365-branded plan or suite: Exchange Online, Exchange Online Archiving, Exchange Online Protection, Advanced Threat Protection, SharePoint Online, OneDrive for Business, Project Online, Skype for Business Online, Sway, Office Online, Delve Analytics, Customer Lockbox, and Yammer Enterprise. Office 365 Services do not include Office 365 ProPlus, any portion of PSTN Services that operate outside of Microsoft’s control, any client software, or any separately branded service made available with an Office 365-branded plan or suite, such as a Bing or a service branded “for Office 365.”
Microsoft Azure Core Services	Cloud Services (web and worker roles), Virtual Machines (including with SQL Server), Storage (Blobs, Tables, Queues), Virtual Network, Traffic Manager, Batch, Web Sites, BizTalk Services, Media Services, Mobile Services, Service Bus, Notification Hub, Workflow Manager, Express Route, Scheduler, Multi-Factor Authentication, Active Directory, Rights Management Service, SQL Database, and HDInsight.
Microsoft Intune Online Services	The cloud service portion of Microsoft Intune such as the Microsoft Intune Add-on Product or a management service provided by Microsoft Intune such as Mobile Device Management for Office 365.

“Government Community Cloud Services” means the Microsoft Online Services identified as such (a subset of DPT Services) in the DPT and Product Terms.

“In-Scope Services” means the Microsoft Online Services for which subscription licenses are sold to Purchasing Entity by Contractor. In-Scope Services include, but are not limited to DPT Services. Except for Azure Consumption-Based Services (which are billed based upon the quantities of such services consumed by Purchasing Entity), Purchasing Entity will order individual licenses from Contractor for all In-Scope Services Licenses.

“Master Agreement” means the agreement (described therein as “NASPO ValuePoint Master Agreement for Cloud Services”) between Contractor and the Lead State, subject to the terms and conditions of the Participating Addendum between Contractor and Purchasing Entity’s State (if different from Utah), and which resulted from award of contract under Utah Solicitation Number CH16012.

“Microsoft” means Microsoft Corporation, which is subcontractor to Contractor and delivers the In-Scope Services.

“Online Services” means the Microsoft-hosted services for which Purchasing Entity acquires Licenses from Contractor hereunder.

“OST” means the Microsoft “Online Services Terms,” which are additional terms that apply to Purchasing Entity’s use of Online Services published on the Volume Licensing Site and updated from time to time, and is subject to terms and conditions in the OST which govern which version of OST is applicable to a particular subscription order.

“Portal” means the Online Services Portal for each In-Scope Service.

“Product Terms” means the document that provides information about Microsoft Products and Professional Services available through volume licensing. The Product Terms document is available on the Volume Licensing Site and is updated from time to time.

“Security Incident” means any unlawful access, use, theft or destruction to any Customer Data stored on Microsoft’s equipment or in Microsoft’s facilities, or unauthorized access to such equipment or facilities resulting in use, theft, loss, disclosure, alteration or destruction of Customer Data. All references to “Data Breach” in the Master Agreement shall be deemed to mean Security Incident.

“Service Level Agreement” (“SLA”) means the document which specifies the standards to which Microsoft agrees to adhere and by which it measures the level of service for an In-Scope Service. Microsoft’s current and archived prior version SLAs are available at the Volume Licensing Site. With respect to any In-Scope Service ordered under the this Agreement, the most current SLA available at the onset of a subscription License term shall apply to that In-Scope Service for the duration of that subscription License Term, after which (for any subsequent renewal subscription License term) it will be superseded by the version current at the time of renewal.

“Subscription” means an order for a quantity of an In-Scope Service, including but not limited to (a) orders for User Subscription Licenses for one or more In-Scope Services; and (b) an order for the purpose of establishing a subscription for Azure Consumption-Based Services.

“Volume Licensing Site” means <http://www.microsoft.com/licensing/contracts> or a successor site.

APPENDIX 13-GOOGLE-1

Exhibit A13.1 lists Google's cloud services.

Exhibit A13.1. Google Cloud Services.

Category	Product Name	Product Description
SaaS	Google Apps Auth User	<i>Google Apps Auth User includes the ability to define Google Apps user identities for managing Android applications and devices only.</i>
SaaS	Google Apps Message Encryption	<i>With Google Apps for Work, Edu, or Unlimited, customers can add Google Apps Message Encryption to set policies to route their outbound email to an encryption server for external recipients to securely receive and send email to their end users.</i>
SaaS	Google Apps No Gmail	<i>Google Apps No Gmail includes Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps for Work	<i>Google Apps for Work includes mail (30 GB of storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps for EDU, Unlimited	<i>For qualifying K-12 and higher education organizations, Google Apps for EDU, Unlimited includes mail, Google Vault eDiscovery and Archiving (unlimited storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps Deskless	<i>Google Apps Deskless includes the same services above; this offering is intended for nontechnical workers and infrequent users.</i>
SaaS	Google Apps Unlimited	<i>Google Apps Unlimited includes mail, Google Vault eDiscovery and Archiving (unlimited storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps Unlimited Deskless	<i>Google Apps Unlimited customers can blend this product into their contract. This offering, which includes the same services above, is intended for nontechnical workers and infrequent users.</i>
SaaS	Google Apps Vault	<i>Google Apps Vault provides Vault eDiscovery and Archiving (unlimited storage).</i>
SaaS	Google Apps Vault Deskless	<i>Google Apps customers can blend this product into their contract. This offering, which includes the same services above, is intended for nontechnical workers and infrequent users.</i>

APPENDIX-UNISYS-1

Exhibits X and Y list the Unisys solutions and services provided by Unisys.

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud.

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
Unisys Justice Law Enforcement and Border Security Solutions for the Cloud		
Secure Image Management Solution	Unisys developed a highly secure solution to meet the forensic integrity standards and policy requirements of several of the largest police departments in the world. The secure storage of images and other multimedia content occurs in a tamperproof information archive – in effect, a digital evidence vault.	All
U-LEAF	U-LEAF provides a POLE-type data model (Person, Object, Location and Event) for the storage and recording of incidents and entities. The POLE model allows entities to be recorded in the system once. The recorded entities, however, can be linked to other entities and events as many times as necessary, to build the picture of an incident, or a network of associations.	All
Unisys Social Services Solutions for the Cloud		
Unisys 311	The Unisys “311 solution” is a Multi- Channel Citizen Service Delivery and Engagement solution that can be used by a Local Government or any Governmental Agency to deliver and manage non-emergency services, respond to inquiries and engage with their constituents. This offers multi-channel interaction including social media and mobile applications, supported by a knowledgebase and GIS systems and also interfaces with existing systems to manage work orders and service requests. The solution can be configured to meet specific needs and extend its functionality and flexibility through other apps, and integration to mobile apps, social media and other systems.	All
State/Local Government Enterprise Regulatory System (AMANDA)	The Unisys AMANDA Platform (from CSDC) is designed to provide a collection of back office functions such as Licensing, Permitting, Inspections, Land Use, Planning and a number of other functions in various form factors (desktop, mobile and tablet) and can also be deployed either on premise, hosted, or delivered via the Cloud. The solution includes the design, delivery, deployment and ongoing support and maintenance of the solution. The AMANDA platform can be configured specific to client needs in order to extend its functionality and interface with other systems as required.	All
Unisys Enterprise Content Management for the Cloud		
Infomage	Unisys Infomage is the Enterprise Content Management (ECM), Business Process Management (BPM) and Record Management (RM) solution for organizations looking to significantly improve business processes that depend heavily high-volume paper documents, documents generated from Internet transactions, Office documents, and other electronic documents that need to be accessed for automated and manual processes. Infomage can easily capture, manage, store, and access the content required for cases, inquiries, and process-centric work, regardless of data structure or document origination, with a single	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	intuitive user interface. Unisys Infolmage brings together ECM, imaging, workflow, document management, record management and integration technologies to form an integrated end-to-end solution.	
Unisys Horizontal Solutions for the Cloud		
Stealth	The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. We built a better way to deal with advanced threats for our clients by applying novel approaches to new threats. By substituting traditional hardware topology for software-based cryptography, our Stealth Microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users.	Private and Hybrid Cloud Only. For Public and Community Cloud, see AWS offering above. IaaS, PaaS and SaaS
VantagePoint	VantagePoint can also extend beyond the boundaries of IT with uses cases to support relevant business scenarios in, • Security: Strengthens the security posture of organizations through greater visibility into the Stealth enabled infrastructure. • Advanced Data Analytics: Cuts across data silos and sources relevant data to optimize the performance of Advanced Analytics solution. • Cloud and Infrastructure Services: Enables hybrid IT by seamlessly integrating with a variety of data sources and services and presenting them via a digital control panel, irrespective of technology underpinnings. • End User Services: Supports millennial users via personalized and secure access to relevant data and services, across endpoints. • Service Management: Cuts across ITSM platforms to improve service delivery through service automation, orchestration, and aggregation capabilities. • Facilities/Crisis Management: Helps organizations monitor their facilities and keep their Facilities Managers abreast of any security threats and risks. Also acts as a standard communication vehicle for users at times of crisis. (178) With VantagePoint, our clients gain a personalized, secure, and intuitive data and service aggregation platform that cuts across strategic and operational dimensions of business and accelerates digital transformation.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys Cloud Hosting	Unisys Hosted Private Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can provide, integrate and scale the cloud infrastructure to meet client's needs. Our management platform enables a single pane of glass to effectively manage client's workloads. Unisys Hybrid Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can source, integrate and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our management platform enables a single pane of glass to effectively	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	manage a variety of clouds within an organization, whether they be public or private.	
ServiceNow	ServiceNow offers enterprise service management software for human resources, law, facilities management, finance, marketing, and field operations in the cloud. ServiceNow has its STAR Self-Assessment available on the CSA's website. ServiceNow specializes in ITSM applications and provides forms-based workflow application development. ServiceNow has open integration options to variety platforms such as: Salesforce, SharePoint, and BMC Remedy Action Request System.	All
Unisys ClearPath Forward!	Unisys ClearPath Forward is an Intel based fabric computing platform designed to run mission critical applications that require predictable performance and low transport latency. The ClearPath Forward platform provides hardware partitioning technology similar to the Mainframes, but is designed to run Windows and Linux operating environments on commodity server hardware. The fabric interconnect included with the ClearPath Forward platform provides transport speeds at 56Gbps or higher. The ClearPath Forward platform fabric can be used in a hybrid mode with VMware.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys White Label Offerings	Unisys has strategic, long lasting relationships in the industry and is pleased to offer our partner products surrounding the cloud market space. This is an evolving list. Our current partner list includes: SalesForce, Oracle, Verizon, SHI, Decision Lens, Aptio, NetApp, VMware, EMC, NetApp, Birst, Okta, Box, WatchDox. We have created webpage on our public site that is updated with our list of current partners. The page is available here: http://www.unisys.com/ms/wsca-cloud-hosting/	All

Exhibit Y. Unisys Services Offerings.

Unisys Professional Services	Description
Cloud Advisory Services	Unisys Cloud Advisory Services provide strategic and financial guidance on aligning IT with business objectives. This starts with a roadmap that outlines the vision of a hybrid IT based on a combination of existing data center, internal cloud, and external cloud resources to provide agility, flexibility, and control.
Data Center Planning, Design, and Implementation Services	Unisys Data Center Planning, Design, and Implementation Services offer a complete range of services that delivers cohesive, end-to-end optimization of data centers. With a wide range of services for discovery, analysis, optimization, virtualization, consolidation, and migration of data centers that can complement client efforts and fill gaps in skills and capacities and a combination of world-class people, processes, and technology with program and project management expertise, Unisys transforms clients' existing data centers to a business engine that provides agility at a lower cost.
CloudBuild Services	Unisys CloudBuild Services enable organizations to successfully build a cloud that is integrated with the overall business process, transforming their existing infrastructure to an agile IT-as-a-service model. The Unisys "8-Tracks" model, a 360-degree cloud view approach, covers eight critical data center domains. Together with ConOps, which includes industry best practices, Unisys enables a client's cloud infrastructure to meet the security, regulatory, and compliance requirements that enable us to deliver the most secure and reliable cloud in the marketplace.
Hybrid Cloud Strategy	Unisys Hybrid Cloud strategy helps organizations overcome key challenges when planning to implement a cloud environment. Unisys provides governments and businesses with a comprehensive cloud architecture that gives clients cloud on their

Unisys Professional Services	Description
	terms. We can source, integrate, and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our homegrown VantagePoint management platform enables a single pane of glass to effectively manage a variety of clouds in an organization, whether they be public or private.
Unisys Platform Services	
Platform Assessment Services	Unisys Platform Assessment Services enable governments and businesses to foster their platform adoption initiatives. Unisys analyzes organizations' business goals and objectives along with their technology landscape of their enterprise application portfolio and recommend a best-fit Platform Suitability Analysis to help them make the right decision in their move to PaaS.
Architecture Design Services	Unisys Architecture Design Services assist governments and businesses with designing an effective architecture suited for cloud. We provide several concepts and best practices that are essential to build highly scalable applications in the cloud – be it on-premises, public, or hybrid.
PaaS Enablement Services	Unisys PaaS Enablement Services analyze an organization's current set of applications and provide a strategy to use the best practices and engagement models of various platform providers such as Microsoft Azure, IBM Bluemix, SAP HCP, and SFDC AppCloud.
Application Migration Services	Unisys Application Migration Services assist organizations in their cloud migration activities by using tools and frameworks adhering to Unisys best practices that promise minimal downtime without affecting the day-to-day business processes.
Application Portability Services	Unisys Application Portability Services help organizations to develop cross-platform applications that can be scaled across multiple cloud platforms.
Application Development Services	Unisys Application Development Services help organizations to build rich, interactive applications focused on business logic and workflows using visual tools as well as cloud-based tools, architectures, and services that make their applications cloud ready.
IoT Development Services	Unisys IoT Development Services enable organizations to use IoT Services across various cloud platforms that enable apps to communicate and consume data collected by the connected devices, sensors, and gateways.
API Management Services	Unisys API Management Services provide a solution that addresses the aspects of the application programming interface (API) life cycle for on-premises and cloud environments and offer capabilities to create, run, manage, secure, and monetize APIs and microservices that deliver an integrated user experience and enable rapid deployment and simplified administration of APIs.
Horizontal Technology Integration Services	
Horizontal Technology Integration Services	Unisys is a world-recognized leader in integrating technology infrastructures across platforms and vendors. We particularly specialize in integrating solutions surrounding service desks, cloud provisioning, data center management, and the ITIL framework.
Unisys Service Management Services	
Maturity and Platform Technology Assessments	Unisys Service Management Consulting Services leverage subject matter experts with more than 15 years of experience in working and delivering services for ITIL and service management. We leverage that knowledge to provide best practices, lessons learned, and roadmap development for Service Management disciplines and benchmark them with industry and operational best practices.
Service Management Platform Implementation	Unisys brings experience in moving existing ITSM tool information, processes, and requirements to new cloud enterprise Service Management platforms. By leveraging industry best practices and structured methodologies, Unisys can establish transition and transformation plans that minimize risk to clients' ongoing operations and deliver speed to value with the new cloud-based Service Management platform.
Enterprise Cloud Service Management Platform Support	Unisys can provide ongoing day-to-day support in the ongoing management of a client's Service Management platform. This service allows the organization to focus on the business aspects instead of the routine activities necessary to support and maintain a platform. Unisys provides cost-effective solutions that leverage staff in many locations, driving the cost of servicing the environment to effective and efficient levels.

