Contract # AR2480

STATE OF UTAH COOPERATIVE CONTRACT

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

DLT Solutions		
Name		
2411 Dulles Corner Park, Suite 800		
Address		
Herndon	VA	20171
City	State	Zip

LEGAL STATUS OF CONTRACTOR
- ☐ Sole Proprietor
- ☐ Non-Profit Corporation
- ☒ For-Profit Corporation
- ☐ Partnership
- ☐ Government Agency

Contact Person Duane Thorpe Phone #703-773-1163 Email duane.thorpe@dlt.com
Vendor #_____ 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 the 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: Contractor's Terms and Conditions

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

Contractor's signature — Date: 12/8/16

ALAN M SMITH PRESIDENT & CEO
Type or Print Name and Title

STATE

Director, Division of Purchasing — Date: 12.30.16

Christopher Hughes	801-538-3254		christopherhughes@utah.gov
Division of Purchasing Contact Person	Telephone Number	Fax Number	Email

(Revision 16 June 2016)

APPROVED

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 Addendum[1] ("PA");
(2) NASPO ValuePoint Master Agreement Terms & Conditions, including the applicable Exhibits[2] to the Master Agreement;
(3) The Solicitation;
(4) Contractor's response to the Solicitation, as revised (if permitted) and accepted by the Lead State; and
(5) 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 obtained by Contractor 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, and (3) information concerning individuals, is confidential information of Purchasing Entity.

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.

Data means all information, whether in oral or written (including electronic) form,

[1] A Sample Participating Addendum will be published after the contracts have been awarded.
[2] The Exhibits comprise the terms and conditions for the service models: PaaS, IaaS, and PaaS.

created by or in any way originating with a Participating Entity or Purchasing Entity, and all information 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 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 consumer 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 consumer 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 consumer-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 consumer 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, software, and any incidental tangible goods.

Protected Health Information (PHI) means individually identifiable health information transmitted by electronic media, maintained in electronic media, or transmitted or maintained in any other form or medium. 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 possible or 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 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 Contractor's applications 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 discounts 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. A price or rate reduction will apply automatically to the Master Agreement and an amendment is not necessary.

7. Termination: Unless otherwise stated, this Master Agreement may be terminated by either party upon 60 days written notice prior to the effective date of the termination. Further, any Participating Entity may terminate its participation upon 30 days written notice, unless otherwise limited or stated in the Participating Addendum. Termination may be in whole or in part. 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.

8. Confidentiality, Non-Disclosure, and Injunctive Relief

a. Confidentiality. Contractor 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 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 Contractor 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 Contractor) publicly known; (2) is furnished by Purchasing Entity to others without restrictions similar to those imposed by this Master Agreement; (3) is rightfully in Contractor'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 Purchasing Entity without the obligation of confidentiality, (5) is disclosed with the written consent of Purchasing Entity or; (6) is independently developed by employees, agents or subcontractors of Contractor who can be shown to have had no access to the Confidential Information.

b. Non-Disclosure. Contractor 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. 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 person 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 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.

d. Purchasing Entity Law. These provisions shall be applicable only to extent they are not in conflict with the applicable public disclosure laws of any Purchasing Entity.

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

(4) Institution of proceedings under any bankruptcy, insolvency, reorganization or similar law, by or against Contractor, or the appointment of a receiver or similar officer for Contractor or any of its property, which is not vacated or fully stayed within thirty (30) calendar days after the institution or occurrence thereof; or

(5) Any default specified in another section of this Master Agreement.

b. Upon the occurrence of an event of default, Lead State shall issue a written notice of default, identifying the nature of the default, and providing a period of 30 calendar days in which Contractor 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 this Master Agreement in whole or in part 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 Contractor is afforded an opportunity to cure and fails to cure the default within the period specified in the written notice of default, Contractor shall be in breach of its obligations under this Master Agreement and Lead State 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 any related Contracts or portions thereof; and

(3) Suspend Contractor from being able to respond to future bid solicitations; and

(4) Suspend Contractor's performance; and

(5) Withhold payment 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. Nothing in these Master Agreement Terms and Conditions shall be construed to limit the rights and remedies available to a Purchasing Entity under the applicable commercial code.

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 in accordance with 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, 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, indemnify and hold harmless 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, from and against any and all third party claims, direct damages or causes of action including reasonable attorneys' fees and related costs for any death, injury, or damage to property arising directly from act(s), error(s), or omission(s) of the Contractor, its employees or subcontractors or volunteers, at any tier, relating to the performance under the Master Agreement.

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 claims, damages or causes of action including reasonable attorneys' fees and related costs arising out of the claim that the Product or its use, infringes Intellectual Property rights ("Intellectual Property Claim") of another person or entity.

(1) The Contractor's obligations under this section shall not extend to any claims arising from the combination of the Product with any other product, system or method, unless the Product, system or method is:

(a) provided by the Contractor or the Contractor's subsidiaries or affiliates;

(b) specified by the Contractor to work with the Product; or

(c) reasonably required, in order to use the Product in its intended manner, and the infringement could not have been avoided by substituting another reasonably available product, system or method capable of performing the same function; or

(d) It would be reasonably expected to use the Product in combination with such product, system or method.

(2) The Indemnified Party shall notify the Contractor within a reasonable time 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. Contractor shall have control over the defense and settlement of it. 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.

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, with no deductible for each of the following categories:

(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; Contractor will meet these requirements with a combination of primary and umbrella liability limits.

(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

Lead State acknowledges and accepts Contractor has limits of $3,000,000 each claim / policy aggregate for technology errors and Omissions, and Cyber with a $50,000 deductible paid by Contractor.

(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 an occurrence 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 to the Lead State that (1) Contractor has additional insured endorsement that automatically includes those entities requiring to be additional insureds when required by written contract and does not specifically name additional insured, (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, provided Contractor has been given such written notice by its insurance carrier, and (3) provides that the Contractor's liability insurance policy shall be primary, 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); a description of coverage in detailed standard terminology (including policy period, policy number, limits of liability, exclusions and endorsements); and an acknowledgment of the requirement for notice of cancellation provided Contractor has been given such notice by its insurance carrier. 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: Any and all Services offered and furnished shall comply fully with all applicable Federal and State laws and regulations.

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 Office[3].

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.

[3] 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.

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 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: Purchasing Entity understands and acknowledges only the Purchasing Entity controls access to Purchasing Entity data. Purchasing Entity further acknowledges Contractor has no access to Purchasing Entity's data, unless Purchasing Entity grants Contractor access in order to fulfill an obligation 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: Contractor shall maintain the administrative, physical, technical, and procedural infrastructure associated with the provision of the Product in a manner that is, at all times during the term of this Master Agreement, at a level equal to or more stringent than those specified in the Solicitation.

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 and any Data derived therefrom, including metadata.
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.

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 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. 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.

d. The Contractor shall allow the Purchasing Entity to audit conformance to the Master Agreement and applicable Participating Addendum terms. The purchasing entity may perform this audit or contract with a third party at its discretion and at the purchasing entity’s expense.

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. System Failure or Damage: In the event of system failure or damage caused by Contractor or its Services, the Contractor agrees to use its best efforts to restore or assist in restoring the system to operational capacity.

29. Title to Product: 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.

30. Data Privacy: The Contractor must comply with all applicable laws related to data privacy and security, including IRS Pub 1075. 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. Contractor represents and warrants that the representations contained in its response to the Solicitation by the Lead State.

d. The Contractor will not interfere with a Purchasing Entity’s access to and use of the Services it acquires from this Master Agreement.

e. The Services provided by the Contractor are compatible with and will operate successfully with any environment (including web browser and operating system) specified by the Contractor in its response to the Solicitation by the Lead State.

f. The Products or services Contractor provides under this Master Agreement are

provided on an "as is" basis. Except for the express warranty that the Products provided under this Master Agreement are free of malware, Contractor makes no representations or warranties of any kind, whether express or implied, statutory or otherwise regarding the services or that any third party materials will be uninterrupted, error free or free of harmful components, or that any materials, including Participating Entities materials or third party materials, will be secure or not otherwise lost or damaged. Except to the extent prohibited by law, Contractor disclaims 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. The Contractor warrants that is will use appropriate industry-leading technology to detect and remove worms, Trojans, rootkits, rogues, dialers, spyware, etc.

32. Transition Assistance:

a. 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 usable without the use of the Services and as agreed by a Purchasing Entity, at no additional cost to the Purchasing Entity. Any transition services requested by a Purchasing Entity involving additional knowledge transfer and support may be subject to an additional cost and 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, including a Statement of Work if applicable, and subject to additional costs. 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, or Purchasing Entity 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, or Purchasing Entity must be in writing. Waiver by the Lead State or Participating Entity 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. The Contractor acknowledges and agrees that it has conducted its own due diligence prior to entering into this Master Agreement as to all the foregoing matters.

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 support, facility space, materials, special access, personnel 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 in 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. No click-through, or other end user terms and conditions or agreements required by the Contractor ("Additional Terms") provided with any Services hereunder shall be binding on Participating Entities or Purchasing Entities, even if use of such Services requires an affirmative "acceptance" of those Additional Terms before access is permitted.

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 or (ii) one million dollars ($1,000,000), whichever is greater.

b. CONTRACTOR SHALL HAVE NO LIABILITY TO LEAD STATE OR PARTICIPATING STATES OR PARTICIPATING ENTITIES FOR ANY INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING LOST PROFITS, REVENUE, DATA OR DATA USE), REGARDLESS OF THE FORM OF ACTION AND INCLUDING CONTRACT, TORT, STRICT LIABILITY OR OTHERWISE EVEN IF SUCH PARTY HAS BEEN ADVISED OF THE POSSIBILITY IN ADVANCE.

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 limitations of liability in Section 44 will not apply to claims for bodily injury or death, Section 8, Section 13(b), and Section 30.

Exhibit 1 to the Master Agreement: Software-as-a-Service

1. Data Ownership: The Purchasing Entity will own all right, title and interest in its data that is related to the Services provided by this Master Agreement. The Contractor shall not access Purchasing Entity user accounts or Purchasing Entity data, except (1) in the course of data center operations, (2) in response to service or technical issues, (3) as required by the express terms of this Master Agreement, Participating Addendum, SLA, and/or other contract documents, or (4) at the Purchasing Entity's written request.

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 a 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. This obligation shall survive and extend beyond the term of this Master Agreement.

2. Data Protection: Protection of personal privacy and data shall be an integral part of the business activities of the Contractor to ensure there is no inappropriate or unauthorized use of Purchasing Entity information at any time. To this end, the Contractor shall safeguard the confidentiality, integrity and availability of Purchasing Entity information and comply with the following conditions:

a. The Contractor shall implement and maintain appropriate administrative, technical and organizational security measures to safeguard against unauthorized access, disclosure or theft of Personal Data and Non-Public Data. Such security measures shall be in accordance with recognized industry practice and not less stringent than the measures the Contractor applies to its own Personal Data and Non-Public Data of similar kind.

b. All data obtained by the Contractor in the performance of the Master Agreement shall become and remain the property of the Purchasing Entity.

c. All Personal Data shall be encrypted at rest and in transit with controlled access. Unless otherwise stipulated, the Contractor is responsible for encryption of the Personal Data. Any stipulation of responsibilities will identify specific roles and responsibilities and shall be included in the service level agreement (SLA), or otherwise made a part of the Master Agreement.

d. Unless otherwise stipulated, the Contractor shall encrypt all Non-Public Data at rest and in transit. The Purchasing Entity shall identify data it deems as Non-Public Data to the Contractor. The level of protection and encryption for all Non-Public Data shall be identified in the SLA.

e. At no time shall any data or processes — that either belong to or are intended for the use of a Purchasing Entity or its officers, agents or employees — be copied, disclosed or retained by the Contractor or any party related to the Contractor for subsequent use in any transaction that does not include the Purchasing Entity.

f. The Contractor shall not use any information collected in connection with the Services issued from this Master Agreement for any purpose other than fulfilling the Services.

3. Data Location: The Contractor shall provide its services to the Purchasing Entity and its end users solely from data centers in the U.S. Storage of Purchasing Entity data at rest shall be located solely in data centers in the U.S. The Contractor shall not allow its personnel or contractors to store Purchasing Entity data on portable devices, including personal computers, except for devices that are used and kept only at its U.S. data centers. The Contractor shall permit its personnel and contractors to access Purchasing Entity data remotely only as required to provide technical support. The Contractor may provide technical user support on a 24/7 basis using a Follow the Sun model, unless otherwise prohibited in a Participating Addendum.

4. Security Incident or Data Breach Notification:

a. Incident Response: Contractor may need to communicate with outside parties regarding a security incident, which may include contacting law enforcement, fielding media inquiries and seeking external expertise as mutually agreed upon, defined by law or contained in the contract. Discussing security incidents with the Purchasing Entity should be handled on an urgent as-needed basis, as part of Contractor's communication and mitigation processes as mutually agreed upon, defined by law or contained in the Master Agreement.

b. Security Incident Reporting Requirements: The Contractor shall report a security incident to the Purchasing Entity identified contact immediately as soon as possible or promptly without out reasonable delay, or as defined in the SLA.

c. Breach Reporting Requirements: If the Contractor has actual knowledge of a confirmed data breach that affects the security of any purchasing entity's content that is subject to applicable data breach notification law, the Contractor shall (1) as soon as possible or promptly without out reasonable delay notify the Purchasing Entity, unless shorter time is required by applicable law, and (2) take commercially reasonable measures to address the data breach in a timely manner.

5. Personal Data Breach Responsibilities: This section only applies when a Data Breach occurs with respect to Personal Data within the possession or control of the Contractor.

a. The Contractor, unless stipulated otherwise, shall immediately notify the appropriate Purchasing Entity identified contact by telephone in accordance with the agreed upon security plan or security procedures if it reasonably believes there has been a security incident.

b. The Contractor, unless stipulated otherwise, shall promptly notify the appropriate Purchasing Entity identified contact within 24 hours or sooner by telephone, unless shorter time is required by applicable law, if it has confirmed that there is, or reasonably believes that there has been a Data Breach. The Contractor shall (1) cooperate with the Purchasing Entity as reasonably requested by the Purchasing Entity to investigate and resolve the Data Breach, (2) promptly implement necessary remedial measures, if necessary, and (3) document responsive actions taken related to the Data Breach, including any post-incident review of events and actions taken to make changes in business practices in providing the services, if necessary.

c. Unless otherwise stipulated, if a data breach is a direct result of Contractor's breach of its contractual obligation to encrypt personal data or otherwise prevent its release as reasonably determined by the Purchasing Entity, the Contractor shall bear the costs associated with (1) the investigation and resolution of the data breach; (2) notifications to individuals, regulators or others required by federal and state laws or as otherwise agreed to; (3) a credit monitoring service required by state (or federal) law or as otherwise agreed to; (4) a website or a toll-free number and call center for affected individuals required by federal and state laws — all not to exceed the average per record per person cost calculated for data breaches in the United States (currently $217 per record/person) in the most recent Cost of Data Breach Study: Global Analysis published by the Ponemon Institute at the time of the data breach; and (5) complete all corrective actions as reasonably determined by Contractor based on root cause.

6. Notification of Legal Requests: The Contractor shall contact the Purchasing Entity upon receipt of any electronic discovery, litigation holds, discovery searches and expert testimonies related to the Purchasing Entity's data under the Master Agreement, or which in any way might reasonably require access to the data of the Purchasing Entity. The Contractor shall not respond to subpoenas, service of process and other legal requests related to the Purchasing Entity without first notifying and obtaining the approval of the Purchasing Entity, unless prohibited by law from providing such notice.

7. Termination and Suspension of Service:

a. In the event of a termination of the Master Agreement or applicable Participating Addendum, the Contractor shall implement an orderly return of purchasing entity's data in a CSV or another mutually agreeable format at a time agreed to by the parties or allow the Purchasing Entity to extract it's data and the subsequent secure disposal of purchasing entity's data.

b. During any period of service suspension, the Contractor shall not take any action to intentionally erase or otherwise dispose of any of the Purchasing Entity's data.

c. In the event of termination of any services or agreement in entirety, the Contractor shall not take any action to intentionally erase purchasing entity's data for a period of:

- 10 days after the effective date of termination, if the termination is in accordance with the contract period
- 30 days after the effective date of termination, if the termination is for convenience
- 60 days after the effective date of termination, if the termination is for cause

After such period, the Contractor shall have no obligation to maintain or provide any purchasing entity's data and shall thereafter, unless legally prohibited, delete all purchasing entity's data in its systems or otherwise in its possession or under its control.

d. The purchasing entity shall be entitled to any post termination assistance generally made available with respect to the services, unless a unique data retrieval arrangement has been established as part of an SLA.

e. Upon termination of the Services or the Agreement in its entirety, Contractor shall securely dispose of all Purchasing Entity's data in all of its forms, such as disk, CD/ DVD, backup tape and paper, unless stipulated otherwise by the Purchasing Entity. Data shall be permanently deleted and shall not be recoverable, according to National Institute of Standards and Technology (NIST)-approved methods. Certificates of destruction shall be provided to the Purchasing Entity.

8. Background Checks: Upon the request of the Purchasing Entity, the Contractor shall conduct criminal background checks and not utilize any staff, including subcontractors, to fulfill the obligations of the Master Agreement who have been convicted of any crime of dishonesty, including but not limited to criminal fraud, or otherwise convicted of any felony or misdemeanor offense for which incarceration for up to 1 year is an authorized penalty. The Contractor shall promote and maintain an awareness of the importance of securing the Purchasing Entity's information among the Contractor's employees and agents. If any of the stated personnel providing services under a Participating Addendum is not acceptable to the Purchasing Entity in its sole opinion as a result of the background or criminal history investigation, the Purchasing Entity, in its' sole option shall have the right to either (1) request immediate replacement of the person, or (2) immediately terminate the Participating Addendum and any related service agreement.

9. Access to Security Logs and Reports: The Contractor shall provide reports on a schedule specified in the SLA to the Purchasing Entity in a format as specified in the SLA agreed to by both the Contractor and the Purchasing Entity. Reports shall include latency statistics, user access, user access IP address, user access history and security logs for all public jurisdiction files related to this Master Agreement and applicable Participating Addendum.

10. Contract Audit: The Contractor shall allow the Purchasing Entity to audit conformance to the Master Agreement terms. The Purchasing Entity may perform this audit or contract with a third party at its discretion and at the Purchasing Entity's expense.

11. Data Center Audit: The Contractor shall perform an independent audit of its data centers at least annually at its expense, and provide an unredacted version of the audit report upon request to a Purchasing Entity. The Contractor may remove its proprietary information from the unredacted version. A Service Organization Control (SOC) 2 audit report or approved equivalent sets the minimum level of a third-party audit.

12. Change Control and Advance Notice: The Contractor shall give a minimum forty eight (48) hour advance notice (or as determined by a Purchasing Entity and included in the SLA) to the Purchasing Entity of any upgrades (e.g., major upgrades, minor upgrades, system changes) that may impact service availability and performance. A major upgrade is a replacement of hardware, software or firmware with a newer or better version in order to bring the system up to date or to improve its characteristics. It usually includes a new version number.

Contractor will make updates and upgrades available to Purchasing Entity at no additional costs when Contractor makes such updates and upgrades generally available to its users.

No update, upgrade or other charge to the Service may decrease the Service's functionality, adversely affect Purchasing Entity's use of or access to the Service, or increase the cost of the Service to the Purchasing Entity.

Contractor will notify the Purchasing Entity at least sixty (60) days in advance prior to any major update or upgrade.

13. Security: As requested by a Purchasing Entity, the Contractor shall disclose its non-proprietary system security plans (SSP) or security processes and technical limitations to the Purchasing Entity such that adequate protection and flexibility can be attained between the Purchasing Entity and the Contractor. For example: virus checking and port sniffing — the Purchasing Entity and the Contractor shall understand each other's roles and responsibilities.

14. Non-disclosure and Separation of Duties: The Contractor shall enforce separation of job duties, require commercially reasonable non-disclosure agreements, and limit staff knowledge of Purchasing Entity data to that which is absolutely necessary to perform job duties.

15. Import and Export of Data: The Purchasing Entity shall have the ability to import or export data in piecemeal or in entirety at its discretion without interference from the Contractor at any time during the term of Contractor's contract with the Purchasing Entity. This includes the ability for the Purchasing Entity to import or export data to/from other Contractors. Contractor shall specify if Purchasing Entity is required to provide its' own tools for this purpose, including the optional purchase of Contractors tools if Contractors applications are not able to provide this functionality directly.

16. Responsibilities and Uptime Guarantee: The Contractor shall be responsible for the acquisition and operation of all hardware, software and network support related to the services being provided. The technical and professional activities required for establishing, managing and maintaining the environments are the responsibilities of the Contractor. The system shall be available 24/7/365 (with agreed-upon maintenance downtime), and provide service to customers as defined in the SLA.

17. Subcontractor Disclosure: Contractor shall identify all of its strategic business partners related to services provided under this Master Agreement, including but not limited to all subcontractors or other entities or individuals who may be a party to a joint venture or similar agreement with the Contractor, and who shall be involved in any application development and/or operations.

18. Right to Remove Individuals: The Purchasing Entity shall have the right at any time to require that the Contractor remove from interaction with Purchasing Entity any Contractor representative who the Purchasing Entity believes is detrimental to its working relationship with the Contractor. The Purchasing Entity shall provide the Contractor with notice of its determination, and the reasons it requests the removal. If the Purchasing Entity signifies that a potential security violation exists with respect to the request, the Contractor shall immediately remove such individual. The Contractor shall not assign the

person to any aspect of the Master Agreement or future work orders without the Purchasing Entity's consent.

19. Business Continuity and Disaster Recovery: The Contractor shall provide a business continuity and disaster recovery plan upon request and ensure that the Purchasing Entity's recovery time objective (RTO) of XXX hours/days is met. (XXX hour/days shall be provided to Contractor by the Purchasing Entity.) Contractor must work with the Purchasing Entity to perform an annual Disaster Recovery test and take action to correct any issues detected during the test in a time frame mutually agreed between the Contractor and the Purchasing Entity.

20. Compliance with Accessibility Standards: The Contractor shall comply with and adhere to Accessibility Standards of Section 508 Amendment to the Rehabilitation Act of 1973, or any other state laws or administrative regulations identified by the Participating Entity.

21. Web Services: The Contractor shall use Web services exclusively to interface with the Purchasing Entity's data in near real time.

22. Encryption of Data at Rest: The Contractor shall ensure hard drive encryption consistent with validated cryptography standards as referenced in FIPS 140-2, Security Requirements for Cryptographic Modules for all Personal Data, unless the Purchasing Entity approves in writing for the storage of Personal Data on a Contractor portable device in order to accomplish work as defined in the statement of work.

23. Subscription Terms: Contractor grants to a Purchasing Entity a license to: (i) access and use the Service for its business purposes; (ii) for SaaS, use underlying software as embodied or used in the Service; and (iii) view, copy, upload and download (where applicable), and use Contractor's documentation.

No Contractor terms, including standard click through license or website terms or use of privacy policy, shall apply to Purchasing Entities unless such terms are included in this Master Agreement.

Exhibit 2 to the Master Agreement: Platform-as-a-Service

1. Data Ownership: The Purchasing Entity will own all right, title and interest in its data that is related to the Services provided by this Master Agreement. The Contractor shall not access Purchasing Entity user accounts or Purchasing Entity data, except (1) in the course of data center operations, (2) in response to service or technical issues, (3) as required by the express terms of this Master Agreement, Participating Addendum, SLA, and/or other contract documents, or (4) at the Purchasing Entity's written request.

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 a 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. This obligation shall survive and extend beyond the term of this Master Agreement.

2. Data Protection: Protection of personal privacy and data shall be an integral part of the business activities of the Contractor to ensure there is no inappropriate or unauthorized use of Purchasing Entity information at any time. To this end, the Contractor shall safeguard the confidentiality, integrity and availability of Purchasing Entity information and comply with the following conditions:

a. The Contractor shall implement and maintain appropriate administrative, technical and organizational security measures to safeguard against unauthorized access, disclosure or theft of Personal Data and Non-Public Data. Such security measures shall be in accordance with recognized industry practice and not less stringent than the measures the Contractor applies to its own Personal Data and Non-Public Data of similar kind.

b. All data obtained by the Contractor in the performance of the Master Agreement shall become and remain the property of the Purchasing Entity.

c. All Personal Data shall be encrypted at rest and in transit with controlled access. Unless otherwise stipulated, the Contractor is responsible for encryption of the Personal Data. Any stipulation of responsibilities will identify specific roles and responsibilities and shall be included in the service level agreement (SLA), or otherwise made a part of the Master Agreement.

d. Unless otherwise stipulated, the Contractor shall encrypt all Non-Public Data at rest and in transit. The Purchasing Entity shall identify data it deems as Non-Public Data to the Contractor. The level of protection and encryption for all Non-Public Data shall be identified in the SLA.

e. At no time shall any data or processes — that either belong to or are intended for the use of a Purchasing Entity or its officers, agents or employees — be copied, disclosed or retained by the Contractor or any party related to the Contractor for subsequent use in any transaction that does not include the Purchasing Entity.

f. The Contractor shall not use any information collected in connection with the Services issued from this Master Agreement for any purpose other than fulfilling the Services.

3. Data Location: The Contractor shall provide its services to the Purchasing Entity and its end users solely from data centers in the U.S. Storage of Purchasing Entity data at rest shall be located solely in data centers in the U.S. The Contractor shall not allow its personnel or contractors to store Purchasing Entity data on portable devices, including personal computers, except for devices that are used and kept only at its U.S. data centers. The Contractor shall permit its personnel and contractors to access Purchasing Entity data remotely only as required to provide technical support. The Contractor may provide technical user support on a 24/7 basis using a Follow the Sun model, unless otherwise prohibited in a Participating Addendum.

4. Security Incident or Data Breach Notification: The Contractor shall inform the Purchasing Entity of any security incident or data breach within the possession and control of the Contractor and related to the service provided under the Master Agreement, Participating Addendum, or SLA. Such notice shall include, to the best of Contractor's knowledge at that time, the persons affected, their identities, and the Confidential Information and Data disclosed, or shall include if this information is unknown.

a. Incident Response: The Contractor may need to communicate with outside parties regarding a security incident, which may include contacting law enforcement, fielding media inquiries and seeking external expertise as mutually agreed upon, defined by law or contained in the Master Agreement, Participating Addendum, or SLA. Discussing security incidents with the Purchasing Entity should be handled on an urgent as-needed basis, as part of Contractor's communication and mitigation processes as mutually agreed, defined by law or contained in the Master Agreement, Participating Addendum, or SLA.

b. Security Incident Reporting Requirements: Unless otherwise stipulated, the Contractor shall immediately report a security incident related to its service under the Master Agreement, Participating Addendum, or SLA to the appropriate Purchasing Entity.

c. Breach Reporting Requirements: If the Contractor has actual knowledge of a confirmed data breach that affects the security of any Purchasing Entity data that is subject to applicable data breach notification law, the Contractor shall (1) promptly notify the appropriate Purchasing Entity within 24 hours or sooner, unless shorter time is required by applicable law, and (2) take commercially reasonable measures to address the data breach in a timely manner

5. Breach Responsibilities: This section only applies when a Data Breach occurs with respect to Personal Data within the possession or control of the Contractor.

a. The Contractor, unless stipulated otherwise, shall immediately notify the appropriate Purchasing Entity identified contact by telephone in accordance with the agreed upon security plan or security procedures if it reasonably believes there has been a security incident.

b. The Contractor, unless stipulated otherwise, shall promptly notify the appropriate Purchasing Entity identified contact within 24 hours or sooner by telephone, unless shorter time is required by applicable law, if it has confirmed that there is, or reasonably believes that there has been a data breach. The Contractor shall (1) cooperate with the Purchasing Entity as reasonably

requested by the Purchasing Entity to investigate and resolve the data breach, (2) promptly implement necessary remedial measures, if necessary, and (3) document responsive actions taken related to the data breach, including any post-incident review of events and actions taken to make changes in business practices in providing the services, if necessary.

c. Unless otherwise stipulated, if a Data Breach is a direct result of Contractor's breach of its contractual obligation to encrypt Personal Data or otherwise prevent its release, the Contractor shall bear the costs associated with (1) the investigation and resolution of the data breach; (2) notifications to individuals, regulators or others required by federal and state laws or as otherwise agreed to; (3) a credit monitoring service required by state (or federal) law or as otherwise agreed to; (4) a website or a toll-free number and call center for affected individuals required by federal and state laws — all not to exceed the average per record per person cost calculated for data breaches in the United States (currently $217 per record/person) in the most recent Cost of Data Breach Study: Global Analysis published by the Ponemon Institute at the time of the data breach; and (5) complete all corrective actions as reasonably determined by Contractor based on root cause.

6. Notification of Legal Requests: The Contractor shall contact the Purchasing Entity upon receipt of any electronic discovery, litigation holds, discovery searches and expert testimonies related to the Purchasing Entity's data under the Master Agreement, or which in any way might reasonably require access to the data of the Purchasing Entity. The Contractor shall not respond to subpoenas, service of process and other legal requests related to the Purchasing Entity without first notifying and obtaining the approval of the Purchasing Entity, unless prohibited by law from providing such notice.

7. Termination and Suspension of Service:

a. In the event of an early termination of the Master Agreement, Participating or SLA, Contractor shall allow for the Purchasing Entity to retrieve its digital content and provide for the subsequent secure disposal of the Purchasing Entity's digital content.

b. During any period of service suspension, the Contractor shall not take any action to intentionally erase or otherwise dispose of any of the Purchasing Entity's data.

c. In the event of early termination of any Services or agreement in entirety, the Contractor shall not take any action to intentionally erase any Purchasing Entity's data for a period of 1) 45 days after the effective date of termination, if the termination is for convenience; or 2) 60 days after the effective date of termination, if the termination is for cause. After such day period, the Contractor shall have no obligation to maintain or provide any Purchasing Entity data and shall thereafter, unless legally prohibited, delete all Purchasing Entity data in its systems or otherwise in its possession or under its control. In the event of either termination for cause, the Contractor will impose no fees for access and retrieval of digital content to the Purchasing Entity.

d. The Purchasing Entity shall be entitled to any post termination assistance generally made available with respect to the services, unless a unique data retrieval arrangement has been established as part of an SLA.

e. Upon termination of the Services or the Agreement in its entirety, Contractor shall securely dispose of all Purchasing Entity's data in all of its forms, such as disk, CD/ DVD, backup tape and paper, unless stipulated otherwise by the Purchasing Entity. Data shall be permanently deleted and shall not be recoverable, according to National Institute of Standards and Technology (NIST)-approved methods. Certificates of destruction shall be provided to the Purchasing Entity.

8. Background Checks:

a. Upon the request of the Purchasing Entity, the Contractor shall conduct criminal background checks and not utilize any staff, including subcontractors, to fulfill the obligations of the Master Agreement who have been convicted of any crime of dishonesty, including but not limited to criminal fraud, or otherwise convicted of any felony or misdemeanor offense for which incarceration for up to 1 year is an authorized penalty. The Contractor shall promote and maintain an awareness of the importance of securing the Purchasing Entity's information among the Contractor's employees and agents.

b. The Contractor and the Purchasing Entity recognize that security responsibilities are shared. The Contractor is responsible for providing a secure infrastructure. The Purchasing Entity is responsible for its secure guest operating system, firewalls and other logs captured within the guest operating system. Specific shared responsibilities are identified within the SLA.

c. If any of the stated personnel providing services under a Participating Addendum is not acceptable to the Purchasing Entity in its sole opinion as a result of the background or criminal history investigation, the Purchasing Entity, in its' sole option shall have the right to either (1) request immediate replacement of the person, or (2) immediately terminate the Participating Addendum and any related service agreement.

9. Access to Security Logs and Reports:

a. The Contractor shall provide reports on a schedule specified in the SLA to the Purchasing Entity in a format as specified in the SLA and agreed to by both the Contractor and the Purchasing Entity. Reports will include latency statistics, user access, user access IP address, user access history and security logs for all Purchasing Entity files related to the Master Agreement, Participating Addendum, or SLA.

b. The Contractor and the Purchasing Entity recognize that security responsibilities are shared. The Contractor is responsible for providing a secure infrastructure. The Purchasing Entity is responsible for its secure guest operating system, firewalls and other logs captured within the guest operating system. Specific shared responsibilities are identified within the SLA.

10. Contract Audit: The Contractor shall allow the Purchasing Entity to audit conformance to the Master Agreement terms. The Purchasing Entity may perform this audit or contract with a third party at its discretion and at the Purchasing Entity's expense.

11. Data Center Audit: The Contractor shall perform an independent audit of its data centers at least annually at its expense, and provide an unredacted version of the audit report upon request to a Purchasing Entity. The Contractor may remove its proprietary information from the unredacted version. A Service Organization Control (SOC) 2 audit report or approved equivalent sets the minimum level of a third-party audit.

12. Change Control and Advance Notice: The Contractor shall give a minimum forty eight (48) hour advance notice (or as determined by a Purchasing Entity and included in the SLA) to the Purchasing Entity of any upgrades (e.g., major upgrades, minor upgrades, system changes) that may impact service availability and performance. A major upgrade is a replacement of hardware, software or firmware with a newer or better version in order to bring the system up to date or to improve its characteristics. It usually includes a new version number.

Contractor will make updates and upgrades available to Purchasing Entity at no additional costs when Contractor makes such updates and upgrades generally available to its users.

No update, upgrade or other charge to the Service may decrease the Service's functionality, adversely affect Purchasing Entity's use of or access to the Service, or increase the cost of the Service to the Purchasing Entity.

Contractor will notify the Purchasing Entity at least sixty (60) days in advance prior to any major update or upgrade.

13. Security: As requested by a Purchasing Entity, the Contractor shall disclose its non-proprietary system security plans (SSP) or security processes and technical limitations to the Purchasing Entity such that adequate protection and flexibility can be attained between the Purchasing Entity and the Contractor. For example: virus checking and port sniffing — the Purchasing Entity and the Contractor shall understand each other's roles and responsibilities.

14. Non-disclosure and Separation of Duties: The Contractor shall enforce separation of job duties, require commercially reasonable non-disclosure agreements, and limit staff knowledge of Purchasing Entity data to that which is absolutely necessary to perform job duties.

15. Import and Export of Data: The Purchasing Entity shall have the ability to import or export data in piecemeal or in entirety at its discretion without interference from the Contractor at any time during the term of Contractor's contract with the Purchasing Entity. This includes the ability for the Purchasing Entity to import or export data to/from other Contractors. Contractor shall specify if Purchasing Entity is required to provide its' own tools for this purpose, including the optional purchase of Contractors tools if Contractors applications are not able to provide this functionality directly.

16. Responsibilities and Uptime Guarantee: The Contractor shall be responsible for the acquisition and operation of all hardware, software and network support related to the services being provided. The technical and professional activities required for establishing, managing and maintaining the environments are the responsibilities of the Contractor. The system shall be available 24/7/365 (with agreed-upon maintenance downtime), and provide service to customers as defined in the SLA.

17. Subcontractor Disclosure: Contractor shall identify all of its strategic business partners related to services provided under this Master Agreement, including but not limited to all subcontractors or other entities or individuals who may be a party to a joint venture or similar agreement with the Contractor, and who shall be involved in any application development and/or operations.

18. Business Continuity and Disaster Recovery: The Contractor shall provide a business continuity and disaster recovery plan upon request and ensure that the Purchasing Entity's recovery time objective (RTO) of XXX hours/days is met. (XXX hour/days shall be provided to Contractor by the Purchasing Entity.) Contractor must work with the Purchasing Entity to perform an annual Disaster Recovery test and take action to correct any issues detected during the test in a time frame mutually agreed between the Contractor and the Purchasing Entity.

19. Compliance with Accessibility Standards: The Contractor shall comply with and adhere to Accessibility Standards of Section 508 Amendment to the Rehabilitation Act of 1973 or any other state laws or administrative regulations identified by the Participating Entity..

20. Web Services: The Contractor shall use Web services exclusively to interface with the Purchasing Entity's data in near real time.

21. Encryption of Data at Rest: The Contractor shall ensure hard drive encryption consistent with validated cryptography standards as referenced in FIPS 140-2, Security Requirements for Cryptographic Modules for all Personal Data as identified in the SLA, unless the Contractor presents a justifiable position that is approved by the Purchasing Entity that Personal Data, is required to be stored on a Contractor portable device in order to accomplish work as defined in the scope of work.

22. Subscription Terms: Contractor grants to a Purchasing Entity a license to: (i) access and use the Service for its business purposes; (ii) for PaaS, use underlying software as embodied or used in the Service; and (iii) view, copy, upload and download (where applicable), and use Contractor's documentation.

No Contractor terms, including standard click through license or website terms or use of privacy policy, shall apply to Purchasing Entities unless such terms are included in this Master Agreement.

Exhibit 3 to the Master Agreement: Infrastructure-as-a-Service

1. Data Ownership: The Purchasing Entity will own all right, title and interest in its data that is related to the Services provided by this Master Agreement. The Contractor shall not access Purchasing Entity user accounts or Purchasing Entity data, except (1) in the course of data center operations, (2) in response to service or technical issues, (3) as required by the express terms of this Master Agreement, Participating Addendum, SLA, and/or other contract documents, or (4) at the Purchasing Entity's written request.

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 a 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. This obligation shall survive and extend beyond the term of this Master Agreement.