APPENDIX 13-GOOGLE-1

Exhibit A13.1 lists Google's cloud services.

Exhibit A13.1. Google Cloud Services.

Category	Product Name	Product Description
SaaS	Google Apps Auth User	<i>Google Apps Auth User includes the ability to define Google Apps user identities for managing Android applications and devices only.</i>
SaaS	Google Apps Message Encryption	<i>With Google Apps for Work, Edu, or Unlimited, customers can add Google Apps Message Encryption to set policies to route their outbound email to an encryption server for external recipients to securely receive and send email to their end users.</i>
SaaS	Google Apps No Gmail	<i>Google Apps No Gmail includes Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps for Work	<i>Google Apps for Work includes mail (30 GB of storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps for EDU, Unlimited	<i>For qualifying K-12 and higher education organizations, Google Apps for EDU, Unlimited includes mail, Google Vault eDiscovery and Archiving (unlimited storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps Deskless	<i>Google Apps Deskless includes the same services above; this offering is intended for nontechnical workers and infrequent users.</i>
SaaS	Google Apps Unlimited	<i>Google Apps Unlimited includes mail, Google Vault eDiscovery and Archiving (unlimited storage), a calendar, Google Drive, Mobile Device Management, groups, G+, sites (unlimited), videoconferencing, and VoIP.</i>
SaaS	Google Apps Unlimited Deskless	<i>Google Apps Unlimited customers can blend this product into their contract. This offering, which includes the same services above, is intended for nontechnical workers and infrequent users.</i>
SaaS	Google Apps Vault	<i>Google Apps Vault provides Vault eDiscovery and Archiving (unlimited storage).</i>
SaaS	Google Apps Vault Deskless	<i>Google Apps customers can blend this product into their contract. This offering, which includes the same services above, is intended for nontechnical workers and infrequent users.</i>

APPENDIX-UNISYS-1

Exhibits X and Y list the Unisys solutions and services provided by Unisys.

Exhibit X. Unisys Public Sector Solution Offerings in the Cloud.

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
Unisys Justice Law Enforcement and Border Security Solutions for the Cloud		
Secure Image Management Solution	Unisys developed a highly secure solution to meet the forensic integrity standards and policy requirements of several of the largest police departments in the world. The secure storage of images and other multimedia content occurs in a tamperproof information archive – in effect, a digital evidence vault.	All
U-LEAF	U-LEAF provides a POLE-type data model (Person, Object, Location and Event) for the storage and recording of incidents and entities. The POLE model allows entities to be recorded in the system once. The recorded entities, however, can be linked to other entities and events as many times as necessary, to build the picture of an incident, or a network of associations.	All
Unisys Social Services Solutions for the Cloud		
Unisys 311	The Unisys “311 solution” is a Multi- Channel Citizen Service Delivery and Engagement solution that can be used by a Local Government or any Governmental Agency to deliver and manage non-emergency services, respond to inquiries and engage with their constituents. This offers multi-channel interaction including social media and mobile applications, supported by a knowledgebase and GIS systems and also interfaces with existing systems to manage work orders and service requests. The solution can be configured to meet specific needs and extend its functionality and flexibility through other apps, and integration to mobile apps, social media and other systems.	All
State/Local Government Enterprise Regulatory System (AMANDA)	The Unisys AMANDA Platform (from CSDC) is designed to provide a collection of back office functions such as Licensing, Permitting, Inspections, Land Use, Planning and a number of other functions in various form factors (desktop, mobile and tablet) and can also be deployed either on premise, hosted, or delivered via the Cloud. The solution includes the design, delivery, deployment and ongoing support and maintenance of the solution. The AMANDA platform can be configured specific to client needs in order to extend its functionality and interface with other systems as required.	All
Unisys Enterprise Content Management for the Cloud		
Infomage	Unisys Infomage is the Enterprise Content Management (ECM), Business Process Management (BPM) and Record Management (RM) solution for organizations looking to significantly improve business processes that depend heavily high-volume paper documents, documents generated from Internet transactions, Office documents, and other electronic documents that need to be accessed for automated and manual processes. Infomage can easily capture, manage, store, and access the content required for cases, inquiries, and process-centric work, regardless of data structure or document origination, with a single	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	intuitive user interface. Unisys Infolmage brings together ECM, imaging, workflow, document management, record management and integration technologies to form an integrated end-to-end solution.	
Unisys Horizontal Solutions for the Cloud		
Stealth	The Unisys Stealth software-defined security portfolio delivers consistent, inimitable security for global enterprises focused on protecting data in their data center, cloud, and mobile infrastructures. We built a better way to deal with advanced threats for our clients by applying novel approaches to new threats. By substituting traditional hardware topology for software-based cryptography, our Stealth Microsegmentation solutions prevent unauthorized access to sensitive information and reduce the attack surface, thereby making end points invisible to unauthorized users.	Private and Hybrid Cloud Only. For Public and Community Cloud, see AWS offering above. IaaS, PaaS and SaaS
VantagePoint	VantagePoint can also extend beyond the boundaries of IT with uses cases to support relevant business scenarios in, • Security: Strengthens the security posture of organizations through greater visibility into the Stealth enabled infrastructure. • Advanced Data Analytics: Cuts across data silos and sources relevant data to optimize the performance of Advanced Analytics solution. • Cloud and Infrastructure Services: Enables hybrid IT by seamlessly integrating with a variety of data sources and services and presenting them via a digital control panel, irrespective of technology underpinnings. • End User Services: Supports millennial users via personalized and secure access to relevant data and services, across endpoints. • Service Management: Cuts across ITSM platforms to improve service delivery through service automation, orchestration, and aggregation capabilities. • Facilities/Crisis Management: Helps organizations monitor their facilities and keep their Facilities Managers abreast of any security threats and risks. Also acts as a standard communication vehicle for users at times of crisis. (178) With VantagePoint, our clients gain a personalized, secure, and intuitive data and service aggregation platform that cuts across strategic and operational dimensions of business and accelerates digital transformation.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys Cloud Hosting	Unisys Hosted Private Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can provide, integrate and scale the cloud infrastructure to meet client's needs. Our management platform enables a single pane of glass to effectively manage client's workloads. Unisys Hybrid Cloud Services provide businesses and governments a comprehensive cloud architecture that gives customers cloud on their terms. We can source, integrate and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our management platform enables a single pane of glass to effectively	All

Unisys Public Sector Solutions	Description	Availability (Public, Private, Community, Hybrid) (IaaS, PaaS, SaaS)
	manage a variety of clouds within an organization, whether they be public or private.	
ServiceNow	ServiceNow offers enterprise service management software for human resources, law, facilities management, finance, marketing, and field operations in the cloud. ServiceNow has its STAR Self-Assessment available on the CSA's website. ServiceNow specializes in ITSM applications and provides forms-based workflow application development. ServiceNow has open integration options to variety platforms such as: Salesforce, SharePoint, and BMC Remedy Action Request System.	All
Unisys ClearPath Forward!	Unisys ClearPath Forward is an Intel based fabric computing platform designed to run mission critical applications that require predictable performance and low transport latency. The ClearPath Forward platform provides hardware partitioning technology similar to the Mainframes, but is designed to run Windows and Linux operating environments on commodity server hardware. The fabric interconnect included with the ClearPath Forward platform provides transport speeds at 56Gbps or higher. The ClearPath Forward platform fabric can be used in a hybrid mode with VMware.	Private and Hybrid Only IaaS, PaaS and SaaS
Unisys White Label Offerings	Unisys has strategic, long lasting relationships in the industry and is pleased to offer our partner products surrounding the cloud market space. This is an evolving list. Our current partner list includes: SalesForce, Oracle, Verizon, SHI, Decision Lens, Aptio, NetApp, VMware, EMC, NetApp, Birst, Okta, Box, WatchDox. We have created webpage on our public site that is updated with our list of current partners. The page is available here: http://www.unisys.com/ms/wsca-cloud-hosting/	All

Exhibit Y. Unisys Services Offerings.

Unisys Professional Services	Description
Cloud Advisory Services	Unisys Cloud Advisory Services provide strategic and financial guidance on aligning IT with business objectives. This starts with a roadmap that outlines the vision of a hybrid IT based on a combination of existing data center, internal cloud, and external cloud resources to provide agility, flexibility, and control.
Data Center Planning, Design, and Implementation Services	Unisys Data Center Planning, Design, and Implementation Services offer a complete range of services that delivers cohesive, end-to-end optimization of data centers. With a wide range of services for discovery, analysis, optimization, virtualization, consolidation, and migration of data centers that can complement client efforts and fill gaps in skills and capacities and a combination of world-class people, processes, and technology with program and project management expertise, Unisys transforms clients' existing data centers to a business engine that provides agility at a lower cost.
CloudBuild Services	Unisys CloudBuild Services enable organizations to successfully build a cloud that is integrated with the overall business process, transforming their existing infrastructure to an agile IT-as-a-service model. The Unisys "8-Tracks" model, a 360-degree cloud view approach, covers eight critical data center domains. Together with ConOps, which includes industry best practices, Unisys enables a client's cloud infrastructure to meet the security, regulatory, and compliance requirements that enable us to deliver the most secure and reliable cloud in the marketplace.
Hybrid Cloud Strategy	Unisys Hybrid Cloud strategy helps organizations overcome key challenges when planning to implement a cloud environment. Unisys provides governments and businesses with a comprehensive cloud architecture that gives clients cloud on their

Unisys Professional Services	Description
	terms. We can source, integrate, and scale commodity infrastructure such as Microsoft Azure and Amazon Web Services. Our homegrown VantagePoint management platform enables a single pane of glass to effectively manage a variety of clouds in an organization, whether they be public or private.
Unisys Platform Services	
Platform Assessment Services	Unisys Platform Assessment Services enable governments and businesses to foster their platform adoption initiatives. Unisys analyzes organizations' business goals and objectives along with their technology landscape of their enterprise application portfolio and recommend a best-fit Platform Suitability Analysis to help them make the right decision in their move to PaaS.
Architecture Design Services	Unisys Architecture Design Services assist governments and businesses with designing an effective architecture suited for cloud. We provide several concepts and best practices that are essential to build highly scalable applications in the cloud – be it on-premises, public, or hybrid.
PaaS Enablement Services	Unisys PaaS Enablement Services analyze an organization's current set of applications and provide a strategy to use the best practices and engagement models of various platform providers such as Microsoft Azure, IBM Bluemix, SAP HCP, and SFDC AppCloud.
Application Migration Services	Unisys Application Migration Services assist organizations in their cloud migration activities by using tools and frameworks adhering to Unisys best practices that promise minimal downtime without affecting the day-to-day business processes.
Application Portability Services	Unisys Application Portability Services help organizations to develop cross-platform applications that can be scaled across multiple cloud platforms.
Application Development Services	Unisys Application Development Services help organizations to build rich, interactive applications focused on business logic and workflows using visual tools as well as cloud-based tools, architectures, and services that make their applications cloud ready.
IoT Development Services	Unisys IoT Development Services enable organizations to use IoT Services across various cloud platforms that enable apps to communicate and consume data collected by the connected devices, sensors, and gateways.
API Management Services	Unisys API Management Services provide a solution that addresses the aspects of the application programming interface (API) life cycle for on-premises and cloud environments and offer capabilities to create, run, manage, secure, and monetize APIs and microservices that deliver an integrated user experience and enable rapid deployment and simplified administration of APIs.
Horizontal Technology Integration Services	
Horizontal Technology Integration Services	Unisys is a world-recognized leader in integrating technology infrastructures across platforms and vendors. We particularly specialize in integrating solutions surrounding service desks, cloud provisioning, data center management, and the ITIL framework.
Unisys Service Management Services	
Maturity and Platform Technology Assessments	Unisys Service Management Consulting Services leverage subject matter experts with more than 15 years of experience in working and delivering services for ITIL and service management. We leverage that knowledge to provide best practices, lessons learned, and roadmap development for Service Management disciplines and benchmark them with industry and operational best practices.
Service Management Platform Implementation	Unisys brings experience in moving existing ITSM tool information, processes, and requirements to new cloud enterprise Service Management platforms. By leveraging industry best practices and structured methodologies, Unisys can establish transition and transformation plans that minimize risk to clients' ongoing operations and deliver speed to value with the new cloud-based Service Management platform.
Enterprise Cloud Service Management Platform Support	Unisys can provide ongoing day-to-day support in the ongoing management of a client's Service Management platform. This service allows the organization to focus on the business aspects instead of the routine activities necessary to support and maintain a platform. Unisys provides cost-effective solutions that leverage staff in many locations, driving the cost of servicing the environment to effective and efficient levels.

Appendix A

Additional Use Rights and Restrictions

These Additional Use Rights and Restrictions shall apply to each Enterprise's use of Products and must be incorporated in each Government Contract.

1. General.

Right to use. Purchasing Entity may access and use In-Scope Services, and install and use a Client (if any) included with its Subscription, only as described in this agreement. All other rights are reserved.

Acceptable use. Purchasing Entity will use In-Scope Services only per the AUP. Purchasing Entity will not use In-Scope Services in any way that infringes a third party's patent, copyright, or trademark or misappropriates its trade secret. Purchasing Entity may not reverse engineer, decompile, work around technical limits in, or disassemble In-Scope Services, except if applicable law permits despite this limit. Purchasing Entity may not rent, lease, lend, resell, transfer, or host In-Scope Services to or for third parties.

Compliance.

Purchasing Entity will comply with all laws and regulations applicable to its use of In-Scope Services.

In providing In-Scope Services, Contractor and Microsoft will comply with all laws and regulations (including applicable security breach notification law) that generally apply to IT service providers. Microsoft will comply with all laws and regulations applicable to its provision of the Online Services, including security breach notification law. However, Microsoft is not responsible for compliance with any laws or regulations applicable to Customer or Customer's industry or government function that are not generally applicable to information technology service providers. Microsoft does not determine whether Customer Data includes information subject to any specific law or regulation. All Security Incidents are subject to the Security Incident Notification terms below.

Purchasing Entity will obtain any consents required: (1) to allow it to access, monitor, use, and disclose user data; and (2) for Contractor and Microsoft to provide the In-Scope Services. If Purchasing Entity is an educational institution, it will obtain any parental consent for end users' use of In-Scope Services as required by applicable law.