2. Data Protection: Protection of personal privacy and data shall be an integral part of the business activities of the Contractor to ensure there is no inappropriate or unauthorized use of Purchasing Entity information at any time. To this end, the Contractor shall safeguard the confidentiality, integrity and availability of Purchasing Entity information and comply with the following conditions:

a. The Contractor shall implement and maintain appropriate administrative, technical and organizational security measures to safeguard against unauthorized access, disclosure or theft of Personal Data and Non-Public Data. Such security measures shall be in accordance with recognized industry practice and not less stringent than the measures the Contractor applies to its own Personal Data and Non-Public Data of similar kind.

b. All data obtained by the Contractor in the performance of the Master Agreement shall become and remain the property of the Purchasing Entity.

c. All Personal Data shall be encrypted at rest and in transit with controlled access. Unless otherwise stipulated, the Contractor is responsible for encryption of the Personal Data. Any stipulation of responsibilities will identify specific roles and responsibilities and shall be included in the service level agreement (SLA), or otherwise made a part of the Master Agreement.

d. Unless otherwise stipulated, the Contractor shall encrypt all Non-Public Data at rest and in transit. The Purchasing Entity shall identify data it deems as Non-Public Data to the Contractor. The level of protection and encryption for all Non-Public Data shall be identified in the SLA.

e. At no time shall any data or processes — that either belong to or are intended for the use of a Purchasing Entity or its officers, agents or employees — be copied, disclosed or retained by the Contractor or any party related to the Contractor for subsequent use in any transaction that does not include the Purchasing Entity.

f. The Contractor shall not use any information collected in connection with the Services issued from this Master Agreement for any purpose other than fulfilling the Services.

3. Data Location: The Contractor shall provide its services to the Purchasing Entity and its end users solely from data centers in the U.S. Storage of Purchasing Entity data at rest shall be located solely in data centers in the U.S. The Contractor shall not allow its personnel or contractors to store Purchasing Entity data on portable devices, including personal computers, except for devices that are used and kept only at its U.S. data centers. The Contractor shall permit its personnel and contractors to access Purchasing Entity data remotely only as required to provide technical support. The Contractor may provide technical user support on a 24/7 basis using a Follow the Sun model, unless otherwise prohibited in a Participating Addendum.

4. Security Incident or Data Breach Notification: The Contractor shall inform the Purchasing Entity of any security incident or data breach related to Purchasing Entity's Data within the possession or control of the Contractor and related to the service provided under the Master Agreement, Participating Addendum, or SLA. Such notice shall include, to the best of Contractor's knowledge at that time, the persons affected, their identities, and the Confidential Information and Data disclosed, or shall include if this information is unknown.

a. Security Incident Reporting Requirements: The Contractor shall report a security incident to the Purchasing Entity identified contact immediately as soon as possible or promptly without out reasonable delay, or as defined in the SLA.

b. Breach Reporting Requirements: If the Contractor has actual knowledge of a confirmed data breach that affects the security of any purchasing entity's content that is subject to applicable data breach notification law, the Contractor shall (1) as soon as possible or promptly without out reasonable delay notify the Purchasing Entity, unless shorter time is required by applicable law, and (2) take commercially reasonable measures to address the data breach in a timely manner.

5. Breach Responsibilities: This section only applies when a Data Breach occurs with respect to Personal Data within the possession or control of the Contractor and related to the service provided under the Master Agreement, Participating Addendum, or SLA.

a. The Contractor, unless stipulated otherwise, shall immediately notify the appropriate Purchasing Entity identified contact by telephone in accordance with the agreed upon security plan or security procedures if it reasonably believes there has been a security incident.

b. The Contractor, unless stipulated otherwise, shall promptly notify the appropriate Purchasing Entity identified contact within 24 hours or sooner by telephone, unless shorter time is required by applicable law, if it has confirmed that there is, or reasonably believes that there has been a data breach. The Contractor shall (1) cooperate with the Purchasing Entity as reasonably requested by the Purchasing Entity to investigate and resolve the Data Breach, (2) promptly implement necessary remedial measures, if necessary, and (3) document responsive actions taken related to the Data Breach, including any post-incident review of events and actions taken to make changes in business practices in providing the services, if necessary.

c. Unless otherwise stipulated, if a Data Breach is a direct result of Contractor's breach of its contractual obligation to encrypt Personal Data or otherwise prevent its release, the Contractor shall bear the costs associated with (1) the investigation and resolution of the data breach; (2) notifications to individuals, regulators or others required by federal and state laws or as otherwise agreed to; (3) a credit monitoring service required by state (or federal) law or as otherwise agreed to; (4) a website or a toll-free number and call center for affected individuals required by federal and state laws — all not to exceed the average per record per person cost calculated for data breaches in the United States (currently $217 per record/person) in the most recent Cost of Data Breach Study: Global Analysis published by the Ponemon Institute at the time of the data breach; and (5) complete all corrective actions as reasonably determined by Contractor based on root cause.

6. Notification of Legal Requests: The Contractor shall contact the Purchasing Entity upon receipt of any electronic discovery, litigation holds, discovery searches and expert testimonies related to the Purchasing Entity's data under the Master Agreement, or which in any way might reasonably require access to the data of the Purchasing Entity. The Contractor shall not respond to subpoenas, service of process and other legal requests related to the Purchasing Entity without first notifying and obtaining the approval of the Purchasing Entity, unless prohibited by law from providing such notice.

7. Termination and Suspension of Service:

a. In the event of an early termination of the Master Agreement, Participating or SLA, Contractor shall allow for the Purchasing Entity to retrieve its digital content and provide for the subsequent secure disposal of the Purchasing Entity's digital content.

b. During any period of service suspension, the Contractor shall not take any action to intentionally erase or otherwise dispose of any of the Purchasing Entity's data.

c. In the event of early termination of any Services or agreement in entirety, the Contractor shall not take any action to intentionally erase any Purchasing Entity's data for a period of 1) 45 days after the effective date of termination, if the termination is for convenience; or 2) 60 days after the effective date of termination, if the termination is for cause. After such day period, the Contractor shall have no obligation to maintain or provide any Purchasing Entity data and shall thereafter, unless legally prohibited, delete all Purchasing Entity data in its systems or otherwise in its possession or under its control. In the event of either termination for cause, the Contractor will impose no fees for access and retrieval of digital content to the Purchasing Entity.

d. The Purchasing Entity shall be entitled to any post termination assistance generally made available with respect to the services, unless a unique data retrieval arrangement has been established as part of an SLA.

e. Upon termination of the Services or the Agreement in its entirety, Contractor shall securely dispose of all Purchasing Entity's data in all of its forms, such as disk, CD/ DVD, backup tape and paper, unless stipulated otherwise by the Purchasing Entity. Data shall be permanently deleted

and shall not be recoverable, according to National Institute of Standards and Technology (NIST)-approved methods. Certificates of destruction shall be provided to the Purchasing Entity.

8. Background Checks:

a. Upon the request of the Purchasing Entity, the Contractor shall conduct criminal background checks and not utilize any staff, including subcontractors, to fulfill the obligations of the Master Agreement who have been convicted of any crime of dishonesty, including but not limited to criminal fraud, or otherwise convicted of any felony or misdemeanor offense for which incarceration for up to 1 year is an authorized penalty. The Contractor shall promote and maintain an awareness of the importance of securing the Purchasing Entity's information among the Contractor's employees and agents.

b. The Contractor and the Purchasing Entity recognize that security responsibilities are shared. The Contractor is responsible for providing a secure infrastructure. The Purchasing Entity is responsible for its secure guest operating system, firewalls and other logs captured within the guest operating system. Specific shared responsibilities are identified within the SLA.

c. If any of the stated personnel providing services under a Participating Addendum is not acceptable to the Purchasing Entity in its sole opinion as a result of the background or criminal history investigation, the Purchasing Entity, in its' sole option shall have the right to either (1) request immediate replacement of the person, or (2) immediately terminate the Participating Addendum and any related service agreement.

9. Access to Security Logs and Reports:

a. The Contractor shall provide reports on a schedule specified in the SLA to the Contractor directly related to the infrastructure that the Contractor controls upon which the Purchasing Entity's account resides. Unless otherwise agreed to in the SLA, the Contractor shall provide the public jurisdiction a history or all API calls for the Purchasing Entity account that includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters and the response elements returned by the Contractor. The report will be sufficient to enable the Purchasing Entity to perform security analysis, resource change tracking and compliance auditing

b. The Contractor and the Purchasing Entity recognize that security responsibilities are shared. The Contractor is responsible for providing a secure infrastructure. The Purchasing Entity is responsible for its secure guest operating system, firewalls and other logs captured within the guest operating system. Specific shared responsibilities are identified within the SLA.

10. Contract Audit: The Contractor shall allow the Purchasing Entity to audit conformance to the Master Agreement terms. The Purchasing Entity may perform this audit or contract with a third party at its discretion and at the Purchasing Entity's expense.

11. Data Center Audit: The Contractor shall perform an independent audit of its data centers at least annually and at its own expense, and provide an unredacted version of the audit report upon request. The Contractor may remove its proprietary information from the unredacted version. For example, a Service Organization Control (SOC) 2 audit report would be sufficient.

12. Change Control and Advance Notice: The Contractor shall give a minimum forty eight (48) hour advance notice (or as determined by a Purchasing Entity and included in the SLA) to the Purchasing Entity of any upgrades (e.g., major upgrades, minor upgrades, system changes) that may impact service availability and performance. A major upgrade is a replacement of hardware, software or firmware with a newer or better version in order to bring the system up to date or to improve its characteristics. It usually includes a new version number.

Contractor will make updates and upgrades available to Purchasing Entity at no additional costs when Contractor makes such updates and upgrades generally available to its users.

No update, upgrade or other charge to the Service may decrease the Service's functionality, adversely affect Purchasing Entity's use of or access to the Service, or increase the cost of the Service to the Purchasing Entity.

Contractor will notify the Purchasing Entity at least sixty (60) days in advance prior to any major update or upgrade.

13. Security: As requested by a Purchasing Entity, the Contractor shall disclose its non-proprietary system security plans (SSP) or security processes and technical limitations to the Purchasing Entity such that adequate protection and flexibility can be attained between the Purchasing Entity and the Contractor. For example: virus checking and port sniffing — the Purchasing Entity and the Contractor shall understand each other's roles and responsibilities.

14. Non-disclosure and Separation of Duties: The Contractor shall enforce separation of job duties, require commercially reasonable non-disclosure agreements, and limit staff knowledge of Purchasing Entity data to that which is absolutely necessary to perform job duties.

15. Import and Export of Data: The Purchasing Entity shall have the ability to import or export data in piecemeal or in entirety at its discretion without interference from the Contractor at any time during the term of Contractor's contract with the Purchasing Entity. This includes the ability for the Purchasing Entity to import or export data to/from other Contractors. Contractor shall specify if Purchasing Entity is required to provide its' own tools for this purpose, including the optional purchase of Contractors tools if Contractors applications are not able to provide this functionality directly.

16. Responsibilities and Uptime Guarantee: The Contractor shall be responsible for the acquisition and operation of all hardware, software and network support related to the services being provided. The technical and professional activities required for establishing, managing and maintaining the environments are the responsibilities of the Contractor. The system shall be available 24/7/365 (with agreed-upon maintenance downtime), and provide service to customers as defined in the SLA.

17. Subcontractor Disclosure: Contractor shall identify all of its strategic business partners related to services provided under this Master Agreement, including but not limited to all subcontractors or other entities or individuals who may be a party to a joint venture or similar agreement with the Contractor, and who shall be involved in any application development and/or operations.

18. Business Continuity and Disaster Recovery: The Contractor shall provide a business continuity and disaster recovery plan upon request and ensure that the Purchasing Entity's recovery time objective (RTO) of XXX hours/days is met. (XXX hour/days shall be provided to Contractor by the Purchasing Entity.) Contractor must work with the Purchasing Entity to perform an annual Disaster Recovery test and take action to correct any issues detected during the test in a time frame mutually agreed between the Contractor and the Purchasing Entity.

19. **Subscription Terms:** Contractor grants to a Purchasing Entity a license to: (i) access and use the Service for its business purposes; (ii) for IaaS, use underlying software as embodied or used in the Service; and (iii) view, copy, upload and download (where applicable), and use Contractor's documentation.

No Contractor terms, including standard click through license or website terms or use of privacy policy, shall apply to Purchasing Entities unless such terms are included in this Master Agreement.

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.

Service Model:	**Low Risk Data**	**Moderate Risk Data**	**High Risk Data**	**Deployment Models Offered:**
Saas	AWS Standard & GovCloud Regions (see Appendix 1) Google Apps	AWS Standard & GovCloud Regions (see Appendix 1) Google Apps	N/A	Public Cloud, Hybrid Cloud
IaaS	AWS Standard & GovCloud Regions (see Appendix 1) Akamai	AWS Standard & GovCloud Regions (see Appendix 1) Akamai	AWS GovCloud Region (see Appendix 1)	Public Cloud, Hybrid Cloud
PaaS	AWS Standard & GovCloud Regions (see Appendix 1) Akamai	AWS Standard & GovCloud Regions (see Appendix 1) Akamai	AWS GovCloud Region (see Appendix 1)	Public Cloud, Hybrid Cloud

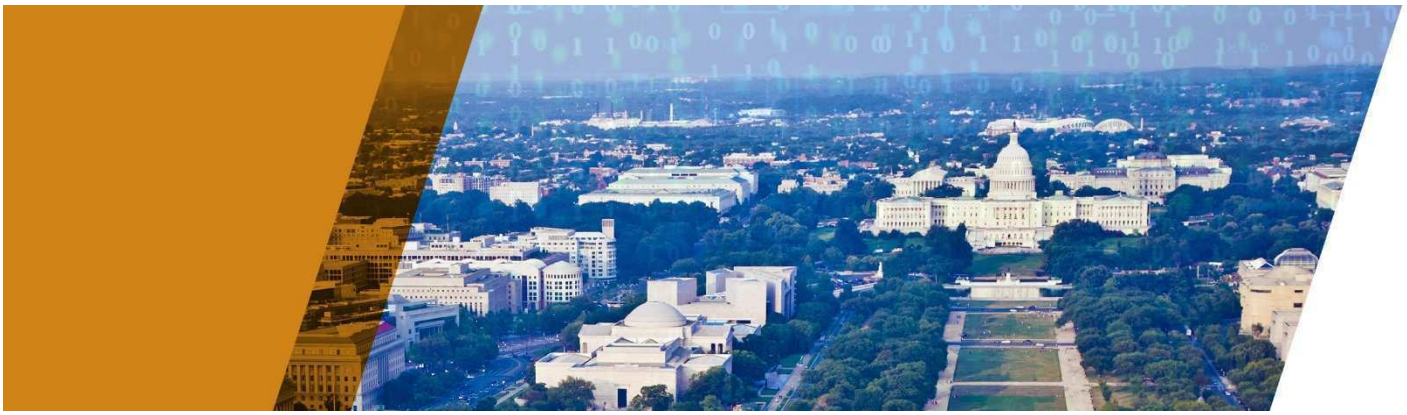

THE STATE OF UTAH
DIVISION OF PURCHASING

NASPO ValuePoint Master Agreement for Cloud Solutions

Utah Solicitation Number CH16012

Technical Proposal – Appendix 1

PRESENTED BY

March 10, 2016

DLT Solutions / 2411 Dulles Corner Park, Suite 800 / Herndon VA 20171 / 703.709.7172 / www.dlt.com
TIN 54-1599882 / DUNS # 78-6468199 / CAGE Code 0S0H9 / GSA Schedule 70 GS-35F-4543G

Product/Service	Description	IaaS	SaaS	PaaS
AWS				
EC2	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 EC2 Container Service	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.			●
Elastic Load Balancing	Elastic Load Balancing automatically distributes your incoming application traffic across multiple Amazon EC2 instances. It detects unhealthy instances and reroutes traffic to healthy instances until the unhealthy instances have been restored. Elastic Load Balancing automatically scales its request handling capacity in response to incoming traffic.	●		
Amazon VPC	Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS.	●		
Amazon Route 53	Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service.	●		
Amazon S3	Amazon Simple Storage Service (Amazon S3) is storage for the Internet. You can use Amazon S3 to store and retrieve any amount of data, at any time, from anywhere on the web. You can accomplish these tasks using the simple and intuitive web interface of the AWS Management Console.	●		
Amazon Glacier	Amazon Glacier is a storage service optimized for infrequently used data, or "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 even 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	●		

Product/Service	Description	IaaS	SaaS	PaaS
Amazon EBS	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 any 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.	●		
Amazon CloudFront	Amazon CloudFront is a content delivery web service. It integrates with other AWS cloud services to give developers and businesses an easy way to distribute content to end users with low latency, high data transfer speeds, and no commitments.	●		
AWS Storage Gateway	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 AWS's storage infrastructure.	●		
Amazon RDS	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.			●
Amazon DynamoDB	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 any amount of data, and serve any level 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.			●
Amazon RedShift	Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse solution that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small for just $0.25 per hour with no commitments or up-front costs and scale to a petabyte or more for $1,000 per terabyte per year, less than a tenth of most other data warehousing solutions.			●
Amazon ElastiCache	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.	●		

Product/Service	Description	IaaS	SaaS	PaaS
Amazon EMR	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 such tasks as web indexing, data mining, log file analysis, machine learning, scientific simulation, and data warehousing.	●		●
Amazon Kinesis	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.”	●		
AWS Data Pipeline	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 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.	●		
AWS Identity & Access Mgt.	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 users can access.	●		
AWS Service Catalog	AWS Service Catalog is a service that allows administrators to create and manage approved catalogs of resources that end users can then access via 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.	●		
AWS Config	AWS Config is a fully managed service that provides you with an AWS 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 all configuration details, and determine how a resource was configured at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.	●		

Product/Service	Description	IaaS	SaaS	PaaS
AWS Key Management Service	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 including Amazon EBS, Amazon S3, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.	●		
AWS CloudWatch	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 trends, and take automated action based on the state of your cloud environment.	●		
AWS Command Line Interface	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.			●
API's	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.			●
AWS Elastic Beanstalk	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 health monitoring.			●
AWS CodeDeploy	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 to one Amazon EC2 instance or thousands.			●
AWS CodeCommit	AWS CodeCommit is a secure, highly scalable, managed source control 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.			●

Product/Service	Description	IaaS	SaaS	PaaS
Amazon WorkDocs	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 resort to emailing multiple versions of their files as attachments.		●	
Amazon Workspaces	Amazon WorkSpaces is a fully managed desktop computing service in the cloud. Amazon WorkSpaces allows customers to easily provision cloud-based desktops that allow end users to access the documents, applications, and resources they need with the device of their choice, including laptops, iPad, Kindle Fire, or Android tablets. With a few clicks in the AWS Management Console, customers can provision a high-quality cloud desktop experience for any number of users at a cost that is highly competitive with traditional desktops and half the cost of most Virtual Desktop Infrastructure (VDI) solutions.		●	
Amazon AppStream	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.			●
Amazon CloudSearch	Amazon CloudSearch is a fully 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. 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.	●		
Amazon SWF	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 and maintaining their state.			●

Product/Service	Description	IaaS	SaaS	PaaS
Amazon SQS	Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles messages or workflows between other components in a system.	●		
Amazon SES	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.	●		
Amazon SNS	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.	●		
Amazon Elastic Transcoder	Amazon Elastic Transcoder lets you convert media files that you have stored in Amazon S3 into media files in the formats required by consumer playback devices. For example, you can convert large, high-quality digital media files into formats that users can play back on mobile devices, tablets, web browsers, and connected televisions.			●
Amazon Cognito	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 also support unauthenticated guests.			●
AKAMAI				
Dynamic Site Accelerator	Akamai Dynamic Site Accelerator (DSA)/DSA Secure is an industry leading scale, performance and intelligence solution to make web experiences fast, reliable and secure across all end-user situations while maximizing infrastructure offload and operational agility and standard and secure websites.	●		
Dynamic Site Delivery	Akamai's Dynamic Site Delivery (DSD) maximizes infrastructure cost savings through offload by caching static and dynamic web content onto Akamai's Intelligent Platform and brings the content closer to users.	●		
Bot Manager	Bot Manager uses a number of different detection techniques to determine if a client making a port 80 HTTP or port 443 HTTPS request at the Edge is a human or a bot. Bot Manager categorized the bots into known bot categories and unknown detected bot categories. Customer may set policies to apply different response actions to different categories of bot traffic.	●		

Product/Service	Description	IaaS	SaaS	PaaS
Client Reputation	The Client Reputation module is designed to help Customers protect their online applications from attacks, improve accuracy, and fight threats. The Client Reputation service provides Customer with an interface to filter out malicious end-user clients of Customer based on a risk score. The Client Reputation service leverages algorithms to compute a risk score of Customer's clients based upon Threat Data compiled on the Akamai network. Customer may set a policy in the Customer Portal to filter and/or take action based on client scores according to Customer's preferences.	●		
Cloudlet	A value-added application which complements Akamai's core delivery solutions to solve specific business challenges.	●		
Ion Standard/ Ion Standard Secure	Ion enables organizations to offload expensive web infrastructure, including bandwidth, compute and storage while providing web and mobile application availability and scalability.	●		
Fast DNS/ Fast DNS Secure Option	DNS has become a favorite target for DDoS attacks, not just because of its critical role in the Web infrastructure, but also because it is typically one of its least scalable components. Fast DNS combines global scalability with IP throttling to provide an effective defense against DNS-based DDoS attacks.	●		
Kona DDoS Defender	Kona DDoS Defender helps organizations better protect their websites and web applications from the increasing size, frequency, and sophistication of DoS and DDoS attacks. Leveraging the global scale and resiliency of the Akamai Intelligent Platform™, Kona DDoS Defender stops attacks at the edge of the network, before they can reach the application	●		
Kona Site Defender	Kona Site Defender leverages a multi-layered toolset designed to defend against sophisticated attacks employing multiple methodologies. Akamai DDoS defense capabilities are always on. This allows the solution to adapt to the unique nature of each attack method in real time.	●		
Kona Web Application Firewall Overview	The Kona WAF is designed to detect potential Web application attacks in HTTP traffic before the traffic reaches the customer's data center. It provides customers with vital information that can be used to detect anomalous and potentially malicious patterns in HTTP request headers and body contents, and be set-up to either issue an alert or block the traffic altogether. Should traffic be blocked, the service provides customers with several options for responding to blocked HTTP requests. It provides customers with summary and trend reporting on attacks through the Akamai LUNA Control Center. Customers who need more detailed visibility and information about the attack traffic can use the Akamai Security Monitor, or configure the firewall to send logs to their log management system directly for analysis.	●		

Product/Service	Description	IaaS	SaaS	PaaS
Mobile Detection & Redirect	Akamai Mobile Detection and Redirect offers a highly scalable solution that is based on evaluating incoming HTTP requests at the Akamai Edge – a leading, highly pervasive platform for content delivery and Web site acceleration, comprising more than one hundred thousand servers in 90 countries across thousands of networks.	●		
NetStorage	NetStorage is designed to store content that is delivered to end users by one of Akamai's Intelligent Platform content delivery services. This means that end users obtain NetStorage content from Akamai Edge server caches.	●		
Prolexic Proxy*	As part of the Akamai Intelligent Platform, Prolexic Proxy provides a cloud-based architecture to protect organizations from DDoS attacks before they reach the data center. It employs a globally-distributed DDoS mitigation infrastructure to inspect network traffic for DDoS attacks and then mitigate those attacks, forwarding clean traffic to the data center.*Optional Product*	●		
Prolexic Routed*	Prolexic Routed provides a cloud-based architecture to protect organizations from DDoS attacks before they reach the data center. It employs a globally-distributed DDoS mitigation infrastructure to inspect network traffic for DDoS attacks and then mitigate those attacks, forwarding clean traffic to the data center. *Optional Product*	●		
GOOGLE				
Google Apps	Google Apps is a suite of cloud computing productivity and collaboration software tools and software. It includes Gmail, Google Drive, Google Hangouts, Google Calendar, and Google Docs.		●	

Attachment C

NASPO ValuePoint Master Agreement for Cloud Solutions
Utah Solicitation Number CH16012 (March 10, 2016)

Cloud Solutions by Category

Vendor	Product Group	Discount % Off List
Akamai		
	Infrastructure as a Service	21.00%
	Value Added Services	15.00%
Amazon Web Services*		
	Software as a Service	3.00%
	Infrastructure as a Service	3.00%
	Platform as a Service	3.00%
Google		
	Software as a Service	1.00%
DLT Solutions		
	Value Added Services	30.00%

Additional Value Added Services

Maintenance Services			
		Onsite Hourly Rate $	$ 229.73
		Remote Hourly Rate $	$ 229.73
Professional Services			
	Deployment Services		
		Onsite Hourly Rate $	$ 159.54
		Remote Hourly Rate $	$ 159.54
	Consulting/Advisory Services		
		Onsite Hourly Rate $	$ 192.36
		Remote Hourly Rate $	$ 192.36
	Achitectural Design Services		
		Onsite Hourly Rate $	$ 191.44
		Remote Hourly Rate $	$ 191.44
	Statement of Work Services		
		Onsite Hourly Rate $	$ 174.88
		Remote Hourly Rate $	$ 174.88
Partner Services			
		Onsite Hourly Rate $	$ 229.73
		Remote Hourly Rate $	$ 229.73
Training Deployment Services**			
		Onsite Hourly Rate $	
		Remote Hourly Rate $	

*Support pricing for Amazon Web Services will be calculated and invoiced at the individual Customer Account level as follows:

Support Services Pricing Tiers	% of Gross Fees[1]
Support Pricing Tier 1 – The First $10K	10%
Support Pricing Tier 2 – The Next $10K-$80K	7%
Support Pricing Tier 3 – The Next $80K-$250K	5%
Support Pricing Tier 4 – Over $250K	3%

[1]Gross Fees refer to list prices

** Please reference the price list for DLT Value Added Services for proposed training offerings.

ATTACHMENT D

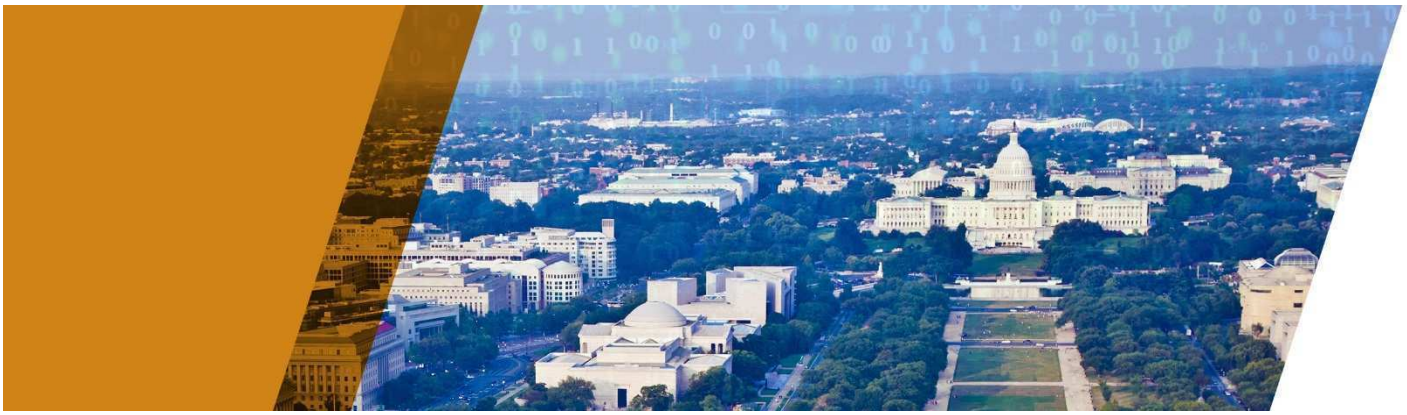

THE STATE OF UTAH • DIVISION OF PURCHASING

NASPO ValuePoint Master Agreement for Cloud Solutions

Utah Solicitation Number CH16012

Technical Proposal

PRESENTED BY

March 10, 2016

DLT Solutions / 2411 Dulles Corner Park, Suite 800 / Herndon VA 20171 / 703.709.7172 / www.dlt.com
TIN 54-1599882 / DUNS # 78-6468199 / CAGE Code 0S0H9 / GSA Schedule 70 GS-35F-4543G

TABLE OF CONTENTS

LIST OF FIGURES

ATTACHMENTS (PROVIDED IN SEPARATE FILES)

DLT File 02 - Attachment A - Redlined Master Agreement Terms and Conditions.docx

DLT File 03 - Attachment B Exhibit 1 – CAIQ – AWS.pdf

DLT File 04 - Attachment B Exhibit 1 – CAIQ – Akamai.xlsx

DLT File 05 - Attachment B Exhibit 1 – CAIQ – Google.xlsx

DLT File 06 - Attachment H - Identification of Service Models Matrix.docx

APPENDICES (PROVIDED IN SEPARATE FILES)

DLT File 07 - Appendix 1 Products/Service Model Matrix

DLT File 08 - Appendix 2 Amazon Web Services Risk and Compliance Whitepaper

DLT File 09 - Appendix 3 SLAs

DLT File 10 - Appendix 4 AWS Cloud Services

DLT File 11 - Appendix 5 AWS Overview of Security Processes

DLT File 12 - Appendix 6 Google Terms of Service and Data Processing Amendment

DLT File 13 - Appendix 7 Google Apps for Business Agreement

DLT File 14 - Appendix 8 Akamai Log Delivery Service User Guide

DLT File 15 - Appendix 9 Google Apps Security Whitepaper

DLT File 16 - Appendix 10 Standard Support Package Overview

DLT File 17 - Appendix 11 Akamai Sample Reports

March 10, 2016

State of Utah Division of Purchasing
3150 State Office Building, Capitol Hill
Salt Lake City, Utah 84114-1061

Attention: Christopher Hughes, Assistant Director
State of Utah, Division of Purchasing

Re: Utah Solicitation Number CH16012 – NASPO ValuePoint Master Agreement for Cloud Solutions

Dear Mr. Hughes:

Thank you for the opportunity to respond to the State of Utah's Request for Proposal (RFP) for NASPO ValuePoint Master Agreement for Cloud Solutions. DLT Solutions, LLC (DLT) is confident that we will provide the State of Utah (as the Lead State), NASPO ValuePoint, and Participating Entities with an unsurpassed combination of cloud solutions, service, support, billing and contract administration.

Per RFP Section 5.2, DLT presents the following:

- [RFP Reference 5.2.1] DLT understands that we may be required to negotiate additional terms and conditions, including additional administrative fees, with Participating Entities when executing a Participating Addendum.

- [RFP Reference 5.2.2] The following individuals, all DLT employees, were responsible for writing this proposal:
    - Brian Strosser, Chief Operating Officer
    - Gary Danoff, Senior Vice President, Cloud Solutions
    - Duane Thorpe, Sales Manager, Cloud Solutions
    - Greg Agana, Senior Systems Engineer, Cloud Solutions
    - Danni Schuder, Contracts Administrator

- [RFP Reference 5.2.3] DLT Solutions, LLC is not currently suspended, debarred or otherwise excluded from federal or state procurement and non-procurement programs.

- [RFP Reference 5.2.4] DLT acknowledges 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.

- [RFP Reference 5.2.5] Please see Appendix 1, Products/Service Model Matrix.

- [RFP Reference 5.2.6] Please see completed Attachment H.

If you have any questions or require additional information of a technical nature, please contact Mr. Duane Thorpe at 703-773-1163 or Duane.Thorpe@dlt.com. If you have any questions or require additional information of a contractual nature, please contact Ms. Danni Schuder at 703-708-9661 or Danni.Schuder@dlt.com.

We look forward to your analysis of our proposal.

Sincerely,

Brian T. Strosser
Chief Operating Officer
DLT Solutions, LLC

1.0 RFP SIGNATURE PAGE [RFP REFERENCE 5.1]

Please see next page.

State of Utah Vendor Information Form

Legal Company Name (include d/b/a if applicable)	Federal Tax Identification Number	State of Utah Sales Tax ID Number
DLT Solutions, LLC	54-1599882	Not applicable

Ordering Address	City	State	Zip Code
2411 Dulles Corner Park, Suite 800	Herndon	VA	20171-6168

Remittance Address (if different from ordering address)	City	State	Zip Code
PO Box 102549	Atlanta	GA	30368-2549

Type	Company Contact Person
☐ Proprietorship ☐ Partnership ☐ Government ☒ For-Profit Corporation ☐ Non-Profit Corporation	Duane Thorpe

Telephone Number (include area code)	Fax Number (include area code)
703-773-1163	866-708-7033

Company's Internet Web Address	Email Address
http://www.dlt.com	Duane.Thorpe@dlt.com

Offeror's Authorized Representative's Signature

Brian T. Strosser

Type or Print Name

Brian T. Strosser

Position or Title of Authorized Representative

Chief Operating Officer

Date: 3/9/2016

2.0 EXECUTIVE SUMMARY [RFP REFERENCE 5.4]

For 25 years, DLT Solutions, LLC (DLT) has been dedicated to working with the U.S. Public Sector to solve IT challenges including migration to the cloud, protect against cybersecurity threats, and optimize and consolidate data center environments. DLT is a US-based, award-winning provider of Government IT and Cloud Solutions, distinguished for its industry expertise and the process integration it has developed with its world-class software and hardware manufacturer partners. By selectively partnering with only the top innovators in the industry, DLT is differentiated by acquiring deep subject matter knowledge and providing solution delivery to better serve public sector customers.

DLT was one of Amazon Web Services' (AWS) first Government partners. Starting in 2011, DLT invested in staff and infrastructure across the company to best support the Government's migration into the cloud, culminating in:

- Contracts/Legal team with expertise in the specifics and nuances of cloud technology contract terms and conditions
- Enhanced information technology systems to support utility pricing model
- Cloud-certified engineering staff across all aspects of information technology lifecycle including cloud architecture, cloud consulting, and cloud technical support
- Cloud billing and invoicing system
- Webinars, electronic books (including Cloud Computing for Govies), and other cloud support tools to assist migration to the cloud
- Custom cloud governance capabilities to better support our customers

DLT has partnered with Day1 Solutions to form Team DLT, creating a strong, comprehensive delivery model in support of NASPO. As a Premier Consulting Partner for Amazon Web Services, DLT Solutions helps public sector customers identify appropriate use cases for cloud computing, as well as design, implement and manage cloud solutions that meet their unique needs. Day1, an AWS Advanced Consulting partner, brings additional expertise to the table in large-scale deployments and professional services bench strength. Together, Team DLT offers a premier product portfolio and delivery team to meet the varied and ever-changing cloud solution requirements of the NASPO community. Additionally, Day1 is a current holder of the NASPO Value Point/WSCA contract and as such has an intimate working knowledge of the existing customer base and capabilities needed to create a deeper penetration into existing customers' cloud migration projects.

Team DLT is offering Amazon Web Services, Google and Akamai, as well as a strong complement of Expert Services and support to round out a complete solution for NASPO's Purchasing Entities. This strong combination of technologies more than meets the minimum mandatory requirements as required by the solicitation with the added benefit of opening the door to access the wider DLT portfolio of cloud technologies as they exist today or as innovation dictates in the future. Our corporate motto of *fast, flexible and responsive* enables us to respond and support the myriad and diverse needs of our customers while providing the best value.

Team DLT's support offerings include pre-sales architecture, migration and implementation support, and a 24x7x365 US Citizen/US Soil ITAR-compliant, vendor authorized help desk. We

have the full complement of professional services needed to implement cloud deployments of any size to our portfolio of cloud technologies. Our support team also provides AWS Certified Managed Services, enabling our customers to realize the full benefit of the AWS offerings while alleviating them of the need to have, or develop, in-house expertise. By supporting the customer at whichever stage they are in the cloud adoption lifecycle, we enable end-to-end solution capabilities to fit the wide array of the NASPO community's needs.

By providing best in class technologies and supporting the customer's procurement, architecting, implementation, and full lifecycle support at any level required, Team DLT will provide the most complete solution set to meet the needs of NASPO Purchasing Entities. The Team DLT model of continuous consultative interaction with our customers, combined with the various premier cloud services contained within our solicitation response, will serve to enable the NASPO community in the adoption of the innovative and ever changing technology sector known as *cloud*.

PLEASE NOTE: DLT has provided complete responses to all RFP requirements in the text of this proposal that fully explain how we meet or exceed the minimum mandatory requirements. We have also provided footnotes containing links to additional information that is supplemental to our response but may be of interest to evaluators—the information available at these links is not required for evaluation but instead is meant as reference material.

3.0 MANDATORY MINIMUM REQUIREMENTS [RFP REFERENCE 5]

If applicable to an Offerors offering, an Offeror must provide a point by point responses to each mandatory minimum requirement. If a mandatory minimum requirement is not applicable to an Offeror's offering then the Offeror must explain why the mandatory minimum 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.

PLEASE NOTE: DLT has provided complete responses to all RFP requirements in the text of this proposal that fully explain how we meet or exceed the minimum mandatory requirements. We have also provided footnotes containing links to additional information that is supplemental to our response but may be of interest to evaluators—the information available at these links is not required for evaluation but instead is meant as reference material.

3.1 SIGNATURE PAGE [RFP REFERENCE 5.1]

Please see section 1.0 on page 3.

3.2 COVER LETTER [RFP REFERENCE 5.2]

Please see page 1.

3.3 ACKNOWLEDGEMENT OF AMENDMENTS [RFP REFERENCE 5.3]

Please see next page.

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 attests to reviewing the documents listed above.

DLT Solutions, LLC
Offeror

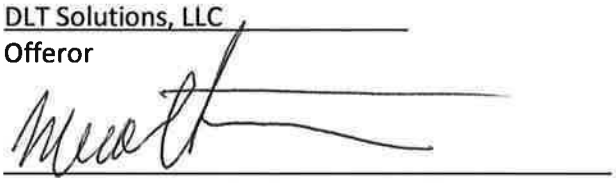

Brian T. Strosser, Chief Operating Officer

3.4 EXECUTIVE SUMMARY [RFP REFERENCE 5.4]

Please see section 2.0 on page 5.

3.5 GENERAL REQUIREMENTS [RFP REFERENCE 5.5]

3.5.1 RFP Reference 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.

DLT Response:

DLT Solutions will, upon award, provide a Usage Report Administrator to ensure compliance with the scope and terms of this contract. DLT's Contracts department currently performs this function for all of the many contracts held by DLT and has back-end systems in place to ensure current and ongoing compliance.

3.5.2 RFP Reference 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.

> [RFP Reference 2.14] 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.