Customer Data. Customer Data is used only to provide Purchasing Entity In-Scope Services. This use may include troubleshooting to prevent, find and fix problems with those In-Scope Services' operation. It may also include improving features for finding and protecting against threats to users. Neither Contractor nor Microsoft will derive information from Customer Data for any advertising or other commercial purposes. As between the parties, the Purchasing Entity retains all right, title and interest in and to Customer Data. Neither Contractor nor Microsoft acquires rights in Customer Data, other than the rights Purchasing Entity grants to Microsoft through Contractor to provide the Online Services to Purchasing Entity. This paragraph does not affect Microsoft's rights in software or Online Services Contractor and Microsoft license to Purchasing Entity.

Contractor and Microsoft will enable Purchasing Entity to keep Customer Data in DPT Services separate from Microsoft's consumer services.

Customer Data will not be disclosed unless required by law or allowed by this agreement. Purchasing Entity's contact information may be provided so that a requestor can contact it. If law requires disclosure, Contractor and Microsoft will use commercially reasonable efforts to notify Purchasing Entity, if permitted.

Customer Data may be transferred to, and stored and processed in, any country Contractor or Microsoft maintain facilities, except when a tenant for DPT Services is provisioned in the United States, in which case Customer Content shall be processed and stored at rest in the Continental United States.

Security Incident Notification for DPT Services. Solely for DPT Services, the following terms and conditions apply:

If Microsoft becomes aware of any Security Incident (as defined herein), Microsoft will promptly (1) notify Purchasing Entity of the Security Incident; (2) investigate the Security Incident and provide Purchasing Entity with detailed information about the Security Incident; and (3) take reasonable steps to mitigate the effects and to minimize any damage resulting from the Security Incident.

Notification(s) of Security Incidents will be delivered to one or more of Purchasing Entity's administrators by any means Microsoft selects, including via email. It is Purchasing Entity's sole responsibility to ensure Purchasing Entity's administrators maintain accurate contact information on each applicable Online Services portal. Microsoft's obligation to report or respond to a Security Incident under this section is not an acknowledgement by Microsoft of any fault or liability with respect to the Security Incident.

Purchasing Entity must notify Microsoft promptly about any possible misuse of its accounts or authentication credentials or any Security Incident related to an Online Service.

5-day Security Incident notification for Government Community Cloud Services. As an exception to the foregoing, notification of Security Incident solely pertaining to Government Community Cloud Services will be delivered within 5 days after Microsoft determines that a Security Incident has occurred, provided that Purchasing Entity must comply with the following requirements:

For each Online Service Tenant or Azure subscription, as applicable, as a condition of receiving notifications within 5 days, as set forth in the preceding paragraph, Purchasing Entity must register the following information by sending email to ols-notifications@microsoft.com, and must keep such information current at all times:

- 1) Purchasing Entity's Microsoft Online Direct Routing Domain (MODRD);
- 2) For one or more individual(s) to be contacted, each of whom must be registered as an administrator on the applicable Online Services, each of the following:
 - a. Name;
 - b. Title;
 - c. Email address registered as an administrator on the Online Services;
 - d. Email address not registered as a user on the Online Services;
- 3) Name of Purchasing Entity;
- 4) Enrollment number (provided by Contractor and assigned by Microsoft) to represent Purchasing Entity's Tenant or Azure Subscription, on Contractor's subcontract with Microsoft.

Microsoft expects to change the above process by which the Purchasing Entity for each tenant or subscription will be able to register their MODRD and other information for five-day Security Incident notification pursuant to these terms and conditions. In the event that a Purchasing Entity is notified by Microsoft, in the administrative console or otherwise, of revised instructions necessary to ensure five-day Security Incident notification, the Purchasing Entity must comply with such revised instructions.

Reimbursement of Security Incident Remediation Costs. To the extent that a Security Incident for a DPT Service results from Microsoft's failure to comply with its obligations under the Master Agreement (and, where applicable, Participating Addenda), and subject to the Limitation of Liability section in this Agreement, Microsoft will reimburse Purchasing Entities for reasonable out-of-pocket remediation costs incurred by such Purchasing Entities in connection with that Security Incident. "Reasonable out-of-pocket remediation costs" are costs that (a) are customary, reasonable and expected to be paid by entities similar to Purchasing Entity, based on the nature and scope of the Security Incident, and (b) do not arise from or relate to Purchasing Entity's violation of (i) laws applicable to Purchasing Entity or (ii) Purchasing Entity's obligations to third parties, and (c) in no event include costs arising related to compliance with laws applicable to Purchasing Entity or its industry or government function that are not generally applicable to information technology services providers. Purchasing Entity must document all such expenditures and, upon Microsoft's request, those expenditures must be validated by an independent, internationally-recognized

third party industry expert chosen by both parties. For avoidance of doubt, the costs reimbursed by Microsoft under this paragraph will be characterized as direct damages subject to the limitation on liability set forth in this Section, and not as special damages excluded under the "EXCLUSION OF CERTAIN DAMAGES" in the Limitation of Liability section.

Changes. In-Scope Services may be changed periodically. As such, Purchasing Entity's use of them will be subject to a revised OST, then-current as of the date it renews its subscription. For In-Scope Services that include client software, Purchasing Entity may be required to run a client software upgrade on its devices after a change to maintain full functionality, in accordance with the terms and conditions of the OST.

Use rights. Use rights specific to each In-Scope Service are posted online at the link to the OST, which is hereby incorporated by reference into this agreement

2. Confidentiality and Security.

For DPT Services, Contractor and Microsoft will (a) implement and maintain all appropriate administrative, physical, technical, procedural safeguards and organizational measures, internal controls, and data security routines intended to protect Customer Data against accidental loss or change, unauthorized disclosure or access, unlawful destruction, hacks, introduction of viruses, disabling devices, malware, and other forms of malicious or inadvertent acts that can disrupt Purchasing Entity's access to its Customer Data; and (b) not disclose Customer Data, except as required by law or expressly allowed. Neither party will make any public statement about this agreement's terms without the other's prior written consent.

3. Term, Termination, and Suspension.

Term and termination. This agreement will remain in effect for three years subject to each party's right under the Master Agreement and applicable law to terminate for convenience. It may be extended, based upon the mutual consent of the parties, during the term of the Master Agreement.

Customer Data. Purchasing Entity may extract Customer Data at any time. If Purchasing Entity's Subscription expires or terminates, Contractor and Microsoft will keep your Customer Data in a limited account for at least 90 days so Purchasing Entity may extract it. Contractor and Microsoft may delete Purchasing Entity's Customer Data after that, and specifically with regard to DPT Services will delete Purchasing Entity's Customer Data after no more than 180 days.

Regulatory. If a government rule or regulation applies to Contractor or Microsoft, but not generally to other businesses, and makes it difficult to operate In-Scope Services without change, or Contractor or Microsoft believe this agreement or an In-Scope Service may conflict with the rule or regulation, Microsoft may change the applicable In-Scope Service(s) or Contractor may terminate the agreement. If Microsoft changes In-Scope Services to come into compliance, and Purchasing Entity does not like the change, Purchasing Entity may terminate.

Suspension. Microsoft may suspend use of In-Scope Services: (1) if reasonably needed to prevent unauthorized Customer Data access; (2) if Purchasing Entity does not promptly respond under §5 to intellectual property claims; (3) for non-payment; or (4) if Purchasing Entity violates the AUP. A suspension will be in effect only while the condition or need exists and, if under clause (1) or (2), will apply to the minimum extent necessary. Contractor or Microsoft will notify Purchasing Entity before suspension, unless doing so may increase damages. Contractor will notify Purchasing Entity at least 30 days before suspending for non-payment. If Purchasing Entity does not fully address the reasons for suspension within 60 days after suspension, Contractor may terminate Purchasing Entity's Subscription.

4. Limited warranty; disclaimer.

Contractor warrants that In-Scope Services will meet the SLA terms during the Subscription; Purchasing Entity's only remedy for breach of warranty is stated in the SLA. *Contractor provides no (and disclaims to*

the extent permitted by law any) other warranties, express, implied, or statutory, including warranties of merchantability or fitness for a particular purpose.

5. Indemnification.

a. Subject to the exceptions below and the NASPO Participants' (as defined below) compliance with the notice and defense provisions below, in the event of any defect or deficiency in any Microsoft Products or Services purchased by a Participating Entity or Purchasing Entity, Microsoft agrees to defend NASPO, NASPO ValuePoint, the Lead State, Participating Entities, and Purchasing Entities, along with their officers, agents, and employees (collectively, the "NASPO Participants") against third party claims, damages or causes of action including reasonable attorneys' fees and related costs for any death, injury, or damage to tangible property suffered by such third party and caused by the negligence, or willful misconduct of Microsoft, its employees or subcontractors or volunteers, at any tier, during the performance of this Master Agreement (a "PI Claim"). This clause shall not be construed to bar any legal remedies Microsoft may have with respect to the NASPO Participants' failure to fulfill their obligations pursuant to the Master Agreement or any Participating Addendum.

To qualify for such defense, the Participating Entities/Purchasing Entities shall promptly notify Microsoft of any PI Claim of which the Participating Entities/Purchasing Entities become aware which may give rise to a right of defense pursuant to this Section. Notice of any PI Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Participating Entities'/Purchasing Entities' first learning of such proceeding. If the Participating Entity's/Purchasing Entity's laws require approval of a third party to defend Participating Entity/Purchasing Entity, Participating Entity/Purchasing Entity will seek such approval and if approval is not received, Microsoft is not required to defend that Participating Entity/Purchasing Entity. If a PI Claim is settled, to the extent permitted by law, the Participating Entities/Purchasing Entities shall not publicize the settlement and will cooperate with Microsoft so that Microsoft can make every effort to ensure the settlement agreement contains a non-disclosure provision.

Notwithstanding anything to the contrary contained herein, Participating Entities/Purchasing Entities agree that Microsoft has no obligation for any PI Claim covered by this Section arising out of or resulting from the Participating Entities'/Purchasing Entities' or any of their respective employees', contractors' or agents' acts of negligence, gross negligence or misconduct. THE FOREGOING SHALL CONSTITUTE EACH AND EVERY PARTICIPATING ENTITY'S/PURCHASING ENTITY'S SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL PI CLAIMS.

b. **Indemnification** – Intellectual Property. Subject to the NASPO Participant's compliance with the notice and defense requirements and exceptions set forth below, Microsoft agrees to defend the NASPO Participants against any claims made by an unaffiliated third party that any Product or Service or its proper or reasonably expected or acceptable use infringes that third party's patent, copyright, or trademark or makes unlawful use of its trade secret (an "Intellectual Property Claim").

(1) Microsoft's obligations under this section shall not extend to any claims based on:

- (a) Microsoft's compliance with a Participating Entity's/Purchasing Entity's designs, specifications or instructions; or
- (b) Microsoft's use of technical information or technology provided by the Participating Entity/Purchasing Entity; or
- (c) Non-Microsoft software, modifications a Participating Entity/Purchasing Entity makes to, or any specifications or materials a Participating Entity/Purchasing Entity provides or makes available for, a Product; or
- (d) Participating Entity's/Purchasing Entity's combination of the Product or Service with a Non-Microsoft product, data or business process; or damages based on the use of a Non-Microsoft product, data or business process; or

(e) Participating Entity's/Purchasing Entity's use of either Microsoft's trademarks or the use or redistribution of a Product or Service in violation of this Master Agreement, Participating Addendum, or any other agreement incorporating its terms; or

(f) Participating Entity's/Purchasing Entity's use of a Product or Service after Microsoft notifies Participating Entity/Purchasing Entity to discontinue that use due to a third party claim.

(2) To qualify for such defense, the involved NASPO Participants (the "Indemnified Party") shall promptly notify Microsoft of any Intellectual Property Claim of which the Indemnified Party become aware which may give rise to right of defense pursuant to this Section. Notice of any Intellectual Property Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Indemnified Party's learning of such proceeding. If the Indemnified Party's laws require approval of a third party to defend the Indemnified Party, the Indemnified Party will seek such approval and if approval is not received, Microsoft is not required to defend that Indemnified Party. In the event the Indemnified Party does not authorize sole control to Microsoft over any claims that may arise under this subsection, then the parties agree that Microsoft will be granted authorization to equally participate in any proceeding subject to this subsection. The Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Microsoft's reasonable request and expense, information and assistance necessary for such defense. If Microsoft fails to vigorously pursue the defense or settlement of the Intellectual Property Claim, the Indemnified Party may assume the defense or settlement of it and Microsoft shall be liable for all costs and expenses, including reasonable attorneys' fees and related costs, incurred by the Indemnified Party in the pursuit of the Intellectual Property Claim. Unless otherwise agreed in writing, this section is not subject to any limitations of liability in this Master Agreement or in any other document executed in conjunction with this Master Agreement.

If Microsoft reasonably believes that a Product or Service may infringe or misappropriate a third-party's intellectual property rights, Microsoft will seek to:

- i. procure for Participating Entities/Purchasing Entities the right to continue to use the Product or Service; or
- ii. modify or replace it with a functional equivalent to make it non-infringing and notify Participating Entities/Purchasing Entities to discontinue use of the prior version, which Participating Entities/Purchasing Entities must do immediately.

If the foregoing options are not commercially reasonable for Microsoft, or if required by a valid judicial or government order, Microsoft may cause Contractor to terminate Participating Entities'/Purchasing Entities' license or access rights in the Product or Service. In such a case, Microsoft will provide Participating Entities'/Purchasing Entities with notice and refund any amounts Participating Entities'/Purchasing Entities have paid for those rights to the Product or Service.

THE FOREGOING SHALL CONSTITUTE THE LEAD STATE'S AND EACH AND EVERY PARTICIPATING AND/OR PURCHASING ENTITIES' SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL INTELLECTUAL PROPERTY CLAIMS.