DLT Response:

If awarded a contract, DLT will cooperate with NASPO ValuePoint and SciQuest to ensure that DLT's ordering instructions are quickly and correctly added to the NASPO ValuePoint eMarket Center. Additionally, DLT agrees that we will provide our ordering instructions and any other applicable information to the NASPO ValuePoint eMarket Center Site Admin within thirty (30) days from receipt of written request to provide such information.

3.5.3 RFP Reference 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 to 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 documents.

DLT Response:

DLT has provided completed CSA Consensus Assessment Initiative Questionnaires for all our proposed technologies in the following files:

- DLT File 03 – Attachment B Exhibit 1 – CAIQ – AWS.pdf
- DLT File 04 – Attachment B Exhibit 1 – CAIQ – Akamai.xlsx
- DLT File 05 – Attachment B Exhibit 1 – CAIQ – Google.xlsx

3.5.4 RFP Reference 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.

DLT Response:

Please see Appendix 3 for Service Level Agreements (SLAs) for AWS, Akamai, and Google products and services.

3.6 RFP SECTION 5.6 DELETED

3.7 RECERTIFICATION OF MANDATORY MINIMUMS AND TECHNICAL SPECIFICATIONS [RFP REFERENCE 5.7]

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.

DLT Response:

DLT acknowledges that, if awarded a contract, we will annually certify to the Lead State that our offerings still meet or exceed the technical capabilities discussed in this proposal.

4.0 BUSINESS PROFILE [RFP REFERENCE 6]

4.1 BUSINESS PROFILE [RFP REFERENCE 6.1]

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.**

PLEASE NOTE: DLT has provided complete responses to all RFP requirements in the text of this proposal that fully explain how we meet or exceed the minimum mandatory requirements. We have also provided footnotes containing links to additional information that is supplemental to our response but may be of interest to evaluators—the information available at these links is not required for evaluation but instead is meant as reference material.

DLT Response:

DLT Solutions is a Value Added Reseller (VAR) delivering best in class technology solutions to federal, state and municipal government agencies. Since its founding in 1991, DLT has tightly integrated itself with some of the most prominent IT software and hardware manufacturers in the world including Amazon Web Services, Google, Akamai, Autodesk, Dell Software, NetApp, Oracle, Red Hat, Symantec, SolarWinds and Veritas, among others. With its product portfolio, multiple procurement vehicles and award-winning track record, DLT confidently supports public sector clients with the technology they need, when they need it.

DLT works exclusively with the Public Sector providing innovative and cost-effective commercial-off-the-shelf (COTS) information technology products, solutions, and services to assist agencies in achieving their missions. Our ability to broker vendor and customer relationships, along with our commitment to unparalleled customer satisfaction, enabled DLT to achieve nearly $1 billion in 2015 bookings across more than 28,000 customer delivery orders. DLT is a financially sound company and has averaged 10% year-over-year growth for the past four years. DLT's Cloud Solutions Group, established in 2011, has experienced 100% plus year-over-year growth since inception. Team DLT has delivered enterprise level cloud solutions to our public sector clients, enabling DLT to achieve AWS Premier status, which is the highest status level within the AWS Partner Network.

Our corporate approach to vendor relations has proven extremely successful. DLT works with more than 150 vendors and nearly 500,000 unique line items on active vendor pricelists. DLT stringently ensures that our offerings comply with government regulations and certifications.

DLT Solutions, Inc. was incorporated in 1991 in the State of Virginia. In 2009, DLT Solutions, Inc. became a privately held limited liability company, DLT Solutions, LLC. In 2015, DLT Solutions, LLC was acquired by DLT MergerCo, LLC. As of 3/1/2016, DLT has 260 employees. For the period 2011-2015, our average annual employee retention rate was 88%.

Figure 1 depicts DLT's organizational structure.

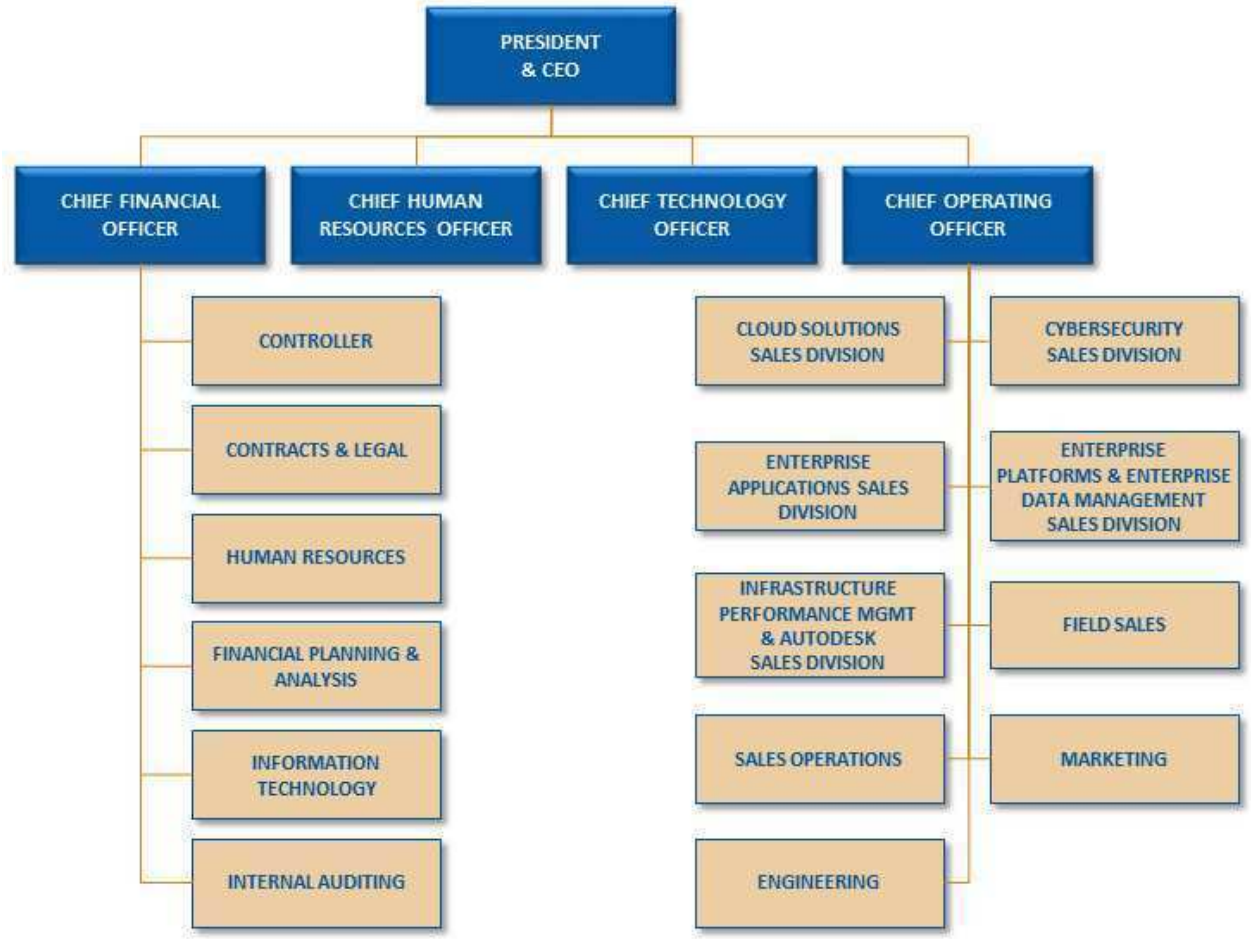

Figure 1. DLT's Organizational Structure

4.2 SCOPE OF EXPERIENCE [RFP REFERENCE 6.2]

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 business provided offerings identical or very similar to those required by this RFP. Government experience is preferred.

DLT Response:

CONTRACT EXPERIENCE

DLT holds numerous large-scale Federal, State and Local Government-wide Acquisition Contracts (GWACs) through which we have provided software, cloud services and professional services—transacting hundreds of millions of dollars annually through those contracts while maintaining strict compliance with contractual requirements. These contracts include:

- GSA IT Schedule 70 Contract Number GS-35F-4543G
- National Institutes of Health (NIH) Chief Information Officers – Commodities and Solutions (NIH CIO-CS)

- Electronic Commodities Store III (ECS III)
- NASA SEWP V
- U.S. Communities (Oracle products)
- Internet2 NET+ (I2)
- National Cooperative Purchasing Alliance (NCPA) - Web Services and Cloud Computing
- The Interlocal Purchasing System (TIPS)
- Texas DIR for Amazon Web Services (TX-DIR-CLOUD)
- King County, Washington
- California Multiple Award Schedule (CMAS)
- California Software Licensing Program (SLP)
- Virginia Association of State College & University Purchasing Professionals (VASCUPP)

In 2014-2015, DLT's largest sales from GWACs were:

Contract No.	Contract Description	2014	2015
GSA	Federal Supply Schedule 70 Commercial IT Equipment	$145,642,151.57	$142,325,137.26
SB ESI BPA-O-R	CHESS DoD BPA (Renewals Only)	$123,981,459.61	$127,538,016.55
NAVY ELA BPA-O-R	Navy Oracle ELA BPA (Support Renewals)	$45,057,810.78	$48,749,423.00
SB ESI BPA - O	DoD Enterprise Software Agreement (ESA) and SmartBuy	$45,072,020.52	$18,603,525.66
QUEST SMARTBUY	DoD ESI and GSA SmartBuy (Quest)	$13,554,592.03	$15,321,426.54

In 2014-2015, DLT's largest sales for ***just*** Amazon Web Services (AWS) from GWACs were:

- I2 $2,314,502.50
- NIH CIO-CS $308,204.47
- ECS-III $811,353.88
- TX-DIR-CLOUD $308,000.00
- NCPA $155,134.88

DLT currently supports scores of customers across the Federal Civilian, DOD, State and Local and Higher Education sectors. The size and scope of projects we support ranges from individual users accessing AWS for a small test project all the way up to large enterprise customers migrating enterprise workloads to the cloud. We provide customer support in the areas of procurement, billing and cost allocation, pre-sales architecture and scoping, training (both onsite and remote), Statement of Work (SOW) --based professional services, 24x7x365 US Citizen, US Soil ITAR-compliant helpdesk support and AWS certified managed services.

One of the first customers that we supported was the National Science Foundation (NSF) beginning in early 2012. The first hurdle that we worked extensively on was the actual procurement. As one of the Federal Government's first ventures into public cloud, there was a great deal of groundbreaking in order to complete a purchase. DLT's Contracts and Legal team

worked extensively with NSF and created a custom procurement path enabling the customer to move forward. NSF had decided to move their Microsoft SharePoint environment to AWS due to data center capacity issues and potential cost savings. DLT and AWS worked hand in hand to enable the migration and continue to support the customer to this day.

One example of large-scale support of an enterprise customer's cloud migration was for a customer that needed to migrate a large amount of data (digitized photos in multiple formats) to AWS and standardize on a single jpg file format. DLT created a custom tool to complete this task. By utilizing our data migration/conversion tool, we were able to upload millions of files while simultaneously converting them to the single format and indexing the files into a searchable format. After the first 20 million files were successfully uploaded in less than half of the time initially anticipated, the customer engaged DLT to upload/transcode many more batches, each of which we successfully completed in incrementally less time than the prior run.

PROGRAM MANAGEMENT

DLT has a proven, mature, services Program Management Delivery Methodology, overseen by the Program Management Office (PMO). A fundamental PMO goal is to improve the execution, management and outcomes of all service engagements—from small short-term projects to large multi-year programs. DLT's proprietary Project Lifecycle Management Methodology (PLM2) is a proven approach that greatly improves project results, mitigates risk, and improves service consistency. DLT's Project Managers expertly lead every project through the five project phases: Initiate, Plan, Execute, Monitor & Control, and Close. PLM2 is based on the concept of "progressive elaboration," which allows the Project Manager to incorporate more details and improve the project management plan as the project evolves. Fundamental to this approach is a continuous improvement process applied to each phase, where all project management documents and other resources are updated regularly based on best practices and lessons learned from current and previous projects.

DLT Project Managers employ a combination of standardized templates and resources to ensure projects are delivered on time, within scope and budget. Projects of any size benefit from use of several key document resources, specifically the Project Plan and Project Dashboard. The Project Dashboard is a central repository for Stakeholder Register, Risk Matrix, Milestones, Deliverable Acceptance Log, Action Item Tracker, Communication Matrix, and Change Control Log. The depth of detail of the Dashboard varies depending on the project's goals, budget and level of risk and is implemented in accordance with the project plan.

DLT's Quality Control Program ensures project Stakeholders are involved with the evaluation and acceptance of deliverables. Project checkpoints strategically placed at the end of each project phase provide Sponsors, Stakeholders and Executive Management oversight on cost, schedule and scope as the project proceeds. This collaborative approach allows Stakeholders the opportunity to review and approve project work and authorize the project team to advance to the next phase. Additionally, the Quality Program provides a defined escalation process to resolve issues. This process is designed to resolve problems in the most efficient manner and at the lowest organizational level possible.

We follow the same Quality Control Program for urgent requirements as described above, however, for urgent requirements that are over and above the standard SLA or contract terms, we assign a quality control analyst to work with DLT's Program Manager to ensure that our quality program is followed while meeting urgent requirements. DLT's ordering system allows for the identification and prioritization of RUSH orders to ensure prompt response for urgent requirements.

DLT's proven quality control processes enabled successful fulfillment of more than 28,000 orders for products and services in 2015. DLT incorporates quality control (QC) checks and quality assurance (QA) processes into all contract and business operations including pricing modifications, order execution, product configuration, warranty support, and unique contract-specific requirements. If any of these aspects fail to meet QC standards, we have well-documented policies and procedures to provide clear guidance for recording, tracking, managing, and resolving quality issues as noted in our Quality Control Program. DLT's Continuous Improvement Program, managed by DLT's PMO, ensures peak scalable operational efficiency while managing competing priorities.

DLT'S INTERNAL SYSTEMS

DLT has developed customized and highly integrated internal systems (Figure 2) to manage and support all aspects of our organizational and accounting controls. Our systems handle the complete end-to-end business model, from quote to order to product delivery to invoicing:

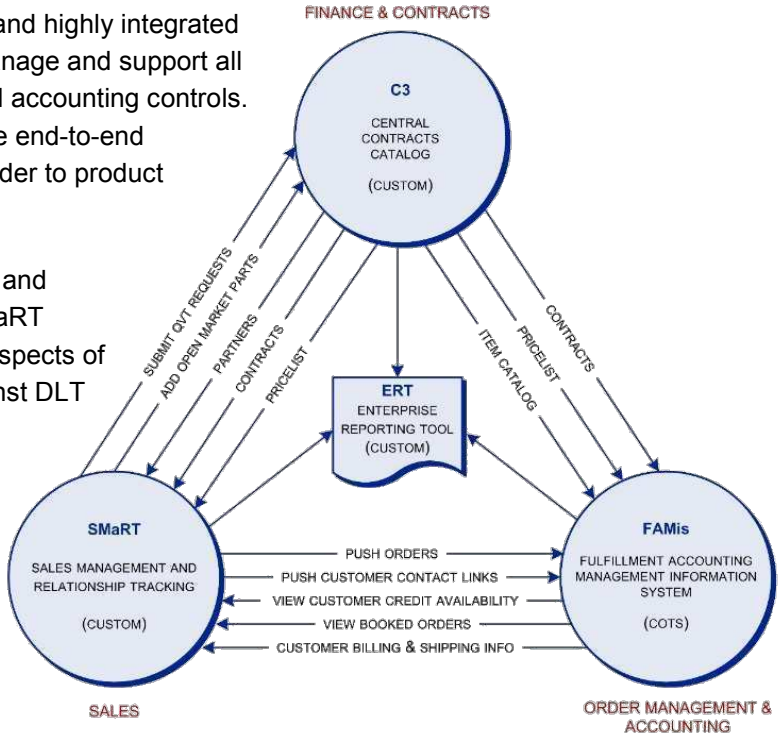

Figure 2. DLT's Internal Systems

- **SMaRT** (Sales Management and Relationship Tracking) – SMaRT maintains and manages all aspects of the lifecycle of an order against DLT contracts from first customer contact through sales quotation through booking of an order. SMaRT is completely integrated with DLT's C3 and FAMis systems to ensure that all quotes adhere to applicable contract terms and conditions. Through various auditable means, SMaRT ensures that customers' orders are accurately and efficiently processed.

- **C3** (Central Contracts Catalog) – The C3 database contains line item pricing for all contracts. It contains systematic rules to ensure that all product pricing is accurate and current, and that all quotes adhere to the respective contract's specific terms and conditions.

- **FAMis** (Fulfillment Accounting Management Information System) – FAMis is DLT's Enterprise-wide accounting and order management system handling both fulfillment and professional services task orders. FAMis also provides all reporting of specific data to our awarded contracts. For order fulfillment and professional services task orders, our Sales Team enters the award information into SMaRT, which transmits the data to FAMis.

- **ERT** (Enterprise Reporting Tool) – ERT is integrated with all DLT systems to provide reporting capabilities to DLT business units.

Figure 3 depicts the flow of a purchase order from when it is received by Sales through shipping and invoicing.

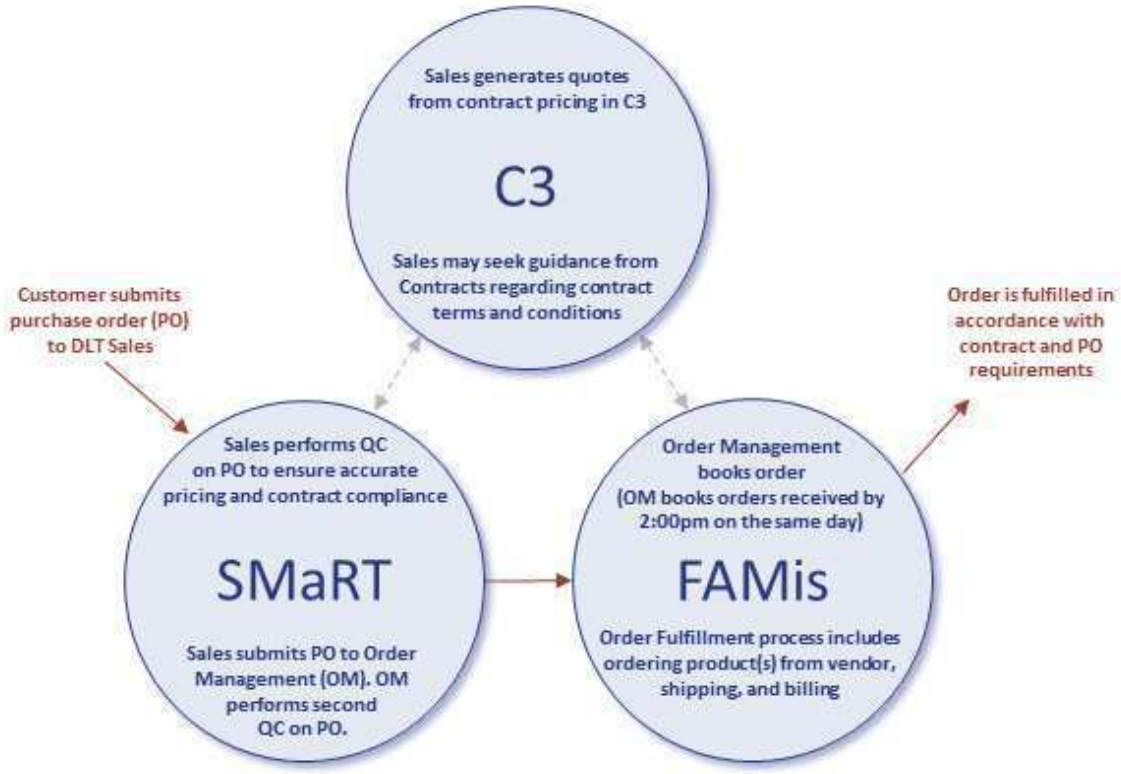

Figure 3. Purchase Order Flow Through DLT's Internal Systems

4.3 FINANCIALS [RFP REFERENCE 6.3]

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.

DLT Response:

DLT Solutions is a financially sound company, averaging annual sales of $863 million for the past five years (Figure 4).

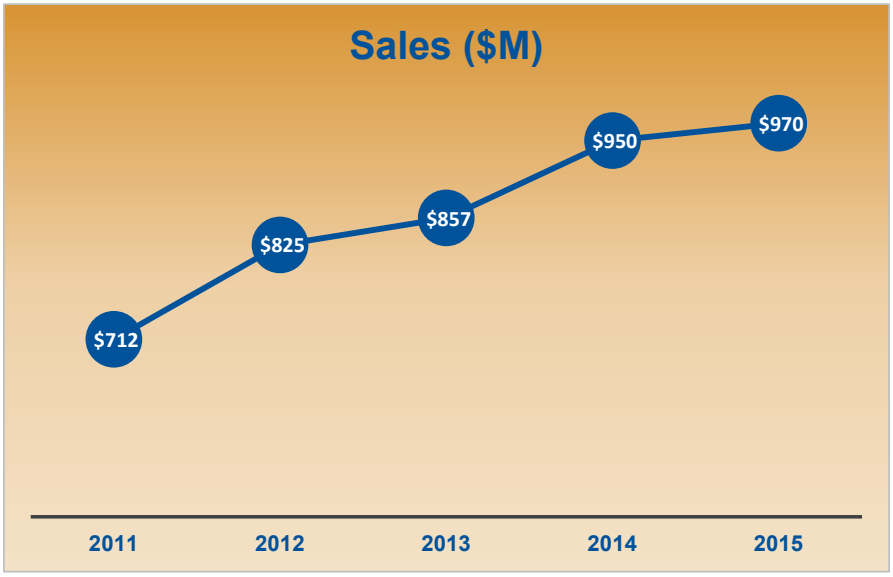

Figure 4. DLT Sales 2011-2015

Our financial systems are reviewed annually for adequate internal controls since we have hundreds of prime and subcontractor awards. We have stringent accounting and financial controls in place to track hours, expenses and costs in accordance with applicable procurement regulations. Additionally, over the years DLT Solutions has gone through pre-award government audits and has been deemed to have adequate billing and accounting systems.

DLT Solutions conforms to all of the standards in FAR 9.104 and will provide specific information upon request. A copy of GSA Form 527 showing DLT's Contractor's Qualifications and Financial Information will be provided upon request.

As a privately held company, the information contained in our financial statements represents proprietary, confidential financial information of DLT Solutions. This information is not subject to FOIA requests.

Please refer to the following files, submitted separately from this technical proposal document:

- DLT File 24 – Utah Claim of Business Confidentiality.pdf
- DLT File 25 – Audited Financial Statements 2013-2014.pdf
- DLT File 26 – Dun and Bradstreet Report 3-1-16.pdf

4.4 GENERAL INFORMATION [RFP REFERENCE 6.4]

4.4.1 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES DIFFERENTIATORS

Below are some features and benefits of AWS that set our cloud infrastructure services apart:

- Pace of Innovation: AWS's pace of innovation is funded and sustained through our economies of scale and commitment to delivering the products and services that matter most to our customers. Our approach to product development and delivery is fundamentally different than that of other Cloud Service Providers (CSPs). We have decentralized, autonomous development teams that work directly with customers. They are empowered to autonomously develop and launch new features based on what they learn from interactions with both commercial and public sector customers. AWS's continual innovation ensures that customers maintain state-of-the-art IT infrastructure without having to make recapitalization investments. As of January 1, 2016, AWS has launched a total of 1,896 new services or major features since inception in 2006 (including 516 in 2014 and 722 in 2015). According to the Gartner, Inc. 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS), Worldwide, "AWS is a thought leader; it is extraordinarily innovative, exceptionally agile, and very responsive to the market."

- Service Breadth and Depth: AWS offers the broadest set of global compute, storage, networking, database, analytics, application, deployment, management, and mobile services that help organizations move faster, lower IT costs, and scale applications. AWS has been continually expanding its services to support virtually any cloud workload, and it now has more than 50 services that serve over one million active customers in more than 190 countries through our 12 regions, 32 Availability Zones, and 54 Edge Locations. Gartner Inc. reported in its 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS), Worldwide report that AWS "has the richest array of IaaS features," "continues to rapidly expand its service offerings and offer higher-level solutions," and has "over 10 times more cloud IaaS compute capacity in use than the aggregate total of the other 14 providers in this Magic Quadrant."

- Partner and Software Ecosystem: According to the Gartner, Inc. 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS), Worldwide report, AWS has attracted "a very large technology partner ecosystem that includes software vendors that have licensed and packaged their software to run on AWS, as well as many vendors that have integrated their software with AWS capabilities. It also has an extensive network of partners that provide application development expertise, managed services, and professional services such as data center migration." AWS has thousands of organizations in the AWS Partner Network (APN) including system integrators, consulting firms, and independent software vendors (ISVs). AWS Marketplace, an online software store, helps customers search over 2,300 listings to buy and immediately start using software that runs on AWS.

- AWS Cloud Security Authorizations and Experience: AWS offers customers a powerful cloud security capability based on cutting-edge security experience and backed by an extensive repertoire of accreditations and authorizations. In The Forrester Wave™:

Public Cloud Platform Service Providers' Security, Q4 2014 report, Forrester Research named AWS as the only provider in the Leader category. Forrester stated, "AWS leads the pack. AWS demonstrated not only a broad set of security capabilities in data center security, certifications, and network security, but also excelled in customer satisfaction, security services partnerships, and a large installed base." AWS has achieved two Provisional Authorizations to Operate (P-ATOs) for mission systems designated by DISA as Cloud Computing Security Requirements Guide (SRG) level 2 (covering all AWS regions in the contiguous United States [CONUS]) and SRG level 4 (covering only the AWS GovCloud (US) Region).

- AWS Pricing: As AWS's cloud computing infrastructure grows, it gains efficiency and economies of scale, which we pass on to our customers in the form of lower prices. The Gartner, Inc. 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS), Worldwide report states that AWS has "over 10 times more cloud IaaS compute capacity in use than the aggregate total of the other 14 providers," demonstrating how AWS's massive economies of scale make it possible to lead the cloud market in lowering prices. The AWS strategy of pricing each service independently gives customers tremendous flexibility to choose the services they need for each project and to pay only for resources used. The economies of scale available with the cloud, and the massive scale at which we operate, allows AWS to constantly purchase and refresh large volumes of infrastructure at very low cost. Consequently, AWS customers reap the benefits of decreased IT costs such as better performance through improved quality and availability of IT infrastructure and enhanced functionality through system-wide innovation in the AWS IaaS platform.

BUSINESS BENEFITS OF AWS CLOUD SERVICES

There are additional business benefits that AWS cloud services can help customers realize. A few of these are listed here:

- Almost Zero Upfront Infrastructure Investment: If a customer wants to build a large-scale system, it may cost a fortune to invest in real estate, physical security, hardware (racks, servers, routers, backup power supplies), hardware management (power management, cooling), and operations personnel. Because of the high upfront costs, the project would typically require several rounds of management approvals before the project could even get started. With AWS's utility-style cloud computing, there is no fixed cost or startup cost.

- Just-In-Time Infrastructure: In the past, if an application became popular and a business' systems or infrastructure did not scale, it became a victim of its own success. Conversely, if a developer invested heavily and did not get popular, it became a victim of failure. By deploying applications in the AWS cloud with just-in-time self-provisioning, customers do not have to worry about pre-procuring capacity for large-scale systems. AWS's cloud increases agility, lowers risk, and lowers operational cost, because customers can scale cloud resources as they grow and only pay for what they use.

- More Efficient Resource Utilization: System administrators usually worry about procuring hardware (when they run out of capacity) and higher infrastructure utilization (when they have excess and idle capacity). With AWS, they can manage resources more effectively and efficiently by having the applications request and relinquish resources on-demand.

- Usage-Based Costing: With utility-style pricing, AWS customers are billed only for the infrastructure that has been used. AWS customers do not pay for allocated but unused infrastructure. This adds a new dimension to cost savings, allowing customers to see immediate cost savings when they deploy an optimization patch to update their cloud application. For example, if a caching layer can reduce data requests by 70%, the savings begin to accrue immediately. Moreover, if customers build platforms on the cloud, they can pass on the same flexible, variable usage-based cost structure to their own customers.

- Reduced Time to Market: Parallelization is the one of the great ways to speed up processing. If one compute-intensive or data-intensive job that can be run in parallel takes 500 hours to process on one machine, with cloud architectures, it would be possible to spawn and launch 500 instances and process the same job in 1 hour. Having available an elastic infrastructure provides the application with the ability to exploit parallelization in a cost-effective manner reducing time to market.

TECHNICAL BENEFITS OF AWS CLOUD SERVICES

Some of the key technical benefits of the AWS cloud are:

- Automation – "Scriptable Infrastructure": AWS customers can create repeatable build and deployment systems by leveraging programmable (API-driven) infrastructure.

- Auto Scaling: AWS customers can scale their applications up and down to match unexpected demand without any human intervention. Auto Scaling encourages automation and drives more efficiency.

- Proactive Scaling: Customers can scale applications up and down to meet anticipated demand with proper planning of traffic patterns so that costs remain low while scaling.

- More Efficient Development Lifecycle: Production systems may be easily cloned for use as development and test environments. Staging environments may be easily promoted to production.

- Improved Testability: Never run out of hardware for testing. Inject and automate testing at every stage during the development process. AWS customers can spin up an "instant test lab" with pre-configured environments only for the duration of testing phase.

- Disaster Recovery and Business Continuity: The cloud provides a lower cost option for maintaining a fleet of disaster recovery servers and data storage. With the cloud, customers can take advantage of geo-distribution and replicate the environment in other locations within minutes.

- Overflow Traffic to the Cloud: With a few clicks and effective load balancing tactics, customers can create a complete overflow-proof application by routing excess traffic to the cloud.

- Analyst Reports:[1] Gartner, Inc., a leading information technology research company, reported in its 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS), Worldwide report that "AWS is a thought leader; it is extraordinarily innovative, exceptionally agile, and very responsive to the market. It has the richest array of IaaS features and PaaS-like capabilities. It continues to rapidly expand its service offerings and offer higher-level solutions." The Gartner Magic Quadrant for May 2015 (Figure 5) depicts AWS in the Leaders Quadrant.

Additionally, Gartner positions AWS in the Leaders Quadrant of the new Magic Quadrant for Public Cloud Storage Services (Figure 6). Gartner defines leaders as offering innovative storage offerings built on a hardened platform, with global data centers and established credibility as a business.

- The Forrester Wave: Public Cloud Platform Service Providers' Security, Q4 2014 report (Figure 7) evaluated four of the leading public clouds along 15 key security criteria, detailing the findings about how well each vendor fulfilled their criteria and where they stand in relation to each other. Forrester's evaluation states "AWS leads the pack. AWS demonstrated not only a broad set of security capabilities in data center security, certifications, and network security, but also excelled in customer satisfaction, security services partnerships, and a large installed base. AWS led with the size of its development and technical support staff as well."

[1] More analyst reports can be found at http://aws.amazon.com/resources/analyst-reports/.

Figure 5. 2015 Gartner Magic Quadrant for Cloud Infrastructure as a Service

Figure 6. 2015 Gartner Magic Quadrant for Public Cloud Storage Services

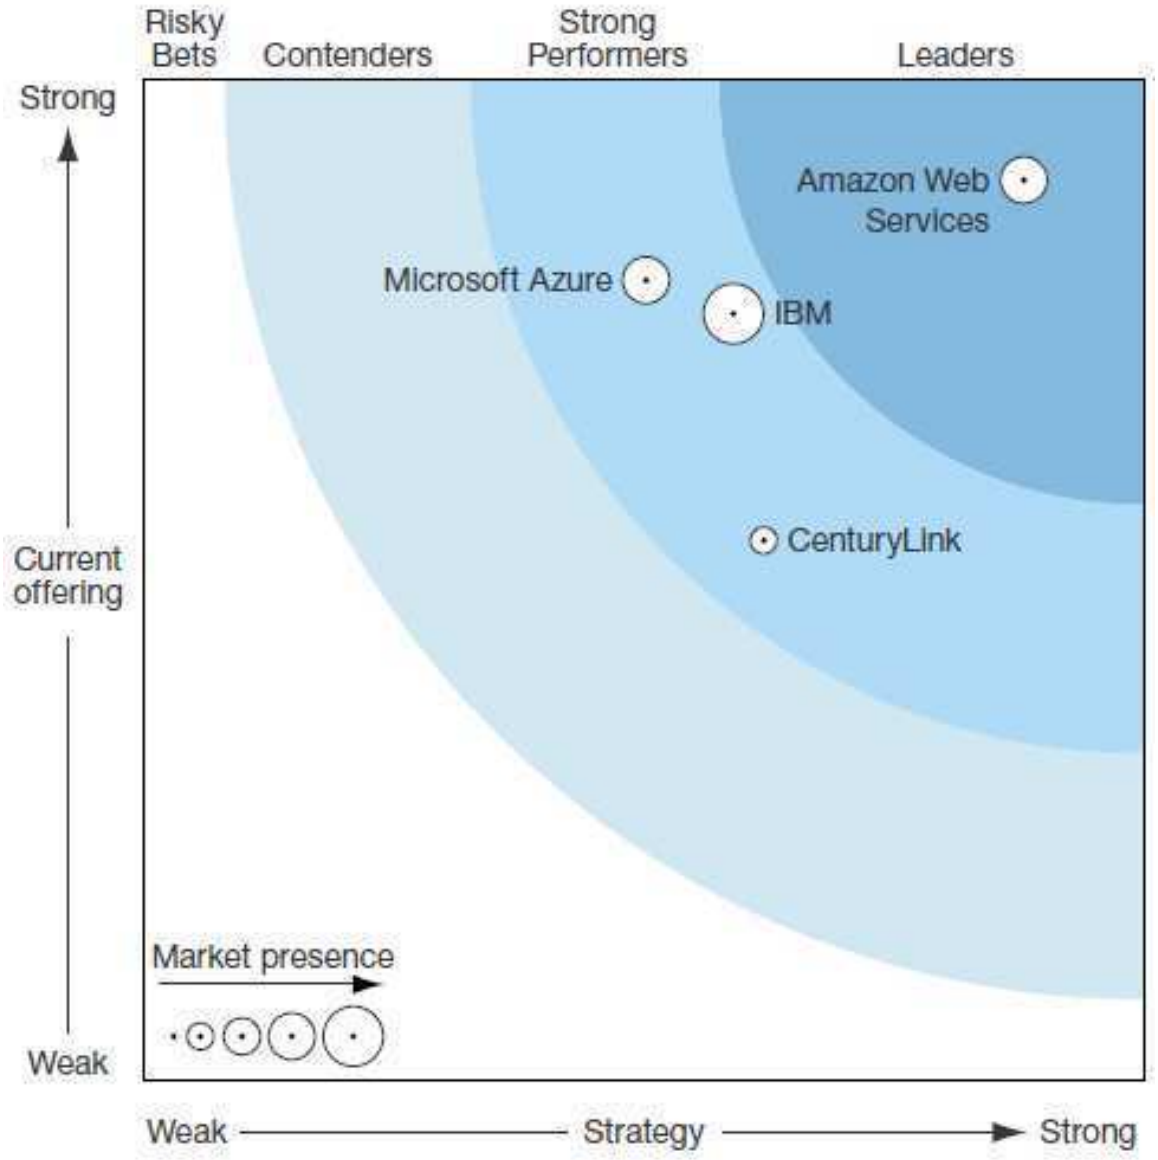

Figure 7. Forrester Wave™: Public Cloud Service Providers' Security, Q4 '14

AKAMAI

Akamai is the leading content delivery network and cloud security platform for helping enterprises across the globe to provide secure, high-performing user experiences on any device, anywhere. At the core of the Company's solutions is the Akamai Intelligent Platform providing extensive reach, coupled with unmatched reliability, security, visibility, and expertise. Akamai removes the complexities of connecting the increasingly mobile world, supporting 24/7 consumer demand, and enabling enterprises to securely deliver applications over the Internet. Akamai is the most "Trusted Choice for Online Business." As a result, Akamai enables secure online operations for over 5,500 top Global customers around the world. Akamai delivers almost 30% of all web traffic for one out of every three Global 500 companies, 10 of the top 10 global banks, all 15 U.S. Cabinet level agencies, all branches of the U.S. Military, over $300 Billion in secure e-commerce transactions, all top 20 global eCommerce sites, 97 of the top 100 online U.S. retailers, 29 of the

top 30 global media and entertainment companies, 9 out of 10 top social media sites, all of the top anti-virus companies, and multiple U.S. Federal and State Health Information Exchange Portals.

GOOGLE

Google's SaaS offering, Google Apps, has been commercially available since 2006 and presently has ~5.3 million paying customers covering more than 26 million end-user licenses. Google Apps represents a combination of the most popular consumer products that have been prepared for Enterprise use through the means of adding Administrative and Compliance Controls to the suite.

In the government space, the States of Wyoming, Utah, Colorado, Maryland and Iowa have selected Google Apps as their Messaging and Collaboration Platform. Several Federal agencies and several dozen Cities and Counties are also successfully using Google Apps.

4.4.2 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

The AWS cloud infrastructure has been designed and is managed in alignment with regulations, standards, and best practices, including:

- 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 [SSAE] 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)

Additional information on all of the security regulations and standards with which AWS complies is available on the Amazon website.[2]

AKAMAI

Akamai does not provide origin data center hosting services, but maintains its server infrastructure in Data Centers that leverage SAS70 Tier II and greater standards. Our DDoS Mitigation subsidiary Prolexic recently announced that it has successfully completed its Type 2 SOC 1 examination, commonly referred to as SSAE (Statement on Standards for Attestation Engagements) 16, and its Type 2 SOC 2 examination, formally known as a Reporting on Controls at a Service Organization Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy. Prolexic also simultaneously completed the ISAE 3402 examination, which is the international equivalent of SSAE 16. Akamai also maintains PCI Level 1, HIPAA Title II Subtitle F, Section 261-264 compliance as well as FedRAMP Level II JAB P-ATO Compliance for the entire Akamai platform offered in this procurement.