6. *Limitation of Liability.*

The parties agree that, other than In-Scope Services, no other Microsoft Products shall be purchased hereunder. As applicable to In-Scope Services, the following terms and conditions shall apply:

a. General. The total liability of each party for In-Scope Services hereunder (Microsoft, Contractor and each Purchasing Entity, including their Affiliates and contractors), for claims arising under this Agreement, is limited to direct damages up to the amount Purchasing Entity paid for the In-Scope Service during the prior 12 months before the cause of action arose; but in no event will a party's aggregate liability for any In-Scope Service exceed the total amount paid for that In-Scope Service under this Agreement. In the case of In-Scope Services provided free of charge, previews, or code that a Purchasing Entity is authorized to redistribute to third parties without separate payment to Contractor, Microsoft's liability is limited to direct damages up to U.S. \$5,000. These limitations apply regardless of whether the liability is based on breach of contract, tort (including negligence), strict liability, breach of warranties, or any other legal theory.

b. Affiliates and contractors. Contractor and Purchasing Entity each agree not to bring any action against the other's Affiliates or contractors in respect of any matter disclaimed on their behalf in this Agreement. Each party will be responsible for its actions in the event of any breach of this provision.

c. EXCLUSION OF CERTAIN DAMAGES. Neither party nor their Affiliates or contractors will be liable for any indirect, consequential, special or incidental damages, or damages for lost profits, revenues, business interruption, or loss of business information in connection with this agreement, even if advised of the possibility of such damages or if such possibility was reasonably foreseeable.

d. Limits. *The limits and exclusions in this section titled "Limitation of liability" do not apply to either party's (1) obligations under the section titled "Defense of third party claims", (2) liability for damages caused by either party's gross negligence or willful misconduct, or that of its employees or its agents, and awarded by a court of final adjudication (provided that, in jurisdictions that do not recognize a legal distinction between "gross negligence" and "negligence," "gross negligence" as used in this subsection shall mean "recklessness"); and (3) liability for violation of its confidentiality obligations (except obligations related to Customer Data) or the other party's intellectual property rights.*

7. Agreement mechanics

Purchasing Entity must send notice by regular mail, return receipt requested, to the address on the Portal (effective when delivered). Contractor or Microsoft may email notice to your account administrators (effective when sent). Purchasing Entity may not assign this agreement, or any right or duty under it. If part of this agreement is held unenforceable, the rest remains in force. Failure to enforce this agreement is not a waiver. The parties are independent contractors. This agreement does not create an agency, partnership, or joint venture. This agreement is governed by the laws applicable to Purchasing Entity, without regard to conflict of laws. This agreement (including the SLA and OST) and Contractor's price sheet are the parties' entire agreement on this subject and supersedes any concurrent or prior communications. Agreement terms that require performance, or apply to events that may occur, after termination or expiration will survive, including §5. In-Scope Service and any Client associated therewith are subject to U.S. export jurisdiction. Purchasing Entity must comply with the U.S. Export Administration Regulations, the International Traffic in Arms Regulations, and end-user, end-use, and destination restrictions. For more information, see <http://www.microsoft.com/exporting/>. Microsoft delivers In-Scope Services, and the rights granted to us also apply to them.

8. Miscellaneous Microsoft Terms and Conditions

A. Accessibility. Microsoft supports the government's obligation to provide accessible technologies to its citizens with disabilities as required by Section 508 of the Rehabilitation Act of 1973, and its state law counterparts (including applicable California provisions). Microsoft encourages its customers (including Purchasing Entities under the Master Agreement and Participating Addenda) to judiciously compare product accessibility performance. The Voluntary Product Accessibility Templates ("VPATs") for the Microsoft technologies used in providing the online services can be found at Microsoft's VPAT page. Further information regarding Microsoft's commitment to accessibility can be found at www.microsoft.com/enable.

B. DPT Services Information Security Policy. For each DPT Service, Microsoft follows a written data security policy ("Information Security Policy") that complies with the control standards and frameworks shown in the table below.

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Office 365 Services	Yes	Yes	Yes	Yes*	Yes*
Microsoft Dynamics CRM Online Services	Yes	Yes	Yes	Yes	Yes

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Microsoft Azure Core Services	Yes	Yes	Yes	Varies**	Varies**
Microsoft Intune Online Services	Yes	Yes	Yes	Yes	Yes

*Does not include Yammer Enterprise.

**Current scope is detailed in the audit report and summarized in the Microsoft Azure Trust Center.

Microsoft may add industry or government standards at any time. Microsoft will not eliminate a standard or framework in the table above, unless it is no longer used in the industry and it is replaced with a successor (if any). Azure Government Services meet a separate set of control standards and frameworks, as detailed on the Microsoft Azure Trust Center.

Subject to non-disclosure obligations, Microsoft will make each Information Security Policy available to Purchasing Entity, along with other information reasonably requested by Purchasing Entity regarding Microsoft security practices and policies.

Purchasing Entity is solely responsible for reviewing each Information Security Policy and making an independent determination as to whether it meets Purchasing Entity's requirements.

C. **Microsoft Audits of DPT Services.** For each DPT Service, Microsoft will conduct audits of the security of the computers, computing environment and physical data centers that it uses in processing Customer Data (including personal data), as follows:

- Where a standard or framework provides for audits, an audit of such control standard or framework will be initiated at least annually for each DPT Service.
- Each audit will be performed according to the standards and rules of the regulatory or accreditation body for each applicable control standard or framework.
- Each audit will be performed by qualified, independent, third party security auditors at Microsoft's selection and expense.

Each audit will result in the generation of an audit report ("Microsoft Audit Report"), which will be Microsoft's Confidential Information. The Microsoft Audit Report will clearly disclose any material findings by the auditor. Microsoft will promptly remediate issues raised in any Microsoft Audit Report to the satisfaction of the auditor.

If Purchasing Entity requests, Microsoft will provide Purchasing Entity with each Microsoft Audit Report so that Purchasing Entity can verify Microsoft's compliance with the security obligations under the DPT. The Microsoft Audit Report will be subject to non-disclosure and distribution limitations of Microsoft and the auditor.

D. **HIPAA/HITECH Business Associate.** If Purchasing Entity is a "covered entity" or a "business associate" and includes "protected health information" in Customer Data, as those terms are defined in 45 CFR § 160.103, this agreement includes execution of the HIPAA Business Associate Agreement ("BAA"), the full text of which identifies the DPT Services to which it applies and is available at <http://aka.ms/BAA>. Purchasing Entity may opt out of the BAA by instructing Contractor to send the following information to Microsoft in a written notice (under the terms of the Contractor's subcontract with Microsoft):

- the full legal name of Purchasing Entity and any Affiliate that is opting out;
- the volume licensing agreement number to which the opt out applies.

E. **FedRAMP.** During the term of a Purchasing Entity's subscription for Government Community Cloud Services, those services will be operated in accordance with a written data security policy and control framework that is consistent with the requirements of NIST 800-53 Revision 4, or successor standards and guidelines (if any), established to support Federal Risk and Authorization Management Program (FedRAMP) accreditation at a Moderate Impact level. Microsoft intends for Government Community Cloud Services to support FedRAMP Authority to Operate ("ATO"), and Microsoft will use commercially reasonable efforts to obtain an ATO from a Federal agency, and to

maintain such ATO through continuous monitoring processes and by conducting regular FedRAMP audits.

- F. **Background Checks.** Microsoft performs the following background checks on all US personnel who have potential to access Customer Data. Adherence to this policy is one of the control procedures addressed by the Microsoft Audit Report per the section of the Microsoft Online Services Terms titled "Microsoft Audits of Online Services." Such Background Checks will be performed in accordance with the Fair Credit Reporting Act and will consist of Social Security Number trace, seven (7) year felony and misdemeanor criminal records check of federal, state, or local records (as applicable) for job related crimes, Office of Foreign Assets Control List (OFAC) check, Bureau of Industry and Security List (BIS) check and Office of Defense Trade Controls Debarred Persons List (DDTC) check.

Additionally, for States with which Microsoft has entered into a Microsoft CJIS Information Agreement and provided an FBI CJIS Addendum Certification to the applicable CJIS Systems Agency (CSA, e.g. the California Department of Justice), for its CJIS Covered Services (which are most of the Microsoft Government Community Cloud Services), Microsoft submits its personnel that have access to unencrypted Criminal Justice Information and other Customer Content (including both employees and subcontractors, where applicable) for FBI NCIC fingerprint background checks and adjudication by each such CSA.

- G. **Click-Through Terms.** For the Microsoft In-Scope Services sold and licensed to Purchasing Entities hereunder, no click-through licensing terms presented to end users or administrators, as they pertain to the delivery, operation or use of such services, shall be binding. For clarity, to the extent that certain services may present terms of use for a web portal used to administer and configure the In-Scope Services, or to download software in conjunction with the In-Scope Services, such terms of use shall be binding, to the extent that there are no equivalent terms and conditions in the Master Agreement which pertain to the use of such web portals.

9. **Definitions.**

"AUP" means the section of the OST titled "Acceptable Use Policy" "Azure Consumption-Based Services" means any Microsoft Online Service made available for Purchasing Entity's use and consumption under the Microsoft Azure Administrative Portal, following Purchasing Entity's initial subscription order for such services, subject to consumption fees billed to Purchasing Entity (or debited against its prepaid Azure monetary commitment balance, if applicable) as described in the applicable Product Terms.

"Client" means device software, if any, that Contractor or Microsoft provides to Purchasing Entity with In-Scope Services.

"Customer Content" means the subset of Customer Data created by users, which includes the following:

- For Office 365 Services, Customer Content shall at least include Exchange Online mailbox content (e-mail body, calendar entries, and the content of e-mail attachments), SharePoint Online site content and the files stored within that site, and Skype for Business Online archived conversations.
- For Microsoft Dynamics CRM Online Services, Customer Content means the entities of Customer Data managed by the Microsoft Dynamics CRM Online Services.
- For Microsoft Azure Core Services, Customer Content includes all Customer Data except as may be noted in the Microsoft Azure Trust Center (which Microsoft may update from time to time, but will not add exceptions for existing Microsoft Azure Core Services in general release).
- For Microsoft Intune Online Services, Customer Content includes all Customer Data except as may be noted in the "Data Location" section of the Microsoft Intune Trust Center.

"Customer Data" means all data, including all text, sound, or image files that are provided to Microsoft by, or on behalf of, Purchasing Entity through its use of In-Scope Services.

"DPT" means the section of the OST titled "Data Processing Terms."

“DPT Services” means the Microsoft Online Services to which the DPT applies. As of the date this agreement was executed by the parties, DPT Services include “Microsoft Dynamics CRM Online Services,” “Office 365 Services,” “Microsoft Intune Online Services,” and “Microsoft Azure Core Services,” which consist of the following component Microsoft Online Services:

DPT Services	
Microsoft Dynamics CRM Online Services	Microsoft Dynamics CRM Online services made available through volume licensing or the Microsoft online services portal, excluding (1) Microsoft Dynamics CRM for supported devices, which includes but is not limited to Microsoft Dynamics CRM Online services for tablets and/or smartphones and (2) any separately-branded service made available with or connected to Microsoft Dynamics CRM Online, such as Microsoft Social Engagement, Parature, from Microsoft, and Microsoft Dynamics Marketing.
Office 365 Services	The following services, each as a standalone service or as included in an Office 365-branded plan or suite: Exchange Online, Exchange Online Archiving, Exchange Online Protection, Advanced Threat Protection, SharePoint Online, OneDrive for Business, Project Online, Skype for Business Online, Sway, Office Online, Delve Analytics, Customer Lockbox, and Yammer Enterprise. Office 365 Services do not include Office 365 ProPlus, any portion of PSTN Services that operate outside of Microsoft’s control, any client software, or any separately branded service made available with an Office 365-branded plan or suite, such as a Bing or a service branded “for Office 365.”
Microsoft Azure Core Services	Cloud Services (web and worker roles), Virtual Machines (including with SQL Server), Storage (Blobs, Tables, Queues), Virtual Network, Traffic Manager, Batch, Web Sites, BizTalk Services, Media Services, Mobile Services, Service Bus, Notification Hub, Workflow Manager, Express Route, Scheduler, Multi-Factor Authentication, Active Directory, Rights Management Service, SQL Database, and HDInsight.
Microsoft Intune Online Services	The cloud service portion of Microsoft Intune such as the Microsoft Intune Add-on Product or a management service provided by Microsoft Intune such as Mobile Device Management for Office 365.

“Government Community Cloud Services” means the Microsoft Online Services identified as such (a subset of DPT Services) in the DPT and Product Terms.

“In-Scope Services” means the Microsoft Online Services for which subscription licenses are sold to Purchasing Entity by Contractor. In-Scope Services include, but are not limited to DPT Services. Except for Azure Consumption-Based Services (which are billed based upon the quantities of such services consumed by Purchasing Entity), Purchasing Entity will order individual licenses from Contractor for all In-Scope Services Licenses.

“Master Agreement” means the agreement (described therein as “NASPO ValuePoint Master Agreement for Cloud Services”) between Contractor and the Lead State, subject to the terms and conditions of the Participating Addendum between Contractor and Purchasing Entity’s State (if different from Utah), and which resulted from award of contract under Utah Solicitation Number CH16012.

“Microsoft” means Microsoft Corporation, which is subcontractor to Contractor and delivers the In-Scope Services.

“Online Services” means the Microsoft-hosted services for which Purchasing Entity acquires Licenses from Contractor hereunder.

“OST” means the Microsoft “Online Services Terms,” which are additional terms that apply to Purchasing Entity’s use of Online Services published on the Volume Licensing Site and updated from time to time, and is subject to terms and conditions in the OST which govern which version of OST is applicable to a particular subscription order.

“Portal” means the Online Services Portal for each In-Scope Service.

“Product Terms” means the document that provides information about Microsoft Products and Professional Services available through volume licensing. The Product Terms document is available on the Volume Licensing Site and is updated from time to time.

“Security Incident” means any unlawful access, use, theft or destruction to any Customer Data stored on Microsoft’s equipment or in Microsoft’s facilities, or unauthorized access to such equipment or facilities resulting in use, theft, loss, disclosure, alteration or destruction of Customer Data. All references to “Data Breach” in the Master Agreement shall be deemed to mean Security Incident.

“Service Level Agreement” (“SLA”) means the document which specifies the standards to which Microsoft agrees to adhere and by which it measures the level of service for an In-Scope Service. Microsoft’s current and archived prior version SLAs are available at the Volume Licensing Site. With respect to any In-Scope Service ordered under the this Agreement, the most current SLA available at the onset of a subscription License term shall apply to that In-Scope Service for the duration of that subscription License Term, after which (for any subsequent renewal subscription License term) it will be superseded by the version current at the time of renewal.

“Subscription” means an order for a quantity of an In-Scope Service, including but not limited to (a) orders for User Subscription Licenses for one or more In-Scope Services; and (b) an order for the purpose of establishing a subscription for Azure Consumption-Based Services.

“Volume Licensing Site” means <http://www.microsoft.com/licensing/contracts> or a successor site.

Appendix A

Additional Use Rights and Restrictions

These Additional Use Rights and Restrictions shall apply to each Enterprise's use of Products and must be incorporated in each Government Contract.

1. General.

Right to use. Purchasing Entity may access and use In-Scope Services, and install and use a Client (if any) included with its Subscription, only as described in this agreement. All other rights are reserved.

Acceptable use. Purchasing Entity will use In-Scope Services only per the AUP. Purchasing Entity will not use In-Scope Services in any way that infringes a third party's patent, copyright, or trademark or misappropriates its trade secret. Purchasing Entity may not reverse engineer, decompile, work around technical limits in, or disassemble In-Scope Services, except if applicable law permits despite this limit. Purchasing Entity may not rent, lease, lend, resell, transfer, or host In-Scope Services to or for third parties.

Compliance.

Purchasing Entity will comply with all laws and regulations applicable to its use of In-Scope Services.

In providing In-Scope Services, Contractor and Microsoft will comply with all laws and regulations (including applicable security breach notification law) that generally apply to IT service providers. Microsoft will comply with all laws and regulations applicable to its provision of the Online Services, including security breach notification law. However, Microsoft is not responsible for compliance with any laws or regulations applicable to Customer or Customer's industry or government function that are not generally applicable to information technology service providers. Microsoft does not determine whether Customer Data includes information subject to any specific law or regulation. All Security Incidents are subject to the Security Incident Notification terms below.