GOOGLE

Google Apps auditing capabilities are in line with these requirements. Google provides a SSAE 16 report and a SOC 2 for Google Apps customers upon request. Google Apps also maintains a FedRAMP Authorization to Operate and ISO 27001 certification. Information about the security regulations and standards with which Google Apps complies is continually updated and available on Google's website.[3]

Google Apps for Work have been deemed FedRAMP compliant (Figure 8). The most current information can be found on the FedRAMP website.[4]

[2] https://aws.amazon.com/compliance/
[3] https://docs.google.com/a/google.com/viewer?url=https://services.google.com/fh/files/blogs/btd-sec-op-2014-grey.pdf
[4] https://www.fedramp.gov/marketplace/compliant-systems/google-google-services/

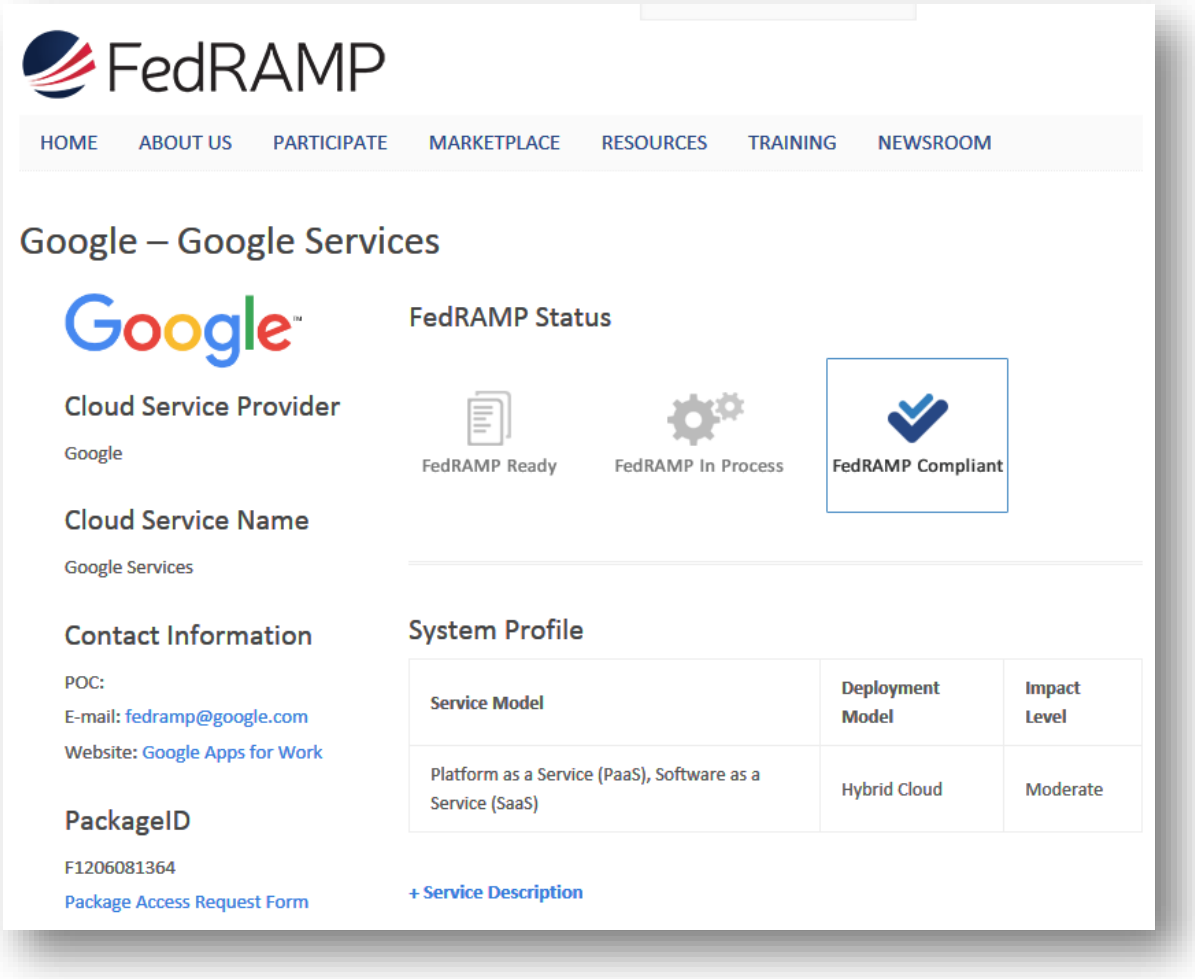

Figure 8. Google's FedRAMP Compliance Status (March 9, 2016)

4.5 BILLING AND PRICING PRACTICES [RFP REFERENCE 6.5]

4.5.1 RFP Reference 6.5.1

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

DLT Response:

AMAZON WEB SERVICES

With AWS, customers can incorporate a utility-style pricing model, only paying for the resources consumed. AWS continues to lower the cost of cloud computing for its customers. In 2014, AWS reduced the cost of compute by an average of 30%, storage by an average of 51%, and relational databases by an average of 28%. AWS continues to drive down the cost of customer IT infrastructure. AWS's utility-style pricing model is explained below:

- Pay as You Go: No minimum commitment or long-term contract is required. Customers can turn off cloud resources and stop paying for them when they are not needed, maximizing Return on Investment (ROI) through full utilization.

- Pay Less When You Reserve: For certain AWS products, customers can invest in reserved capacity, paying a low up-front fee to receive a significant discount. This results in overall savings of up to 60% (depending on the type of instance reserved) over the equivalent on-demand capacity.

- Pay Even Less Per Unit by Using More: AWS pricing is tiered for storage and data transfer, so the more customers use, the less they pay per gigabyte.

- Pay Even Less as AWS Grows: We continually focus on reducing our data center hardware costs, improving our operational efficiencies, lowering our power consumption, and passing savings back to customers. AWS has a history of continually lowering prices and has reduced prices 51 times since AWS launched in 2006.

- Transparency: AWS provides transparent, publicly available, and up-to-date pricing and tools that allow customers to evaluate AWS pricing against other Cloud Service Providers (CSPs) using the Simple Monthly Calculator.[5]

- Governance: AWS provides tools to generate detailed and customizable billing reports to meet customer business and compliance needs. Additionally, AWS Partner Network (APN) Consulting Partners can help customers manage and control cost utilization/tracking tools in order to provide customized billing reports.

The AWS Total Cost of Ownership (TCO) Calculator allows organizations to compare AWS to the cost of running applications in an on-premises or traditional hosting environment. The AWS Whitepaper *The Total Cost of (Non) Ownership of Web Applications in the Cloud* [6] also helps customers understand the different cost factors involved in moving to AWS. Our analysis shows that AWS offers significant cost savings (up to 80%) compared to equivalent on-premises deployments.

AWS is a highly cost-effective alternative to on-premises infrastructure solutions, delivering significantly reduced IT system and management costs. In 2015, AWS commissioned the International Data Corporation (IDC) to interview 11 organizations that deployed applications on AWS. IDC set out to understand the long-term economic implications of moving workloads onto Amazon cloud infrastructure services, the impact of moving applications on developer productivity and business agility, and the new opportunities that businesses could address by moving resources onto AWS.

[5] http://calculator.s3.amazonaws.com/index.html
[6] http://media.amazonwebservices.com/AWS_TCO_Web_Applications.pdf

IDC discovered that developing, deploying, and managing critical applications in AWS delivered a five-year TCO savings of 64.3% when compared with deploying the same resources on-premises or in hosted environments. The findings also showed a 560% ROI over five years.

AKAMAI

The Billing Center (also available via API) allows you to access your website's monthly usage over the Akamai network. One-time fees are billed in advance on the applicable Billing Effective Date. Monthly recurring fees are billed each calendar month in advance starting on the applicable Billing Effective Date. For the avoidance of doubt, the Term on the applicable Transaction Document begins on the Billing Effective Date.

GOOGLE

DLT Solutions published our Google Apps pricing on our contract price list. Customers are billed on per unit basis as stipulated in the pricing schedule. The payment plan is flexible. You can add and delete users at any time. You only pay for the service used during that month. You're billed monthly.

4.5.2 RFP Reference 6.5.2

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

DLT Response:

Please see section 4.5.1.

DLT works closely with our customers to help determine the best possible cloud solution to help optimized cost savings while maximizing the utility of the cloud technologies. Although this is a function best performed at the time of application migration or creation, DLT can support this effort in flight as well.

4.5.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

AWS's industry-leading security strength benefits you in many ways, one of which is by using a platform that is audited extensively by independent third-party assessors. At times, these audits confirm we can meet new requirements, even as they are issued, and this is the case for the National Institute of Standards and Technology (NIST) guidelines 800-171, which were released in June 2015. 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 we have already been audited under our 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 Appendix D of NIST Special Publication 800-171.[7]

AKAMAI

Akamai meets this characteristic by providing a basic monthly rate service plan and a metered bursting charge to provide flexibility to the Purchasing Entity based on scalable utilization. Akamai measurement tools will provide both the customer and the provider with an account of what has been used/consumed.

GOOGLE

Google Apps is optimized for unlimited use by all end users. Google ensures high availability, low latency and fault tolerance as a part of the contracted services. Metering is performed at the User level for this service.

4.6 SCOPE AND VARIETY OF CLOUD SOLUTIONS [RFP REFERENCE 6.6]

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.

DLT Response:

Please see Appendix 1.

4.7 BEST PRACTICES [RFP REFERENCE 6.7]

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.

DLT Response:

AMAZON WEB SERVICES

The AWS virtual infrastructure has been designed to provide optimum availability while ensuring complete customer privacy and segregation. AWS's highly secure data centers use state-of-the-art electronic surveillance and multi-factor access control systems and maintain strict, least-privileged-based access authorizations. Our environmental systems are designed to minimize the impact of disruptions to operations, and our multiple geographic regions and Availability Zones allow customers to remain resilient in the face of most failure modes, including natural disasters or system failures. AWS manages over 1,800 security controls to provide an optimally secure environment for all of our customers.

Additionally, network traffic between AWS regions, Availability Zones, and individual data centers travels over private network segments by default. These private network segments are fully isolated from the public Internet and not routable externally. AWS resources can be configured to

[7] http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-171.pdf

reside only on isolated AWS network segments and thus avoid utilizing any public IP addresses or routing over the public Internet.

AWS security engineers and solution architects have developed whitepapers and operational checklists[8] to help customers select the best options for their needs and to recommend security best practices, such as storing secret keys and passwords in a secure manner and rotating or changing them frequently.

BUILT-IN SECURITY FEATURES

Not only are applications and data protected by highly secure facilities and infrastructure, they are also protected by extensive network and security monitoring systems. AWS and its partners offer over 700 tools and features to help customers meet their security objectives concerning visibility, auditability, controllability, and agility. These tools and features provide basic but important security measures such as Distributed Denial of Service (DDoS) protection and password brute-force detection on AWS accounts. AWS-provided security features include:

- Secure Access: Customer access points, also called Application Programming Interface (API) endpoints, allow secure HTTP access (HTTPS) so that customers can establish secure communication sessions with their AWS cloud services using Secure Socket Layer (SSL)/Transport Layer Security (TSL).

- Built-In Firewalls: Customers can control how accessible their instances are by configuring built-in firewall rules—from totally public to completely private or somewhere in between. And when instances reside within an Amazon Virtual Private Cloud (Amazon VPC) subnet, customers can control egress as well as ingress.

- Unique Users: The AWS Identity and Access Management (IAM) tool allows AWS customers to control the level of access their own users have to AWS infrastructure services. With AWS IAM, each user can have unique security credentials, eliminating the need for shared passwords or keys and allowing the security best practices of role separation and least privilege.

- Multi-Factor Authentication (MFA): AWS provides built-in support for MFA for use with AWS accounts as well as individual AWS IAM user accounts.

- Private Subnets: The Amazon VPC service allows customers to add another layer of network security to instances by creating private subnets and even adding an Internet Protocol Security (IPsec) Virtual Private Network (VPN) tunnel between a home network and Amazon VPC.

- Encrypted Data Storage: Customers can have the data and objects they store in Amazon Elastic Block Store (Amazon EBS), Amazon Simple Storage Service (Amazon S3), Amazon Glacier, Amazon Redshift, and Amazon Relational Database Service (Amazon RDS) on Oracle and SQL Server encrypted automatically using Advanced Encryption

[8] http://aws.amazon.com/security/security-resources/

Standard (AES) 256, a secure symmetric-key encryption standard using 256-bit encryption keys.

- Dedicated Connection Option: The AWS Direct Connect service allows customers to establish a dedicated network connection from their premises to AWS. Using industry-standard 802.1q VLANs, this dedicated connection can be partitioned into multiple logical connections to enable access to both public and private IP environments within the AWS cloud.

- Isolated GovCloud: For customers who require additional measures in order to comply with US International Traffic in Arms Regulations (ITAR), AWS offers an entirely separate region called AWS GovCloud (US). This isolated region provides an environment where customers can run ITAR-compliant applications and provides special endpoints that utilize only Federal Information Processing Standard (FIPS) 140-2 encryption.

- Dedicated, Hardware-Based Crypto Key Storage Option: For customers who must use Hardware Security Module (HSM) appliances for cryptographic key storage, AWS CloudHSM provides a highly secure and convenient way to store and manage keys.

- Centralized Key Management: For customers who use encryption extensively and require strict control of their keys, the AWS Key Management Service (KMS) provides a convenient management option for creating and administering the keys used to encrypt data at rest.

- Perfect Forward Secrecy: For even greater communication privacy, several AWS cloud services such as Elastic Load Balancing and Amazon CloudFront offer newer, stronger cipher suites. These cipher suites allow SSL/TLS clients to use Perfect Forward Secrecy, a technique that 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.

Several of AWS's built-in cloud security features focus on providing visibility into data, performance, and resource usage. The tools listed below help customers gain more insight into their cloud operations, giving them the means to better control their security and providing information for data-driven decisions.[9]

- AWS Trusted Advisor: Provided automatically when AWS customers sign up for premium support, the AWS Trusted Advisor service is a convenient way for customers to see where they could use a little more security. It monitors AWS resources and alerts customers to security configuration gaps such as overly permissive access to certain Amazon Elastic Compute Cloud (Amazon EC2) instance ports and Amazon S3 storage buckets, minimal use of role segregation using AWS IAM, and weak password policies.

[9] More information on these and other features is available at http://aws.amazon.com/security/aws-security-features/.

- Amazon CloudWatch: Amazon CloudWatch enables customers 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 a customer's applications and services and any log files their applications generate. Customers can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health, using these insights to react intelligently and keep applications running smoothly.

- AWS CloudTrail: AWS CloudTrail provides logs of all user activity within an AWS account. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS cloud service. The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

- AWS Config: With the AWS Config service, customers can immediately discover all of their AWS resources and view the configuration of each. Customers can receive notifications each time a configuration changes as well as dig into the configuration history to perform incident analysis.

AKAMAI

Many organizations find it difficult to maintain compliance due to a lack of resources and expertise and would benefit from having a partner in security compliance. Akamai proactively maintains compliance of our SSL network with PCI standards. As the operators of the world's largest content distribution network (CDN), Akamai has deep expertise in security compliance and is able to help our CDN clients meet their own compliance challenges.

Akamai's network, management infrastructure, and associated processes and procedures are consistent with the best practice security requirements of PCI, ISO, BITS, FISMA, and HIPAA. Akamai's customers leveraging the Akamai Intelligent Platform are able to significantly reduce their compliance burden, achieving the following key benefits:

- Reduced risk: Pre-established certification of our SSL network and our best practice guidelines reduce risk for our customers and expedite the overall compliance process.

- Time and cost savings: Our Compliance Management modules help simplify the process of meeting information security reporting and auditing requirements.

- Accelerated validation: Our Compliance Management solution provides our clients with the services, documentation and service terms they need to facilitate fast compliance validation.

Akamai has developed a variety of innovative solutions ranging from media and delivery services including video transcoding and analytics to web experience solutions like our Dynamic Site Accelerator. With the support of our dedicated Information Security (InfoSec) and Computer

Security Incident Response Teams (CSIRT), we are constantly working to make the web a safer place to work, learn and do business.

GOOGLE

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.

Google undergoes audits against key Security and Privacy Controls that cover Corporate Security Policies, Operation Security, Access Control, Organization Security, Asset Classification and Control, Personnel Security, Systems Development and Maintenance, Disaster Recovery and Business Continuity and Regulatory Compliance.

The results of these audits have led to SOC2, ISO2001 and FedRAMP Certifications. The audit reports are confidential records that can be reviewed by interested parties under Non-Disclosure Agreement (NDA).

These security assertions are also backed by contractual obligations in the Google Master Terms of Service and Data Processing Amendment (please refer to Appendix 6) together which outline such things as Customer ownership of their data, confidentiality of all customer data, restricted scope of data use pertaining to the delivery of the contracted services only (meaning no advertising profiles) and the proper handling of data during the contract period and following the contract period.

5.0 ORGANIZATION PROFILE AND STAFFING

PLEASE NOTE: DLT has provided complete responses to all RFP requirements in the text of this proposal that fully explain how we meet or exceed the minimum mandatory requirements. We have also provided footnotes containing links to additional information that is supplemental to our response but may be of interest to evaluators—the information available at these links is not required for evaluation but instead is meant as reference material.

5.1 CONTRACT MANAGER [RFP REFERENCE 7.1]

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.**

5.1.1 RFP Reference 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.

DLT Response:

DLT's Contract Manager for the NASPO ValuePoint Master Agreement is:

Daryl Savage
Vice President, Contracts, and General Counsel
Phone: 703-708-9631
Email: Daryl.Savage@dlt.com
Office Hours: 9:00AM to 6:00PM Eastern, Monday through Friday

5.1.2 RFP Reference 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.

DLT Response:

As DLT's Vice President of Contracts, Mr. Savage has management oversight for all of DLT's contracting vehicles, such as:

- GSA IT Schedule 70 Contract Number GS-35F-4543G
- National Institutes of Health (NIH) Chief Information Officers – Commodities and Solutions (NIH CIO-CS)
- Electronic Commodities Store III (ECS III)
- NASA SEWP V
- U.S. Communities (Oracle products)
- Internet2 (I2)
- National Cooperative Purchasing Alliance (NCPA) - Web Services and Cloud Computing
- The Interlocal Purchasing System (TIPS)
- Texas DIR for Amazon Web Services (TX-DIR-CLOUD)

Please see Mr. Savage's resume on page 37.

5.1.3 RFP Reference 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.

DLT Response:

Mr. Savage has complete authority to commit and manage the corporate resource(s) necessary for the overall contract administration (to include proposals, contract compliance, task/delivery order administration, etc.) of the NASPO ValuePoint Master Agreement.

Mr. Savage will:

- Act as the single point of contact between DLT and NASPO ValuePoint on contractual issues;
- Ensure that DLT's published ordering instructions are clear, complete, and current;
- Provide training and guidance to DLT's sales and support teams regarding Master Agreement terms and conditions;
- Ensure that DLT's task/delivery order administration and fulfillment meet or surpass NASPO's expectations and requirements;
- Monitor the Report Usage Administrator's activities to ensure required reports are produced timely and correctly;
- Ensure that contract fees are reported correctly and paid timely;
- Negotiate Participating Entity Addendums as needed;
- Negotiate all changes to the Master Agreement, such as adding vendors and products;
- Ensure that the proper procedures are followed in the event of a data breach (see section 6.3.1);
- Resolve any delivery or billing issues that may arise; and,
- Attend scheduled and ad hoc meetings as requested by NASPO to review the state of our contract and to discuss issues with and improvements to our technical and administrative processes.