Purchasing Entity will obtain any consents required: (1) to allow it to access, monitor, use, and disclose user data; and (2) for Contractor and Microsoft to provide the In-Scope Services. If Purchasing Entity is an educational institution, it will obtain any parental consent for end users' use of In-Scope Services as required by applicable law.

Customer Data. Customer Data is used only to provide Purchasing Entity In-Scope Services. This use may include troubleshooting to prevent, find and fix problems with those In-Scope Services' operation. It may also include improving features for finding and protecting against threats to users. Neither Contractor nor Microsoft will derive information from Customer Data for any advertising or other commercial purposes. As between the parties, the Purchasing Entity retains all right, title and interest in and to Customer Data. Neither Contractor nor Microsoft acquires rights in Customer Data, other than the rights Purchasing Entity grants to Microsoft through Contractor to provide the Online Services to Purchasing Entity. This paragraph does not affect Microsoft's rights in software or Online Services Contractor and Microsoft license to Purchasing Entity.

Contractor and Microsoft will enable Purchasing Entity to keep Customer Data in DPT Services separate from Microsoft's consumer services.

Customer Data will not be disclosed unless required by law or allowed by this agreement. Purchasing Entity's contact information may be provided so that a requestor can contact it. If law requires disclosure, Contractor and Microsoft will use commercially reasonable efforts to notify Purchasing Entity, if permitted.

Customer Data may be transferred to, and stored and processed in, any country Contractor or Microsoft maintain facilities, except when a tenant for DPT Services is provisioned in the United States, in which case Customer Content shall be processed and stored at rest in the Continental United States.

Security Incident Notification for DPT Services. Solely for DPT Services, the following terms and conditions apply:

If Microsoft becomes aware of any Security Incident (as defined herein), Microsoft will promptly (1) notify Purchasing Entity of the Security Incident; (2) investigate the Security Incident and provide Purchasing Entity with detailed information about the Security Incident; and (3) take reasonable steps to mitigate the effects and to minimize any damage resulting from the Security Incident.

Notification(s) of Security Incidents will be delivered to one or more of Purchasing Entity's administrators by any means Microsoft selects, including via email. It is Purchasing Entity's sole responsibility to ensure Purchasing Entity's administrators maintain accurate contact information on each applicable Online Services portal. Microsoft's obligation to report or respond to a Security Incident under this section is not an acknowledgement by Microsoft of any fault or liability with respect to the Security Incident.

Purchasing Entity must notify Microsoft promptly about any possible misuse of its accounts or authentication credentials or any Security Incident related to an Online Service.

5-day Security Incident notification for Government Community Cloud Services. As an exception to the foregoing, notification of Security Incident solely pertaining to Government Community Cloud Services will be delivered within 5 days after Microsoft determines that a Security Incident has occurred, provided that Purchasing Entity must comply with the following requirements:

For each Online Service Tenant or Azure subscription, as applicable, as a condition of receiving notifications within 5 days, as set forth in the preceding paragraph, Purchasing Entity must register the following information by sending email to ols-notifications@microsoft.com, and must keep such information current at all times:

- 1) Purchasing Entity's Microsoft Online Direct Routing Domain (MODRD);
- 2) For one or more individual(s) to be contacted, each of whom must be registered as an administrator on the applicable Online Services, each of the following:
 - a. Name;
 - b. Title;
 - c. Email address registered as an administrator on the Online Services;
 - d. Email address not registered as a user on the Online Services;
- 3) Name of Purchasing Entity;
- 4) Enrollment number (provided by Contractor and assigned by Microsoft) to represent Purchasing Entity's Tenant or Azure Subscription, on Contractor's subcontract with Microsoft.

Microsoft expects to change the above process by which the Purchasing Entity for each tenant or subscription will be able to register their MODRD and other information for five-day Security Incident notification pursuant to these terms and conditions. In the event that a Purchasing Entity is notified by Microsoft, in the administrative console or otherwise, of revised instructions necessary to ensure five-day Security Incident notification, the Purchasing Entity must comply with such revised instructions.

Reimbursement of Security Incident Remediation Costs. To the extent that a Security Incident for a DPT Service results from Microsoft's failure to comply with its obligations under the Master Agreement (and, where applicable, Participating Addenda), and subject to the Limitation of Liability section in this Agreement, Microsoft will reimburse Purchasing Entities for reasonable out-of-pocket remediation costs incurred by such Purchasing Entities in connection with that Security Incident. "Reasonable out-of-pocket remediation costs" are costs that (a) are customary, reasonable and expected to be paid by entities similar to Purchasing Entity, based on the nature and scope of the Security Incident, and (b) do not arise from or relate to Purchasing Entity's violation of (i) laws applicable to Purchasing Entity or (ii) Purchasing Entity's obligations to third parties, and (c) in no event include costs arising related to compliance with laws applicable to Purchasing Entity or its industry or government function that are not generally applicable to information technology services providers. Purchasing Entity must document all such expenditures and, upon Microsoft's request, those expenditures must be validated by an independent, internationally-recognized

third party industry expert chosen by both parties. For avoidance of doubt, the costs reimbursed by Microsoft under this paragraph will be characterized as direct damages subject to the limitation on liability set forth in this Section, and not as special damages excluded under the "EXCLUSION OF CERTAIN DAMAGES" in the Limitation of Liability section.

Changes. In-Scope Services may be changed periodically. As such, Purchasing Entity's use of them will be subject to a revised OST, then-current as of the date it renews its subscription. For In-Scope Services that include client software, Purchasing Entity may be required to run a client software upgrade on its devices after a change to maintain full functionality, in accordance with the terms and conditions of the OST.

Use rights. Use rights specific to each In-Scope Service are posted online at the link to the OST, which is hereby incorporated by reference into this agreement

2. Confidentiality and Security.

For DPT Services, Contractor and Microsoft will (a) implement and maintain all appropriate administrative, physical, technical, procedural safeguards and organizational measures, internal controls, and data security routines intended to protect Customer Data against accidental loss or change, unauthorized disclosure or access, unlawful destruction, hacks, introduction of viruses, disabling devices, malware, and other forms of malicious or inadvertent acts that can disrupt Purchasing Entity's access to its Customer Data; and (b) not disclose Customer Data, except as required by law or expressly allowed. Neither party will make any public statement about this agreement's terms without the other's prior written consent.

3. Term, Termination, and Suspension.

Term and termination. This agreement will remain in effect for three years subject to each party's right under the Master Agreement and applicable law to terminate for convenience. It may be extended, based upon the mutual consent of the parties, during the term of the Master Agreement.

Customer Data. Purchasing Entity may extract Customer Data at any time. If Purchasing Entity's Subscription expires or terminates, Contractor and Microsoft will keep your Customer Data in a limited account for at least 90 days so Purchasing Entity may extract it. Contractor and Microsoft may delete Purchasing Entity's Customer Data after that, and specifically with regard to DPT Services will delete Purchasing Entity's Customer Data after no more than 180 days.

Regulatory. If a government rule or regulation applies to Contractor or Microsoft, but not generally to other businesses, and makes it difficult to operate In-Scope Services without change, or Contractor or Microsoft believe this agreement or an In-Scope Service may conflict with the rule or regulation, Microsoft may change the applicable In-Scope Service(s) or Contractor may terminate the agreement. If Microsoft changes In-Scope Services to come into compliance, and Purchasing Entity does not like the change, Purchasing Entity may terminate.

Suspension. Microsoft may suspend use of In-Scope Services: (1) if reasonably needed to prevent unauthorized Customer Data access; (2) if Purchasing Entity does not promptly respond under §5 to intellectual property claims; (3) for non-payment; or (4) if Purchasing Entity violates the AUP. A suspension will be in effect only while the condition or need exists and, if under clause (1) or (2), will apply to the minimum extent necessary. Contractor or Microsoft will notify Purchasing Entity before suspension, unless doing so may increase damages. Contractor will notify Purchasing Entity at least 30 days before suspending for non-payment. If Purchasing Entity does not fully address the reasons for suspension within 60 days after suspension, Contractor may terminate Purchasing Entity's Subscription.

4. Limited warranty; disclaimer.

Contractor warrants that In-Scope Services will meet the SLA terms during the Subscription; Purchasing Entity's only remedy for breach of warranty is stated in the SLA. *Contractor provides no (and disclaims to*

the extent permitted by law any) other warranties, express, implied, or statutory, including warranties of merchantability or fitness for a particular purpose.

5. Indemnification.

a. Subject to the exceptions below and the NASPO Participants' (as defined below) compliance with the notice and defense provisions below, in the event of any defect or deficiency in any Microsoft Products or Services purchased by a Participating Entity or Purchasing Entity, Microsoft agrees to defend NASPO, NASPO ValuePoint, the Lead State, Participating Entities, and Purchasing Entities, along with their officers, agents, and employees (collectively, the "NASPO Participants") against third party claims, damages or causes of action including reasonable attorneys' fees and related costs for any death, injury, or damage to tangible property suffered by such third party and caused by the negligence, or willful misconduct of Microsoft, its employees or subcontractors or volunteers, at any tier, during the performance of this Master Agreement (a "PI Claim"). This clause shall not be construed to bar any legal remedies Microsoft may have with respect to the NASPO Participants' failure to fulfill their obligations pursuant to the Master Agreement or any Participating Addendum.

To qualify for such defense, the Participating Entities/Purchasing Entities shall promptly notify Microsoft of any PI Claim of which the Participating Entities/Purchasing Entities become aware which may give rise to a right of defense pursuant to this Section. Notice of any PI Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Participating Entities'/Purchasing Entities' first learning of such proceeding. If the Participating Entity's/Purchasing Entity's laws require approval of a third party to defend Participating Entity/Purchasing Entity, Participating Entity/Purchasing Entity will seek such approval and if approval is not received, Microsoft is not required to defend that Participating Entity/Purchasing Entity. If a PI Claim is settled, to the extent permitted by law, the Participating Entities/Purchasing Entities shall not publicize the settlement and will cooperate with Microsoft so that Microsoft can make every effort to ensure the settlement agreement contains a non-disclosure provision.

Notwithstanding anything to the contrary contained herein, Participating Entities/Purchasing Entities agree that Microsoft has no obligation for any PI Claim covered by this Section arising out of or resulting from the Participating Entities'/Purchasing Entities' or any of their respective employees', contractors' or agents' acts of negligence, gross negligence or misconduct. THE FOREGOING SHALL CONSTITUTE EACH AND EVERY PARTICIPATING ENTITY'S/PURCHASING ENTITY'S SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL PI CLAIMS.

b. **Indemnification** – Intellectual Property. Subject to the NASPO Participant's compliance with the notice and defense requirements and exceptions set forth below, Microsoft agrees to defend the NASPO Participants against any claims made by an unaffiliated third party that any Product or Service or its proper or reasonably expected or acceptable use infringes that third party's patent, copyright, or trademark or makes unlawful use of its trade secret (an "Intellectual Property Claim").

(1) Microsoft's obligations under this section shall not extend to any claims based on:

- (a) Microsoft's compliance with a Participating Entity's/Purchasing Entity's designs, specifications or instructions; or
- (b) Microsoft's use of technical information or technology provided by the Participating Entity/Purchasing Entity; or
- (c) Non-Microsoft software, modifications a Participating Entity/Purchasing Entity makes to, or any specifications or materials a Participating Entity/Purchasing Entity provides or makes available for, a Product; or
- (d) Participating Entity's/Purchasing Entity's combination of the Product or Service with a Non-Microsoft product, data or business process; or damages based on the use of a Non-Microsoft product, data or business process; or

(e) Participating Entity's/Purchasing Entity's use of either Microsoft's trademarks or the use or redistribution of a Product or Service in violation of this Master Agreement, Participating Addendum, or any other agreement incorporating its terms; or

(f) Participating Entity's/Purchasing Entity's use of a Product or Service after Microsoft notifies Participating Entity/Purchasing Entity to discontinue that use due to a third party claim.

(2) To qualify for such defense, the involved NASPO Participants (the "Indemnified Party") shall promptly notify Microsoft of any Intellectual Property Claim of which the Indemnified Party become aware which may give rise to right of defense pursuant to this Section. Notice of any Intellectual Property Claim that is a legal proceeding, by suit or otherwise, must be provided to Microsoft within thirty (30) days of the Indemnified Party's learning of such proceeding. If the Indemnified Party's laws require approval of a third party to defend the Indemnified Party, the Indemnified Party will seek such approval and if approval is not received, Microsoft is not required to defend that Indemnified Party. In the event the Indemnified Party does not authorize sole control to Microsoft over any claims that may arise under this subsection, then the parties agree that Microsoft will be granted authorization to equally participate in any proceeding subject to this subsection. The Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Microsoft's reasonable request and expense, information and assistance necessary for such defense. If Microsoft fails to vigorously pursue the defense or settlement of the Intellectual Property Claim, the Indemnified Party may assume the defense or settlement of it and Microsoft shall be liable for all costs and expenses, including reasonable attorneys' fees and related costs, incurred by the Indemnified Party in the pursuit of the Intellectual Property Claim. Unless otherwise agreed in writing, this section is not subject to any limitations of liability in this Master Agreement or in any other document executed in conjunction with this Master Agreement.

If Microsoft reasonably believes that a Product or Service may infringe or misappropriate a third-party's intellectual property rights, Microsoft will seek to:

- i. procure for Participating Entities/Purchasing Entities the right to continue to use the Product or Service; or
- ii. modify or replace it with a functional equivalent to make it non-infringing and notify Participating Entities/Purchasing Entities to discontinue use of the prior version, which Participating Entities/Purchasing Entities must do immediately.

If the foregoing options are not commercially reasonable for Microsoft, or if required by a valid judicial or government order, Microsoft may cause Contractor to terminate Participating Entities'/Purchasing Entities' license or access rights in the Product or Service. In such a case, Microsoft will provide Participating Entities'/Purchasing Entities with notice and refund any amounts Participating Entities'/Purchasing Entities have paid for those rights to the Product or Service.

THE FOREGOING SHALL CONSTITUTE THE LEAD STATE'S AND EACH AND EVERY PARTICIPATING AND/OR PURCHASING ENTITIES' SOLE REMEDY AND MICROSOFT'S SOLE AND EXCLUSIVE LIABILITY FOR ALL INTELLECTUAL PROPERTY CLAIMS.

6. *Limitation of Liability.*

The parties agree that, other than In-Scope Services, no other Microsoft Products shall be purchased hereunder. As applicable to In-Scope Services, the following terms and conditions shall apply:

a. General. The total liability of each party for In-Scope Services hereunder (Microsoft, Contractor and each Purchasing Entity, including their Affiliates and contractors), for claims arising under this Agreement, is limited to direct damages up to the amount Purchasing Entity paid for the In-Scope Service during the prior 12 months before the cause of action arose; but in no event will a party's aggregate liability for any In-Scope Service exceed the total amount paid for that In-Scope Service under this Agreement. In the case of In-Scope Services provided free of charge, previews, or code that a Purchasing Entity is authorized to redistribute to third parties without separate payment to Contractor, Microsoft's liability is limited to direct damages up to U.S. \$5,000. These limitations apply regardless of whether the liability is based on breach of contract, tort (including negligence), strict liability, breach of warranties, or any other legal theory.

b. Affiliates and contractors. Contractor and Purchasing Entity each agree not to bring any action against the other's Affiliates or contractors in respect of any matter disclaimed on their behalf in this Agreement. Each party will be responsible for its actions in the event of any breach of this provision.

c. EXCLUSION OF CERTAIN DAMAGES. Neither party nor their Affiliates or contractors will be liable for any indirect, consequential, special or incidental damages, or damages for lost profits, revenues, business interruption, or loss of business information in connection with this agreement, even if advised of the possibility of such damages or if such possibility was reasonably foreseeable.

d. Limits. *The limits and exclusions in this section titled "Limitation of liability" do not apply to either party's (1) obligations under the section titled "Defense of third party claims", (2) liability for damages caused by either party's gross negligence or willful misconduct, or that of its employees or its agents, and awarded by a court of final adjudication (provided that, in jurisdictions that do not recognize a legal distinction between "gross negligence" and "negligence," "gross negligence" as used in this subsection shall mean "recklessness"); and (3) liability for violation of its confidentiality obligations (except obligations related to Customer Data) or the other party's intellectual property rights.*

7. Agreement mechanics

Purchasing Entity must send notice by regular mail, return receipt requested, to the address on the Portal (effective when delivered). Contractor or Microsoft may email notice to your account administrators (effective when sent). Purchasing Entity may not assign this agreement, or any right or duty under it. If part of this agreement is held unenforceable, the rest remains in force. Failure to enforce this agreement is not a waiver. The parties are independent contractors. This agreement does not create an agency, partnership, or joint venture. This agreement is governed by the laws applicable to Purchasing Entity, without regard to conflict of laws. This agreement (including the SLA and OST) and Contractor's price sheet are the parties' entire agreement on this subject and supersedes any concurrent or prior communications. Agreement terms that require performance, or apply to events that may occur, after termination or expiration will survive, including §5. In-Scope Service and any Client associated therewith are subject to U.S. export jurisdiction. Purchasing Entity must comply with the U.S. Export Administration Regulations, the International Traffic in Arms Regulations, and end-user, end-use, and destination restrictions. For more information, see <http://www.microsoft.com/exporting/>. Microsoft delivers In-Scope Services, and the rights granted to us also apply to them.

8. Miscellaneous Microsoft Terms and Conditions

- A. Accessibility.** Microsoft supports the government's obligation to provide accessible technologies to its citizens with disabilities as required by Section 508 of the Rehabilitation Act of 1973, and its state law counterparts (including applicable California provisions). Microsoft encourages its customers (including Purchasing Entities under the Master Agreement and Participating Addenda) to judiciously compare product accessibility performance. The Voluntary Product Accessibility Templates ("VPATs") for the Microsoft technologies used in providing the online services can be found at Microsoft's VPAT page. Further information regarding Microsoft's commitment to accessibility can be found at www.microsoft.com/enable.
- B. DPT Services Information Security Policy.** For each DPT Service, Microsoft follows a written data security policy ("Information Security Policy") that complies with the control standards and frameworks shown in the table below.

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Office 365 Services	Yes	Yes	Yes	Yes*	Yes*
Microsoft Dynamics CRM Online Services	Yes	Yes	Yes	Yes	Yes

DPT Service	ISO 27001	ISO 27002 Code of Practice	ISO 27018 Code of Practice	SSAE 16 SOC 1 Type II	SSAE 16 SOC 2 Type II
Microsoft Azure Core Services	Yes	Yes	Yes	Varies**	Varies**
Microsoft Intune Online Services	Yes	Yes	Yes	Yes	Yes

*Does not include Yammer Enterprise.

**Current scope is detailed in the audit report and summarized in the Microsoft Azure Trust Center.

Microsoft may add industry or government standards at any time. Microsoft will not eliminate a standard or framework in the table above, unless it is no longer used in the industry and it is replaced with a successor (if any). Azure Government Services meet a separate set of control standards and frameworks, as detailed on the Microsoft Azure Trust Center.

Subject to non-disclosure obligations, Microsoft will make each Information Security Policy available to Purchasing Entity, along with other information reasonably requested by Purchasing Entity regarding Microsoft security practices and policies.

Purchasing Entity is solely responsible for reviewing each Information Security Policy and making an independent determination as to whether it meets Purchasing Entity's requirements.

- C. **Microsoft Audits of DPT Services.** For each DPT Service, Microsoft will conduct audits of the security of the computers, computing environment and physical data centers that it uses in processing Customer Data (including personal data), as follows:

- Where a standard or framework provides for audits, an audit of such control standard or framework will be initiated at least annually for each DPT Service.
- Each audit will be performed according to the standards and rules of the regulatory or accreditation body for each applicable control standard or framework.
- Each audit will be performed by qualified, independent, third party security auditors at Microsoft's selection and expense.

Each audit will result in the generation of an audit report ("Microsoft Audit Report"), which will be Microsoft's Confidential Information. The Microsoft Audit Report will clearly disclose any material findings by the auditor. Microsoft will promptly remediate issues raised in any Microsoft Audit Report to the satisfaction of the auditor.

If Purchasing Entity requests, Microsoft will provide Purchasing Entity with each Microsoft Audit Report so that Purchasing Entity can verify Microsoft's compliance with the security obligations under the DPT. The Microsoft Audit Report will be subject to non-disclosure and distribution limitations of Microsoft and the auditor.

- D. **HIPAA/HITECH Business Associate.** If Purchasing Entity is a "covered entity" or a "business associate" and includes "protected health information" in Customer Data, as those terms are defined in 45 CFR § 160.103, this agreement includes execution of the HIPAA Business Associate Agreement ("BAA"), the full text of which identifies the DPT Services to which it applies and is available at <http://aka.ms/BAA>. Purchasing Entity may opt out of the BAA by instructing Contractor to send the following information to Microsoft in a written notice (under the terms of the Contractor's subcontract with Microsoft):

- the full legal name of Purchasing Entity and any Affiliate that is opting out;
- the volume licensing agreement number to which the opt out applies.

- E. **FedRAMP.** During the term of a Purchasing Entity's subscription for Government Community Cloud Services, those services will be operated in accordance with a written data security policy and control framework that is consistent with the requirements of NIST 800-53 Revision 4, or successor standards and guidelines (if any), established to support Federal Risk and Authorization Management Program (FedRAMP) accreditation at a Moderate Impact level. Microsoft intends for Government Community Cloud Services to support FedRAMP Authority to Operate ("ATO"), and Microsoft will use commercially reasonable efforts to obtain an ATO from a Federal agency, and to

maintain such ATO through continuous monitoring processes and by conducting regular FedRAMP audits.

- F. **Background Checks.** Microsoft performs the following background checks on all US personnel who have potential to access Customer Data. Adherence to this policy is one of the control procedures addressed by the Microsoft Audit Report per the section of the Microsoft Online Services Terms titled "Microsoft Audits of Online Services." Such Background Checks will be performed in accordance with the Fair Credit Reporting Act and will consist of Social Security Number trace, seven (7) year felony and misdemeanor criminal records check of federal, state, or local records (as applicable) for job related crimes, Office of Foreign Assets Control List (OFAC) check, Bureau of Industry and Security List (BIS) check and Office of Defense Trade Controls Debarred Persons List (DDTC) check.

Additionally, for States with which Microsoft has entered into a Microsoft CJIS Information Agreement and provided an FBI CJIS Addendum Certification to the applicable CJIS Systems Agency (CSA, e.g. the California Department of Justice), for its CJIS Covered Services (which are most of the Microsoft Government Community Cloud Services), Microsoft submits its personnel that have access to unencrypted Criminal Justice Information and other Customer Content (including both employees and subcontractors, where applicable) for FBI NCIC fingerprint background checks and adjudication by each such CSA.

- G. **Click-Through Terms.** For the Microsoft In-Scope Services sold and licensed to Purchasing Entities hereunder, no click-through licensing terms presented to end users or administrators, as they pertain to the delivery, operation or use of such services, shall be binding. For clarity, to the extent that certain services may present terms of use for a web portal used to administer and configure the In-Scope Services, or to download software in conjunction with the In-Scope Services, such terms of use shall be binding, to the extent that there are no equivalent terms and conditions in the Master Agreement which pertain to the use of such web portals.

9. **Definitions.**

"AUP" means the section of the OST titled "Acceptable Use Policy" "Azure Consumption-Based Services" means any Microsoft Online Service made available for Purchasing Entity's use and consumption under the Microsoft Azure Administrative Portal, following Purchasing Entity's initial subscription order for such services, subject to consumption fees billed to Purchasing Entity (or debited against its prepaid Azure monetary commitment balance, if applicable) as described in the applicable Product Terms.

"Client" means device software, if any, that Contractor or Microsoft provides to Purchasing Entity with In-Scope Services.

"Customer Content" means the subset of Customer Data created by users, which includes the following:

- For Office 365 Services, Customer Content shall at least include Exchange Online mailbox content (e-mail body, calendar entries, and the content of e-mail attachments), SharePoint Online site content and the files stored within that site, and Skype for Business Online archived conversations.
- For Microsoft Dynamics CRM Online Services, Customer Content means the entities of Customer Data managed by the Microsoft Dynamics CRM Online Services.
- For Microsoft Azure Core Services, Customer Content includes all Customer Data except as may be noted in the Microsoft Azure Trust Center (which Microsoft may update from time to time, but will not add exceptions for existing Microsoft Azure Core Services in general release).
- For Microsoft Intune Online Services, Customer Content includes all Customer Data except as may be noted in the "Data Location" section of the Microsoft Intune Trust Center.

"Customer Data" means all data, including all text, sound, or image files that are provided to Microsoft by, or on behalf of, Purchasing Entity through its use of In-Scope Services.

"DPT" means the section of the OST titled "Data Processing Terms."

“DPT Services” means the Microsoft Online Services to which the DPT applies. As of the date this agreement was executed by the parties, DPT Services include “Microsoft Dynamics CRM Online Services,” “Office 365 Services,” “Microsoft Intune Online Services,” and “Microsoft Azure Core Services,” which consist of the following component Microsoft Online Services:

DPT Services	
Microsoft Dynamics CRM Online Services	Microsoft Dynamics CRM Online services made available through volume licensing or the Microsoft online services portal, excluding (1) Microsoft Dynamics CRM for supported devices, which includes but is not limited to Microsoft Dynamics CRM Online services for tablets and/or smartphones and (2) any separately-branded service made available with or connected to Microsoft Dynamics CRM Online, such as Microsoft Social Engagement, Parature, from Microsoft, and Microsoft Dynamics Marketing.
Office 365 Services	The following services, each as a standalone service or as included in an Office 365-branded plan or suite: Exchange Online, Exchange Online Archiving, Exchange Online Protection, Advanced Threat Protection, SharePoint Online, OneDrive for Business, Project Online, Skype for Business Online, Sway, Office Online, Delve Analytics, Customer Lockbox, and Yammer Enterprise. Office 365 Services do not include Office 365 ProPlus, any portion of PSTN Services that operate outside of Microsoft’s control, any client software, or any separately branded service made available with an Office 365-branded plan or suite, such as a Bing or a service branded “for Office 365.”
Microsoft Azure Core Services	Cloud Services (web and worker roles), Virtual Machines (including with SQL Server), Storage (Blobs, Tables, Queues), Virtual Network, Traffic Manager, Batch, Web Sites, BizTalk Services, Media Services, Mobile Services, Service Bus, Notification Hub, Workflow Manager, Express Route, Scheduler, Multi-Factor Authentication, Active Directory, Rights Management Service, SQL Database, and HDInsight.
Microsoft Intune Online Services	The cloud service portion of Microsoft Intune such as the Microsoft Intune Add-on Product or a management service provided by Microsoft Intune such as Mobile Device Management for Office 365.

“Government Community Cloud Services” means the Microsoft Online Services identified as such (a subset of DPT Services) in the DPT and Product Terms.

“In-Scope Services” means the Microsoft Online Services for which subscription licenses are sold to Purchasing Entity by Contractor. In-Scope Services include, but are not limited to DPT Services. Except for Azure Consumption-Based Services (which are billed based upon the quantities of such services consumed by Purchasing Entity), Purchasing Entity will order individual licenses from Contractor for all In-Scope Services Licenses.

“Master Agreement” means the agreement (described therein as “NASPO ValuePoint Master Agreement for Cloud Services”) between Contractor and the Lead State, subject to the terms and conditions of the Participating Addendum between Contractor and Purchasing Entity’s State (if different from Utah), and which resulted from award of contract under Utah Solicitation Number CH16012.

“Microsoft” means Microsoft Corporation, which is subcontractor to Contractor and delivers the In-Scope Services.

“Online Services” means the Microsoft-hosted services for which Purchasing Entity acquires Licenses from Contractor hereunder.

“OST” means the Microsoft “Online Services Terms,” which are additional terms that apply to Purchasing Entity’s use of Online Services published on the Volume Licensing Site and updated from time to time, and is subject to terms and conditions in the OST which govern which version of OST is applicable to a particular subscription order.

“Portal” means the Online Services Portal for each In-Scope Service.

“Product Terms” means the document that provides information about Microsoft Products and Professional Services available through volume licensing. The Product Terms document is available on the Volume Licensing Site and is updated from time to time.

“Security Incident” means any unlawful access, use, theft or destruction to any Customer Data stored on Microsoft’s equipment or in Microsoft’s facilities, or unauthorized access to such equipment or facilities resulting in use, theft, loss, disclosure, alteration or destruction of Customer Data. All references to “Data Breach” in the Master Agreement shall be deemed to mean Security Incident.

“Service Level Agreement” (“SLA”) means the document which specifies the standards to which Microsoft agrees to adhere and by which it measures the level of service for an In-Scope Service. Microsoft’s current and archived prior version SLAs are available at the Volume Licensing Site. With respect to any In-Scope Service ordered under the this Agreement, the most current SLA available at the onset of a subscription License term shall apply to that In-Scope Service for the duration of that subscription License Term, after which (for any subsequent renewal subscription License term) it will be superseded by the version current at the time of renewal.

“Subscription” means an order for a quantity of an In-Scope Service, including but not limited to (a) orders for User Subscription Licenses for one or more In-Scope Services; and (b) an order for the purpose of establishing a subscription for Azure Consumption-Based Services.

“Volume Licensing Site” means <http://www.microsoft.com/licensing/contracts> or a successor site.

AWS PUBLIC SECTOR ACCESS POLICY (STATE)

(Last Updated January 4, 2013)

This AWS Access Policy ("**Access Policy**") governs your access to and use of the Services (as defined below) of Amazon Web Services, Inc. ("**AWS**") provided to you by your systems integrator, reseller, or services provider ("**Provider**"). It sets out the additional rules, conditions and restrictions that apply to you or the entity you represent ("**you**") for use of the Services. In this Access Policy, "**we**", "**us**", or "**our**" means AWS and any of its affiliates. Please see Section 10 for definitions of capitalized terms.

1. Use of the Services.

1.1 Generally. You are provided access to the Services by your Provider. Your use of and access to the Services are governed by the agreement between you and Provider. This Access Policy supplements the terms of such agreement and may be updated by us from time to time. AWS Service Level Agreements do not apply to your use of the Services. Your continued access to and use of the Services is conditioned on your compliance with all laws, rules, regulations, policies and instructions applicable to your use of the Services, including the Policies.

1.2 Account Keys. Provider may provide you with AWS account keys which will allow you to directly access the Services via Provider's account(s). We are not responsible for any activities that occur under these account keys, regardless of whether the activities are undertaken by you, Provider or a third party (including your employees, contractors or agents) and we are also not responsible for unauthorized access to the account.

1.3 Third Party Materials. Through the use of Provider's AWS account(s), you may have access to Third Party Materials, such as software applications provided by third parties, which are made available directly to you by other companies or individuals under separate terms and conditions, including separate fees and charges. Your use of any Third Party Materials is at your sole risk.

2. Your Responsibilities.

2.1 Your Materials. You are solely responsible for the development, content, operation, maintenance, and use of Your Materials with the Services. For example, you are solely responsible for:

- (a) the technical operation of Your Materials, including ensuring that calls you make to any Service are compatible with then-current application program interfaces for that Service;
- (b) compliance of Your Materials with the Acceptable Use Policy, the other Policies, and the law;
- (c) any claims relating to Your Materials;
- (d) properly handling and processing notices sent to you (or any of your affiliates) by any person claiming that Your Materials violate such person's rights, including notices pursuant to the Digital Millennium Copyright Act;
- (e) any action that you permit, assist or facilitate any person or entity to take related to this Access Policy, Your Materials or use of the Services; and
- (f) End Users' use of Your Materials and the Services and ensuring that End Users comply with your obligations under this Access Policy and that the terms of your agreement with each End User are consistent with this Access Policy.

2.2 Other Security and Backup. You or Provider are solely responsible for properly configuring and using the Services and taking steps to maintain appropriate security, protection and backup of Your Materials, including using encryption technology to protect Your Materials from unauthorized access and routinely archiving Your Materials.

2.3 End User Violations. If you become aware of any violation of your obligations under this Access Policy by an End User, you will immediately terminate such End User's access to Your Materials and the Services.

3. Service Interruption.

3.1 General. We may suspend the AWS account(s) through which you access the Services immediately if we determine your or an End User's use of the Services (i) violates the terms of this Access Policy (including the Acceptable Use Policy or Service Terms); (ii) poses a security risk to the Services or any other AWS customer, (iii) may harm our systems or the systems or Materials of any other AWS customer; or (iv) may subject us to liability as a result of any of the foregoing. We will provide notice of any suspension as soon as practicable to Provider, who is solely responsible for providing any notices to you under your agreement with them.

3.2 Scope of Interruption. To the extent practicable, we will (i) suspend your right to access or use only those instances, data, or portions of the Services that caused the suspension, and (ii) limit the suspension to those accounts that caused the suspension. If commercially feasible, access to the Services will be restored once the conditions or circumstances giving rise to the suspension have been removed or corrected. Nothing in this Section 3 will operate to limit your rights or remedies otherwise available to you against Provider under your agreement with them or applicable law.

4. Proprietary Rights.

4.1 Services. As between you and us, we or our licensors own and reserve all right, title, and interest in and to the Services. You have the right to use the Services solely as a licensee of Provider in accordance with this Access Policy and the agreement between you and Provider. We have no obligation to provide the Service to you under this Access Policy, so you must look exclusively to Provider and your agreement with Provider regarding such obligation. Except as expressly provided in this Section 4, you obtain no rights to the Services, the AWS Materials or any Third Party Materials.

4.2 Materials. As a part of the Services, you may have access to AWS Materials and Third Party Materials, which may be subject to additional terms and conditions (including the Terms of Use and Apache Software License). By using those materials, you are subject to such additional terms. You are solely responsible for securing any necessary approvals for the download and use of such materials.

4.3 Restrictions. Neither you nor any End User may use the Services in any manner or for any purpose other than as expressly permitted by this Access Policy and the agreement between you and Provider. Neither you nor any End User may, or may attempt to, (a) modify, alter, tamper with, repair, or otherwise create derivative works of any software included in the Services (except to the extent software included in the Services are provided to you under a separate license that expressly permits the creation of derivative works), (b) reverse engineer, disassemble, or decompile the software included in the Services or apply any other process or procedure to derive the source code of any software included in the Services, or (c) access or use the Services in a way intended to avoid incurring fees or exceeding usage limits or quotas. All rights and access granted to you with respect to the Services are conditioned on your continued compliance with this Access Policy, and you will immediately discontinue your use of the Services if you cannot comply with this Access Policy. You will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, vendors, business partners, or licensors, any patent infringement or other intellectual property infringement claim regarding any Services or AWS Materials that you have used.

4.4 Suggestions. If you provide any Suggestions to us when using the Services, you hereby grant to AWS and its affiliates a perpetual, irrevocable, non-exclusive, worldwide, royalty-free right and license to reproduce, distribute, make derivative works based upon, publicly display, publicly perform, make, have made, use, sell, offer for sale, and import the Suggestions, including the right to sublicense such rights through multiple tiers, alone or in combination.

4.5 Government Rights. If you are using the Services on behalf of the government and these terms fail to meet the government's needs or are inconsistent in any respect with federal or state law, you will immediately discontinue your use of the Services (including any AWS Materials).

5. Representations and Warranties. You represent and warrant that (a) you and your End Users' use of the Services (including any use by your employees and personnel) will not violate this Access Policy; (b) you or your

licensors own all right, title, and interest in and to Your Materials; (c) Your Materials (including the use, development, design, production, advertising, or marketing of your Materials) or the combination of your Materials with other applications, content or processes, do not and will not violate any applicable laws or infringe or misappropriate any third-party rights; and (d) your use of the Services will not cause harm to any End User.

6. Indemnification. Except to the extent prohibited by applicable state law, you will defend, indemnify, and hold harmless us, our licensors and each of our respective employees, officers, directors, and representatives from and against any claims, damages, losses, liabilities, costs, and expenses (including reasonable attorneys' fees) arising out of or relating to any third party claim concerning: (a) your or any End Users' use of the Services (including use by your employees and personnel); (b) Your Materials or the combination of Your Materials with other applications, content or processes, including any claim involving alleged infringement or misappropriation of third-party rights or the use, development, design, production, advertising or marketing of Your Materials; or (c) a dispute between you and any End User. If your ability to comply with the foregoing provision is limited to any extent by the absence of appropriations or government authorization, you will make good faith efforts to obtain sufficient appropriations or authorization for any liabilities arising under this Section 5.

7. Disclaimers. WE PROVIDE THE SERVICES ON AN "AS IS" BASIS TO PROVIDER. WE AND OUR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND TO YOU, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE REGARDING THE SERVICES OR ANY THIRD PARTY MATERIALS, INCLUDING ANY WARRANTY THAT THE SERVICES OR THIRD PARTY MATERIALS WILL BE UNINTERRUPTED, ERROR FREE OR FREE OF HARMFUL COMPONENTS, OR THAT ANY MATERIALS, INCLUDING YOUR MATERIALS OR THE THIRD PARTY MATERIALS, WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED. EXCEPT TO THE EXTENT PROHIBITED BY LAW, WE AND OUR LICENSORS DISCLAIM ALL WARRANTIES, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR QUIET ENJOYMENT, AND ANY WARRANTIES ARISING OUT OF ANY COURSE OF DEALING OR USAGE OF TRADE.

8. Limitations of Liability. YOU MUST LOOK SOLELY TO PROVIDER AND YOUR AGREEMENT WITH THEM REGARDING ANY CLAIMS OR DAMAGES RELATED TO THE SERVICES. WE AND OUR AFFILIATES OR LICENSORS WILL NOT BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES (INCLUDING DAMAGES FOR LOSS OF PROFITS, GOODWILL, USE, OR DATA), EVEN IF A PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. FURTHER, NEITHER WE NOR ANY OF OUR LICENSORS WILL BE RESPONSIBLE FOR ANY COMPENSATION, REIMBURSEMENT, OR DAMAGES ARISING IN CONNECTION WITH: (A) YOUR INABILITY TO USE THE SERVICES, INCLUDING AS A RESULT OF ANY (I) SUSPENSION OF YOUR USE OF OR ACCESS TO THE SERVICES, (II) OUR DISCONTINUATION OF ANY OR ALL OF THE SERVICES, OR, (III) ANY UNANTICIPATED OR UNSCHEDULED DOWNTIME OF ALL OR A PORTION OF THE SERVICES FOR ANY REASON; OR (B) ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION, DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR MATERIALS OR OTHER DATA THAT YOU OR ANY END USER SUBMITS OR USES IN CONNECTION WITH THE SERVICES (INCLUDING AS A RESULT OF YOUR OR ANY END USERS' ERRORS, ACTS OR OMISSIONS).

9. Miscellaneous.

9.1 Governing Law; Venue. Except to the extent prohibited by applicable state law, the laws of the State of Washington, without reference to conflict of law rules, govern this Access Policy and any dispute of any sort that might arise between you and us. You irrevocably consent to exclusive jurisdiction and venue of the federal courts located in King County, Washington with respect to any dispute arising in connection with the Services or this Access Policy. We may seek injunctive or other relief in any state, federal, or national court of competent jurisdiction for any actual or alleged infringement of our or any third party's intellectual property or other proprietary rights. The United Nations Convention for the International Sale of Goods does not apply.

9.2 Entire Agreement. This Access Policy is the entire agreement between you and us regarding the Services, including any materials. It supersedes all prior or contemporaneous representations, understandings, agreements, or communications between you and us, whether written or verbal, regarding the subject matter of this Access

Policy. If the terms of this document are inconsistent with the terms contained in your agreement with Provider, the terms contained in this document will control. We will not be bound by, and specifically object to, any term, condition or other provision which is different from or in addition to the provisions of this Access Policy (whether or not it would materially alter it) and which is submitted by you in any order, receipt, acceptance, confirmation, correspondence or other document.

9.3 Survival. The following provisions will survive any termination of your use of the Services: Sections 2.1, 4, 5, 6, 7, 8, 9 and 10.

10. Definitions.

“Acceptable Use Policy” means the policy currently available at <http://aws.amazon.com/aup>, as it may be updated by us from time to time.

“AWS Materials” means Materials we make available in connection with the Services or on the AWS Site to allow access to and use of the Services, including WSDLs; Documentation; sample code; software libraries; command line tools; and other related technology. AWS Materials does not include the Services.

“AWS Service Level Agreement” means all service level agreements that we offer with respect to the Services and post on the AWS Site, as they may be updated by us from time to time.

“AWS Site” means <http://aws.amazon.com> and any successor or related site designated by us.

“Documentation” means the developer guides, getting started guides, user guides, quick reference guides, and other technical and operations manuals, instructions and specifications for the Services currently located at <http://aws.amazon.com/documentation>, as such documentation may be updated by us from time to time.

“End User” means any individual or entity that directly or indirectly through another user: (a) accesses or uses Your Materials; or (b) otherwise accesses or uses the Services through you.

“Materials” means software (including machine images), data, text, audio, video, images or other content.

“Policies” means the Acceptable Use Policy, the Terms of Use, the Service Terms, all restrictions described in the AWS Materials and on the AWS Site, and any other policy or terms referenced in or incorporated into this Access Policy.

“Services” means, collectively or individually (as applicable), the web services made commercially available by us to Provider for use under this Access Policy, including (as applicable) those web services described in the Service Terms.

“Service Terms” means the rights and restrictions for particular Services located at <http://aws.amazon.com/serviceterms>, as they may be updated by us from time to time.

“Suggestions” means all suggested improvements to the Services or AWS Materials that you provide to us.

“Terms of Use” means the terms of use located at <http://aws.amazon.com/terms/>, as they may be updated by us from time to time.

“Third Party Materials” means Materials made available to you by any third party on the AWS Site or in conjunction with the Services.

“Your Materials” means Materials you or any End User (a) run on the Services, (b) cause to interface with the Services, or (c) upload to the Services or otherwise transfer, process, use or store in connection with the Services.

Exhibit A

Product Passthrough Terms – Google Apps for Work (for Customers)

Customer has entered into a certain written agreement (the “Agreement”) pursuant to which Customer has purchased the right to access and use the Product. These Product Passthrough Terms set forth the terms and conditions under which Customer may access and use such Product.

1. Product.

- 1.1 Facilities and Data Transfer. All facilities used to store and process Customer Data will adhere to reasonable security standards no less protective than the security standards at facilities where Provider stores and processes its own information of a similar type. Provider has implemented at least industry standard systems and procedures to ensure the security and confidentiality of Customer Data, protect against anticipated threats or hazards to the security or integrity of Customer Data, and protect against unauthorized access to or use of Customer Data. As part of providing the Product, Provider may transfer, store and process Customer Data in the United States or any other country in which Provider or its agents maintain facilities. By using the Product, Customer consents to this transfer, processing and storage of Customer Data.
- 1.2 Modifications.
 - a. To the Product. Provider may make commercially reasonable changes to the Product, from time to time. If Provider makes a material change to the Product, Customer may be informed of such change, provided that Customer has subscribed with Provider to be informed about such change.
 - b. To URL Terms. Provider may make commercially reasonable changes to the URL Terms from time to time. If Provider makes a material change to the URL Terms, Customer may be informed of such change through the Notification Email Address or via the Admin Console. If the change has a material adverse impact on Customer and Customer does not agree to the change, Customer must so notify Partner or Provider via the Help Center within thirty days after receiving notice of the change. Upon such notification, Customer will remain governed by the URL Terms in effect immediately prior to the change until the end of the then-current term for the affected Product. If the affected Product is renewed, they will be renewed under Provider’s then current URL Terms.
- 1.3 Customer Domain Name Ownership. Prior to providing the Product, Provider or Partner may verify that Customer owns or controls the Customer Domain Names. If Customer does not own, or control, the Customer Domain Names, then Provider will have no obligation to provide Customer with the Product.
- 1.4 Federal Information Security Management Act (FISMA). The Provider’s Product known as “Google Apps Core Services” received a FISMA “Authorization to Operate” for a Moderate impact system. Provider will continue to maintain a System Security Plan (SSP) for the Google Apps Core Services, based on NIST 800-53 Rev. 3, or a similarly applicable standard. If Provider does not maintain this SSP as stated, Customer’s sole and exclusive remedy, and Provider’s entire liability, will be Customer’s ability to terminate use of the Product upon thirty days prior written notice.

2. Customer Obligations.

- 2.1 Compliance. Customer will use the Product in accordance with the Acceptable Use Policy. Provider may make new applications, features or functionality for the Product available from time to time, the use of which may be contingent upon Customer’s agreement to additional terms. In addition, Provider will make other Non-Google Apps Products (beyond the Product) available to Customer and its End Users in accordance with the Non-Google Apps Product Terms and the applicable product-specific Provider terms of service. If Customer does not desire to enable any of the Non-Google Apps Products, Customer can enable or disable them at any time through the Admin Console.
- 2.2 Customer Administration of the Product. Customer may specify one or more Administrators through the Admin Console who will have the rights to access Admin Account(s) and to administer the End User Accounts. Customer and Partner are responsible for: (a) maintaining the confidentiality of the password and Admin Account(s); (b) designating those individuals who are authorized to access the Admin Account(s); and (c) ensuring that all activities that occur in connection with the Admin Account(s) comply with these

Product Passthrough Terms. Customer agrees that Provider's responsibilities do not extend to the internal management or administration of the Product for Customer and that Provider is merely a data-processor.

2.3 End User Consent. Customer's Administrators may have the ability to access, monitor, use, or disclose data available to End Users within the End User Accounts. Customer will obtain and maintain all required consents from End Users to allow: (i) Customer's access, monitoring, use and disclosure of this data and Provider providing Customer with the ability to do so, and (ii) Provider to provide the Product.

2.4 Unauthorized Use. Customer will use commercially reasonable efforts to prevent unauthorized use of the Product, and to terminate any unauthorized use. Customer or Reseller will promptly notify Provider of any unauthorized use of, or access to, the Product of which it becomes aware.

2.5 Restrictions on Use. Unless Provider specifically agrees in writing, Customer will not, and will use commercially reasonable efforts to make sure a third party does not: (a) sell, resell, lease, or the functional equivalent, the Product to a third party (unless expressly authorized in these Product Passthrough Terms); (b) attempt to reverse engineer the Product or any component; (c) attempt to create a substitute or similar service through use of, or access to, the Product; (d) use the Product for High Risk Activities; or (e) use the Product to store or transfer any Customer Data that is controlled for export under Export Control Laws. Customer is solely responsible for any applicable compliance with HIPAA.

2.6 Third Party Requests. Customer is responsible for responding to Third Party Requests. Provider will, to the extent allowed by law and by the terms of the Third Party Request: (a) promptly notify Customer of its receipt of a Third Party Request; (b) comply with Customer's reasonable requests regarding its efforts to oppose a Third Party Request; and (c) provide Customer with the information or tools required for Customer to respond to the Third Party Request. Customer will first seek to obtain the information required to respond to the Third Party Request on its own, and will contact Provider only if it cannot reasonably obtain such information.

3. Technical Support Services.

3.1 By Customer. Customer or Partner will, at its own expense, respond to questions and complaints from End Users or third parties relating to Customer's or End Users' use of the Product. Customer or Partner will use commercially reasonable efforts to resolve support issues before escalating them to Provider.

3.2 By Provider. If Customer or Partner cannot resolve a support issue consistent with the above, then Customer or Partner (as applicable based on the agreement between Provider and Partner) may escalate the issue to Provider in accordance with the TSS Guidelines. Provider will provide TSS to Customer or Partner (as applicable) in accordance with the TSS Guidelines.

4. Suspension.

4.1 Of End User Accounts by Provider. If Provider becomes aware of an End User's violation of these Product Passthrough Terms, then Provider may specifically request that Customer Suspend the applicable End User Account. If Customer fails to comply with Provider's request to Suspend an End User Account, then Provider may do so. The duration of any Suspension by Provider will be until the applicable End User has cured the breach which caused the Suspension.

4.2 Emergency Security Issues. Notwithstanding the foregoing, if there is an Emergency Security Issue, then Provider may automatically Suspend the offending use. Suspension will be to the minimum extent and of the minimum duration required to prevent or terminate the Emergency Security Issue. If Provider Suspend an End User Account for any reason without prior notice to Customer, at Customer's request, Provider will provide Customer the reason for the Suspension as soon as is reasonably possible.

5. Confidential Information.

5.1 Obligations. Customer and Provider will: (a) protect the other's Confidential Information with the same standard of care it uses to protect its own Confidential Information; and (b) not disclose the Confidential Information, except to Affiliates, employees and agents who need to know it and who have agreed in writing to keep it confidential. Customer and Provider (and any Affiliates, employees and agents to whom it has disclosed Confidential Information) may use Confidential Information only to exercise rights and fulfill its obligations under these Product Passthrough Terms, while using reasonable care to protect it. Customer

and Provider, respectively, will be responsible for any actions of its Affiliates, employees and agents in violation of this Section.

5.2 Exceptions. Confidential Information does not include information that: (a) the recipient of the Confidential Information already knew; (b) becomes public through no fault of the recipient; (c) was independently developed by the recipient; or (d) was rightfully given to the recipient by a third party.

5.3 Required Disclosure. Customer and Provider may disclose the other's Confidential Information when required by law but only after it, if legally permissible: (a) uses commercially reasonable efforts to notify the other; and (b) gives the other the chance to challenge the disclosure.

6. Intellectual Property Rights; Brand Features.

6.1 Intellectual Property Rights. Except as expressly set forth herein, these Product Passthrough Terms does not grant either Customer or Provider any rights, implied or otherwise, to the other's content or any of the other's intellectual property. As between the Customer and Provider, Customer owns all Intellectual Property Rights in Customer Data, and Provider owns all Intellectual Property Rights in the Product.

6.2 Display of Brand Features. Provider may display those Customer Brand Features authorized by Customer (such authorization is provided by Customer uploading its Brand Features into the Product), and within designated areas of the Service Pages. Customer may specify the nature of this use using the Admin Console. Provider may also display Provider Brand Features on the Service Pages to indicate that the Product are provided by Provider. Neither party may display or use the other party's Brand Features beyond what is allowed in these Product Passthrough Terms without the other's prior written consent.

6.3 Brand Features Limitation. Any use of Brand Features will inure to the benefit of the entity holding Intellectual Property Rights in those Brand Features. Any right to use its Brand Features pursuant to these Product Passthrough Terms may be revoked at any time with written notice to the other and a reasonable period to stop the use.

7. Publicity. Customer agrees that Provider may include Customer's name or Brand Features in a list of Provider customers, online or in promotional materials. Customer also agrees that Provider may verbally reference Customer as a customer of the Provider's Product that is the subject of these Product Passthrough Terms.

8. Representations, Warranties and Disclaimers.

8.1 Representations and Warranties. Customer and Provider represent that each will comply with all laws and regulations applicable to its provision, or use, of the Product, as applicable (including applicable security breach notification law). Provider warrants that it will provide the Product in accordance with the applicable SLA.

8.2 Disclaimers. TO THE FULLEST EXTENT PERMITTED BY APPLICABLE LAW, EXCEPT AS EXPRESSLY PROVIDED FOR HEREIN, NEITHER PARTY MAKES ANY OTHER WARRANTY OF ANY KIND, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING WITHOUT LIMITATION WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR USE AND NONINFRINGEMENT. PROVIDER MAKES NO REPRESENTATIONS ABOUT ANY CONTENT OR INFORMATION MADE ACCESSIBLE BY OR THROUGH THE PRODUCT. CUSTOMER ACKNOWLEDGES THAT THE PRODUCT IS NOT A TELEPHONY SERVICE AND THAT THE PRODUCT IS NOT CAPABLE OF PLACING OR RECEIVING ANY CALLS, INCLUDING EMERGENCY PRODUCT CALLS, OVER PUBLICLY SWITCHED TELEPHONE NETWORKS.

9. Term and Termination.

9.1 Term. The term for the Product will be as decided upon between Partner and Customer. These Product Passthrough Terms will remain in effect for the Term.

9.2 Termination for Breach. The Services may be suspended or terminated, if: (i) Customer is in material breach of these Product Passthrough Terms and fails to cure that breach within thirty days after receipt of written notice; (ii) Customer ceases its business operations or becomes subject to insolvency proceedings and the proceedings are not dismissed within ninety days; or (iii) Customer is in material breach of these Product Passthrough Terms more than two times notwithstanding any cure of such breaches.

9.3 Effects of Termination. If Services are terminated, then: (i) the rights granted by Provider to Customer, and Customer to Provider will cease immediately (except as set forth in this Section); (ii) neither Customer nor Partner will have access to, or the ability to export, the Customer Data; (iii) Provider will begin to delete Customer Data; and (iv) upon request will promptly use commercially reasonable efforts to return or destroy all other Confidential Information of the other.

10. Miscellaneous.

10.1 Product Development. The Product was developed solely at private expense and is commercial computer software and related documentation within the meaning of the applicable civilian and military federal acquisition regulations and any supplements thereto.

10.2 Force Majeure. Due to circumstances beyond Provider's controls, Provider may not be able to provide the Product.

10.3 No Waiver. Failure to enforce any provision of these Product Passthrough Terms will not constitute a waiver.

10.4 No Agency. These Product Passthrough Terms do not create any agency, partnership or joint venture.

10.5 No Third-Party Beneficiaries. There are no third-party beneficiaries to these Product Passthrough Terms. Customer's sole remedies will be set forth in its agreement with the Partner.

10.6 Survival. The following sections will survive expiration or termination of these Product Passthrough Terms: Sections 5, 6, 9.3, 10 and 12.

10.7 Severability. If any term (or part of a term) of these Product Passthrough Terms is invalid, illegal, or unenforceable, the rest of these Product Passthrough Terms will remain in effect.

10.8 Conflicting Terms. If there is a conflict between any terms of these Product Passthrough Terms and any other document that makes up the entire agreement as needed for Customer to use the Product, the terms of these Product Passthrough Terms will take precedence.

11. Additional Product Terms.

11.1 Ads. Provider does not serve Ads in the Product or use Customer Data for Ads purposes.

11.2 Aliases. Customer is solely responsible for monitoring, responding to, and otherwise processing emails sent to the "abuse" and "postmaster" aliases for Customer Domain Names but Provider may monitor emails sent to these aliases for Customer Domain Names to allow Provider to identify Product abuse.

11.3 Google Apps Vault Retention. If Customer is using Google Apps Vault, Provider will have no obligation to retain any archived Customer Data beyond the retention period specified by Customer. If Customer does not renew Google Apps Vault, Provider will have no obligation to retain any archived Customer Data.

12. Definitions.

"Acceptable Use Policy" means the acceptable use policy for the Product available at http://www.google.com/a/help/intl/en/admins/use_policy.html. The Acceptable Use Policy and such URL link may be updated or modified by Provider from time to time.

"Admin Account(s)" means the administrative account(s) provided to Customer for the purpose of administering the Product. The use of the Admin Account(s) requires a password, which Provider will provide to Customer or Partner.

"Admin Console" means the online tool Customer may use in reporting and certain other administration functions.

"Administrators" mean the Customer-designated technical personnel who administer the Product to End Users on Customer's behalf.

“Ads” means online advertisements displayed by Provider to End Users, excluding advertisements provided by any advertising products that are not part of the Product that Customer chooses to use in connection with the Product.

“Affiliate” means any entity that directly or indirectly controls, is controlled by, or is under common control with an entity.

“Brand Features” means the trade names, trademarks, service marks, logos, domain names, and other distinctive brand features of each party, respectively, as secured from time to time.

“Confidential Information” means information disclosed or exchanged under these Product Passthrough Terms that is marked as confidential or would normally be considered confidential under the circumstances. Customer Data is Customer’s Confidential Information.

“Customer Data” means data, including email, provided, generated, transmitted or displayed via the Product by Customer, End Users, or Partner on behalf of Customer.

“Customer Domain Names” mean the domain names owned or controlled by Customer, which will be used in connection with the Product.

“Emergency Security Issue” means either: (a) Customer’s use of the Product in violation of the Acceptable Use Policy, which could disrupt: (i) the Product; (ii) other customers’ use of the Product; or (iii) the Provider network or servers used to provide the Product; or (b) unauthorized third party access to the Product.

“End Users” means the individuals Customer permits to use the Product.

“End User Account” means a Provider-hosted account established by Customer through the Product for an End User.

“Export Control Laws” means all applicable export and reexport control laws and regulations, including trade and economic sanctions maintained by the Treasury Department’s Office of Foreign Assets Control, and the International Traffic in Arms Regulations (“ITAR”) maintained by the Department of State.

“Google Apps Core Services” means the applicable Product purchased by Customer from Partner which are more fully described here: http://www.google.com/a/help/intl/en/users/user_features.html. The Google Apps Core Services and such URL link may be updated or modified by Provider from time to time.

“Help Center” means the Provider help center accessible at <http://www.google.com/support/>. The Help Center and such URL link may be updated or modified by Provider from time to time.

“High Risk Activities” means uses such as the operation of nuclear facilities, air traffic control, or life support systems, where the use or failure of the Product could lead to death, personal injury, or environmental damage.

“HIPAA” means the Health Insurance Portability and Accountability Act of 1996, as may be amended from time to time, and any regulations issued thereunder.

“Intellectual Property Rights” means current and future worldwide rights under patent law, copyright law, trade secret law, trademark law, moral rights law, and other similar rights.

“Non-Google Apps Products” means Provider products which are not part of the Product, but which may be accessed by End Users using their End User Account login and password. The Non-Google Apps Products are set forth at the following URL: <http://www.google.com/support/a/bin/answer.py?hl=en&answer=181865>. The Non-Google Apps Products and such URL link may be updated or modified by Provider from time to time.

“Non-Google Apps Product Terms” means the terms found at the following URL: http://www.google.com/apps/intl/en/terms/additional_services.html. The Non-Google Apps Product Terms and such URL link may be updated or modified by Provider from time to time.

“Notification Email Address” means the email address designated by Customer to receive email notifications from Provider. Customer may provide a Partner email address for this purpose if it so chooses. Customer may change this email address through the Admin Console.

“Provider” means the third party provider of the Product.

“Partner” means the entity Customer is paying to provide access to and use of the Product.

“Service Pages” mean the web pages displaying the Product to End Users.

“Product” means, as applicable, the Google Apps Core Services purchased from Partner.

“SLA” means the Service Level Agreement located here for applicable Google Apps Core Services: http://www.google.com/apps/intl/en/terms/reseller_sla.html. The SLA and such URL link may be updated or modified by Provider from time to time.

“Suspend” means the immediate disabling of access to the Product, or components of the Product, as applicable, to prevent further use of the Product.

“Term” means the term of these Product Passthrough Terms, which will continue for as long as Customer is receiving Product from Provider, unless terminated earlier pursuant to these Product Passthrough Terms, or pursuant to Customer’s agreement with Partner.

“Third Party Request” means a request from a third party for records relating to an End User’s use of the Product. Third Party Requests can be a lawful search warrant, court order, subpoena, other valid legal order, or written consent from the End User permitting the disclosure.

“TSS” means the technical support Product provided by Provider to the Administrators during the Term pursuant to the TSS Guidelines.

“TSS Guidelines” means Provider’s technical support Product guidelines then in effect for the Product. TSS Guidelines are at the following URL: <http://www.google.com/a/help/intl/en/admins/tssg.html>. The TSS Guidelines and such URL link may be updated or modified by Provider from time to time.

“URL Terms” means the Acceptable Use Policy, the SLA and the TSS Guidelines.