Daryl D. Savage, Esq.

Vice President, Contracts and General Counsel

SUMMARY OF QUALIFICATIONS

Corporate legal professional with diverse experience in a broad range of business activities including government contracts, cloud services, hardware and software productizing, telephony, software licensing and compliance, and labor and employment issues. Track record of success in providing practical, business oriented, innovative and responsive legal advice to senior business management in all areas of operations, including foundational business and operational elements (sales processes, contracts and renewals), establishment of business strategies for product development, marketing campaigns and enterprise risk management. Strong negotiation and contract drafting credentials. Licensed to practice law in Pennsylvania and Virginia.

RELEVANT SKILLS

- Public contracting with emphasis in Schedule 70 and GWAC vehicles;
- Supply chain management including negotiation of hardware, software and services agreements;
- Labor and employment;
- Government contract compliance;
- Mergers & acquisitions;
- Contracts and Uniform Commercial Code (UCC);
- Telephony and information services;
- Information security and data privacy regulation;
- Intellectual property, licensing & product distribution;
- Litigation management and strategy;
- Trade regulation law and compliance internet law; and,
- Negotiating or drafting alliance, teaming, joint venture and non-disclosure agreements.

EXPERIENCE SUMMARY

- Rebranded and rebuilt DLT legal/contracts department into a sales oriented, problem-solving team of nine contracts professionals;
- Drafted DLT Storage-as-a-Service (STaaS), Managed Services and Cloud Services Templates;
- Implemented DLT compliance program including drafting Business Ethics, ITAR, Document Retention, Data Privacy and e-Discovery Guidelines;
- Developed and implemented c-Sight and Contracts Dashboard providing management with real-time insight into contracts department resource utilization and performance achievement;
- Successfully prosecuted several contract claims resulting in recovery of $20 million;
- Successfully negotiated resolution of $52 million claim asserted by Microsoft for alleged license non-compliance under terms that improved CSC's competitive position with no cash outlay by CSC;
- Negotiated terms of Binding Letter of Commitment allowing CSC to obtain an advance payment of $38 million from MCI to CSC just prior to the close of the fiscal year;
- Simultaneously negotiated two network services transactions with Level 3 and Verizon in support of $2.1 billion outsourcing agreement;
- Counsel on $360 million Network Services outsourcing agreement and upon conclusion negotiated support agreements with 7 critical vendors to support CSC's outsourcing obligations;

- Developed CSC Cloud Computing Terms of Services Agreement for CSC's Microsoft BPOS offering; and,
- Developed data security and privacy program to comply with state and federal regulations/guidelines.

EXPERIENCE

Vice President, Contracts and General Counsel / DLT Solutions, LLC / Herndon, VA / March 2010 – Present

Serves on the Senior Management Team as Vice President of Contracts, General Counsel and Compliance Officer for $890M Value Added Reseller (VAR) of hardware, software and services. Duties include providing counsel and advice regarding GSA Schedule 70 and GWAC vehicles, serving as counsel for DLT Cloud Advisory Group, contract negotiation, developing and implementing government contract compliance programs and guidelines, managing software licensing and compliance, ensuring compliance with antitrust and resale price maintenance laws, reviewing marketing and sales materials, drafting and implementing compliance policies, guidelines and leading or supporting merger and acquisition activity. Manages all litigation including governmental investigations, bid protests and government contract claims. Develops and manages the legal department budget with an eye toward increasing efficiency and where appropriate, bringing legal services in-house to reduce expenses. Implemented procedures and management systems to increase legal department responsiveness including development of metrics and implementation of a Contracts Dashboard.

Assistant General Counsel / Computer Sciences Corporation (CSC) / Falls Church, VA / January 1994 – March 2010

Reported to the General Counsel of CSC. Provided advice and counsel on government contract matters including contract performance, statutory and regulatory compliance, and corporate policy and training, with particular emphasis in FAR Part 12 issues, contract cost principals, data rights, terminations, changes and disputes. Led CSC public sector activity related to the Next Generation Network Priority Service (NPS) and Cloud Computing. Provided advice and counsel related to federal and commercial telecommunications transactions. Provided advice and counsel in the development, negotiation and problem resolution for CSC Cloud Computing customer and vendor facing agreements for both public entity and commercial customers. Provided advice and counsel, including negotiation support, for enterprise software (including Symantec, IBM, Red Hat, SAS, SAP, Oracle, Microsoft, Computer Associates, BMC, Google, Amazon); hardware (IBM, HP, Cisco, Avaya, EMC); telecommunications (Verizon, AT&T, Sprint, BT, Telmex); and corporate services agreements (Marquis Jet, Xerox). Developed and updated commercial licensing terms for CSC's Catalyst Business Process and ERP Methodology. Drafted CSC privacy policy and policy statements on CSC website. Provided advice and counsel regarding privacy issues relating to CSC operations as well as bid and proposal activity.

General Counsel / ARC Professional Services Group, Inc. / Rockville, MD / January 1993 – December 1993

Provided advice and counsel on a broad range of general corporate matters including contract performance, statutory and regulatory compliance, corporate policy and training, with particular emphasis in government contracts including contract cost principles, data rights, terminations, changes and disputes. Drafted bid protests to GAO and GSBCA, as well as interested party responses. Provided in-

house expertise with respect to Service Contract Act, Fair Labor Standards Act, Contract's Disputes Act and Freedom of Information Act questions. Assisted with acquisition of ARC PSG by Computer Sciences Corporation.

Counsel / Atlantic Research Corporation / Alexandria, VA / October 1987 – January 1993

Associate / Leftwich Moore and Douglas / Washington, D.C. / February 1986 – October 1987

EDUCATION & PROFESSIONAL ORGANIZATIONS

Juris Doctor / Rutgers University School of Law / Camden, NJ / 1981
Bar Memberships: Pennsylvania, Virginia, U.S. Court of Federal Claims and U.S. Supreme Court
Associations: National Bar Association, Corporate Counsel Association, National Contract Management Association

6.0 TECHNICAL RESPONSE [RFP REFERENCE 8]

PLEASE NOTE: DLT has provided complete responses to all RFP requirements in the text of this proposal that fully explain how we meet or exceed the minimum mandatory requirements. We have also provided footnotes containing links to additional information that is supplemental to our response but may be of interest to evaluators—the information available at these links is not required for evaluation but instead is meant as reference material.

6.1 TECHNICAL REQUIREMENTS [RFP REFERENCE 8.1]

6.1.1 RFP Reference 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.**

DLT Response:

Team DLT is composed of industry leaders representing a new breed in the Cloud Service Provider (CSP) model. Team DLT member Day1 Solutions is one of only four companies to currently hold and manage the current Western States Contracting Alliance (WSCA) Public Cloud Hosting Services contract from its inception. Team DLT has extensive experience guiding state, local and educational institutions on best practices when adopting cloud solutions. This extensive experience allows for Team DLT to translate all of our lessons learned from managing, operating, and delivering on the original WSCA contract, as well as our many other contracts listed in section 4.2, to ensure success of the NASPO ValuePoint Master Agreement for Cloud Solutions.

Team DLT will continue to leverage a combination of Amazon Web Services (AWS) and our organizationally approved cloud solutions from our existing ValuePoint partners to further NASPO ValuePoint's mission to bring best value, innovation and competition in the marketplace for all eligible Purchasing Entities.

NASPO requires a business partner with strong cloud services capabilities to guide, oversee, promote, expand and verify the activities associated with this contract vehicle over an extensive period of time. Team DLT is purposely built to serve the needs of state, local, and university IT systems and business requirements and will stay agile, focused and committed to delivering quick measurable results. Our combined elite organizational vendor certifications, business qualifications, and past performances provide NASPO with the necessary assurance that Team DLT has the appropriate certifications, experience, and knowledge required to deliver the highest quality technical and professional cloud-based services.

AMAZON WEB SERVICES

As an AWS Authorized reseller and Premier Consulting Partner, Team DLT provides NASPO with direct access to the expansive catalog of services and solutions provided by AWS. AWS offers scalable, cost-efficient cloud services that NASPO can use to meet governance and security mandates, reduce costs, drive efficiencies, and accelerate innovation for its customers. Additionally, Team DLT holds the necessary partnership levels with AWS to guarantee a full scope of cloud services and solutions beyond resale, which include *Channel Reseller Partner*, *Premier Consulting Partner, Managed Services Partner, Direct Connect Partner, Authorized*

Government Partner, and *Big Data Competency*. The impact of these partnership levels with AWS enables our organization to provide comprehensive end-to-end cloud-based solutions through an operating model that allows Team DLT to serve as a Value Added Reseller (VAR), Systems Integrator (SI), and a Managed Service Provider (MSP) for AWS. Our achievements with these partnership levels demonstrate our ability to provide NASPO with certified staff, organizational competencies, and exclusive experience as an AWS cloud service provider. Leveraging AWS, we can provide NASPO with the ability to implement all types of cloud, whether it's Private, Public, Hybrid and Community cloud.

Team DLT will offer IaaS, SaaS and PaaS service models with the capability to offer these in private, public or hybrid cloud environments. Today we offer all service models in partnership with AWS core framework (compute, storage, and network). We then offer our clients the ability to layer in SaaS-based solutions from our vendor partners in the AWS Marketplace. Finally, we provide PaaS-based offerings such as Lambda, Code Deploy, EC2 Containers to name a few. All of these service models may be available in Public, Private or Hybrid models.

Team DLT has extensive experience customizing cloud architectures for state, local, federal governments and enterprises and will bring our experiences to bear as we guide NASPO and its Purchasing Entities in developing custom cloud solutions based on the deployment model of each customer. The cloud deployment models are described below:

1. Public: A public cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

2. Community: Community cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations).

3. Hybrid: Hybrid cloud infrastructure is a composition of two or more cloud models (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds). Additionally, hybrid solutions may also integrate with on-premise solutions.

4. Private: The cloud infrastructure is operated solely for an organization. Data, information, and security for this cloud model is strictly controlled, processed, and stored in the private infrastructure.

The Team DLT Business Model leverages our partnership levels with AWS and combines the elements of a Value Added Reseller (VAR), Systems Integrator (SI) and Managed Service Provider (MSP). For NASPO, this business model guarantees a simplified approach to cloud delivery from procurement, design, integration, and through operations and maintenance with the use of a highly accredited cloud service provider. Additionally, our business model provides NASPO with decreased administrative overhead and management by ensuring a single point of contact for all cloud services, increased accountability through a single vendor, and complete end-to-end cloud services.

- Value Added Reseller (VAR): As a Channel Reseller Partner, Team DLT can provide NASPO direct access to the expansive catalog of services within AWS. Additionally, our extensive experience and partner level as a reseller provides NASPO with an organization adept at the complexities of vendor management with AWS.

- Systems Integrator (SI): Our competencies as a Premier Consulting Partner (AWS's highest level) ensure that Team DLT can provide NASPO with a high level of consulting, architecture, design, engineering and other delivery-related professional services. Team DLT has an abundance of AWS certified staff who continue to investigate new cloud offerings and pursue new solutions for IT efforts. Our continuous evolution of cloud services will allow NASPO to quickly innovate to new cutting edge IT solutions.

- Managed Service Provider (MSP): Our designation as an AWS Managed Service Provider provides NASPO with a great level of assurance that Team DLT is appropriately equipped to manage all solutions developed within AWS. Our MSP practice ensures that the complexities of managing a cloud infrastructure are removed from NASPO as our teams will monitor operational support and maintenance to meet Service Level Agreements (SLAs).

Our approach is to be the primary reseller of AWS products and services to NASPO and its customers. We understand that NASPO requires direct access to AWS specific resources and we offer our ability to provide NASPO with the following benefits:

- Our costs are transparent to the end customer
- AWS price reductions are passed directly to customers immediately
- Our immediate access to AWS products ensures that NASPO eliminates delays in obtaining AWS services
- Our use of a master account provides transparent, comprehensive, and consolidated billing
- Our portals provide comprehensive usage reporting and dashboarding

AKAMAI

Akamai provides PaaS and IaaS cloud-based services in a Public Cloud deployment model. There is no additional, on site equipment. There are no additional or hidden charges to consider when purchasing Akamai services.

PAAS SERVICES

Web Content Delivery

- Akamai Dynamic Site Accelerator (DSA)/DSA Secure: Industry leading scale, performance and intelligence solutions to make web experiences fast, reliable and secure across all end-user situations while maximizing infrastructure offload and operational agility and standard and secure websites.

- Akamai Ion Standard/Ion Standard Secure: Industry leading web performance solution makes web experiences fast, reliable, and secure across all end-user situations, while maximizing infrastructure offload and operational agility.

- Cloudlets: Value added-apps that enhance your use of the Akamai Intelligent Platform each adding discrete functionality addressing a particular business or operational challenge.

Cloud DNS Services

- Fast DNS: DNS resolution that is fast, reliable and secure
- Fast DNS Secure Option: DNS Sec for Akamai Fast DNS Service
- Global Traffic Management: Business continuity with a cloud-based intelligent traffic manager
- NetStorage: Cloud based storage to support web content delivery
- SMP2.0: Professional Services package in support of Akamai services
- Professional Services-Enterprise (PSE): Statement of Work based professional services supporting Akamai services
- Enhanced Support SLA (eSLA): Akamai Support Services that promise faster response time for customer care support
- Integration Services: Statement of Work-based Professional Services Integration for Akamai services

IAAS/SECURITY/WEB SECURITY

- Kona Site Defender (KSD): Defends against all types of DDoS, web application and direct-to-origin attacks
- Kona DDoS Defender(KDD): DDoS defense to protect websites against the increasing scale, frequency and sophistication of attacks
- Prolexic Routed: DDoS defense for protecting data center infrastructure against large, complex attacks
- Prolexic Connect: DDoS defense for protecting data center infrastructure against large, complex attacks
- Managed Kona Site Defender: Managed security support of KSD
- Kona Rule Update Service: Security support for KSD
- Client Reputation: Defends against application layer and DDoS attacks by identifying malicious IP addresses dynamically and risk scoring the IPs
- Bot Manager: Provides a flexible framework to manage the wide array of bots accessing websites every day.

GOOGLE

The service model is SaaS for Google Apps with a Public Cloud deployment model.

6.1.2 RFP Reference 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:

DLT Response:

Team DLT understands the need for NASPO and all Purchasing Entities to have the ability to obtain cloud-based solutions that meet NIST essential characteristics as it provides an industry accepted baseline for cloud-based services. In sections 6.1.2.1 through 6.1.2.5, we have provided NASPO with a highlight of our understanding of each NIST characteristic plus an explanation of how our solution satisfies the NIST characteristics. All AWS services are FedRAMP compliant and meet the FISMA moderate level.

6.1.2.1 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Amazon Web Services, Inc. (AWS) provides customers of all sizes with on-demand access to a wide range of cloud infrastructure services, charging you only for the resources you actually use. AWS enables you 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, our customers can simply sign up for AWS and immediately begin deployment in the cloud. 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.

NIST characterizes On-Demand Self-Service as the ability for consumers to unilaterally and automatically provision computing capabilities to include server time and network storage without manual intervention for each service. Team DLT understands that self-service of IT resources is a compelling reason to leverage the cloud as it allows organizations to quickly provision IT resources without creating further deployment delays as characterized by an exhaustive procurement cycle.

How Team DLT's Cloud Solutions Satisfy NIST Characteristic and Capability: Team DLT asserts that Amazon Web Services, Inc. (AWS) provides NASPO with the ability to meet these requirements as AWS provides consumers of all sizes with on-demand access to a wide range of cloud infrastructure services.

How Self-Service Technical Capability Is Met: Team DLT will provide NASPO with the ability to leverage on-demand self-service to provision servers, storage, networking components, and bandwidth through the use of AWS Identity and Access Management (IAM). Additionally, the AWS Management console can be used by NASPO staff to access and manage AWS resources through a simple and intuitive web-based user interface secured through Secure Socket Layers (SSL) encryption. Team DLT staff has deep expertise with the different services available on AWS through a combination of lessons learned from previous engagements with customers, strong industry knowledge, understanding best practices, and constant organizational obligations for continued technical training. AWS specific offerings include global compute, storage, database, analytics, application, and deployment services. Figure 9 is a high-level description of the AWS cloud platform categories that Team DLT can bring immediately to NASPO.

DEPLOYMENT & MANAGEMENT		SECURITY & ADMINISTRATION
Services to help with management of credentials for access to AWS services, to monitor NASPO customer applications, to create and update stacks of AWS resource, deploy applications, use hardware security modules (HSMs) and log AWS API activity.		Services to help facilitate the security and administration of NASPO customer resources deployed in the AWS cloud. Implement controls to ensure an optimized shared security model.
APPLICATION		**MOBILE SERVICES**
A variety of managed services to use with organizational applications including services that provide application streaming, queuing, push notification, email delivery, and transcoding.		Unique services that facilitate and enable the development of mobile centric applications. Deploy, analyze, and test across multiple platforms.
DATABASE	**ANALYTICS**	**ENTERPRISE APPLICATIONS**
Fully managed relational and NoSQL database service, in-memory caching as a service and petabyte-scale data-warehouse service.	Cloud based analytics services to process and analyze any volume of data, whether it by managed Hadoop clusters, real-time streaming data, petabyte scale data warehousing, or orchestration.	A wide variety of enterprise level applications that provide NASPO customer with office automation capabilities.
NETWORKING	**COMPUTE**	**STORAGE**
A full range of networking services including logically isolated networks, private network connection to the AWS cloud, and highly available and saleable DNS service and deliver content to end users.	A wide selection of compute instances which can scale up and down automatically to meet the needs of NASPO customer applications, a managed load balancing service as well as fully managed desktops in the cloud.	Low-cost data storage with high durability and availability. AWS offers storage choices for backup, archiving, and disaster recovery, as well as block, file, and object storage.
AWS GLOBAL INFRASTRUCRURE		

Figure 9. AWS Services and Descriptions

AWS has a comprehensive catalog of products that are used to provide the services mentioned in Figure 9 above. As an AWS reseller, Team DLT will provide NASPO with direct access to these products and services without the complications of billing, transfer of services, and delay in

purchasing or obtaining these services through a separate reseller. A sampling of the specific AWS technologies and products that can be directly purchased through Team DLT can be found in Figure 10.

AKAMAI

Akamai's Luna portal allows for practically anything to be self serviceable from the Akamai Web Performance and Web Security product suites. Users can either log in to Luna or take advantage of the available APIs to achieve the required functionality. Akamai is FedRAMP accredited.

GOOGLE

Google Apps administrators have access to an Admin Panel that allows them to self perform user administration, service configurations, and reporting on demand. Additional tools that support this characteristic include the APIs in the Admin Software Developers Kit (SDK) that can automate many common tasks as well as an extensive set of third party tools developed by trusted Independent Software Vendors. Google Apps is FedRAMP accredited.

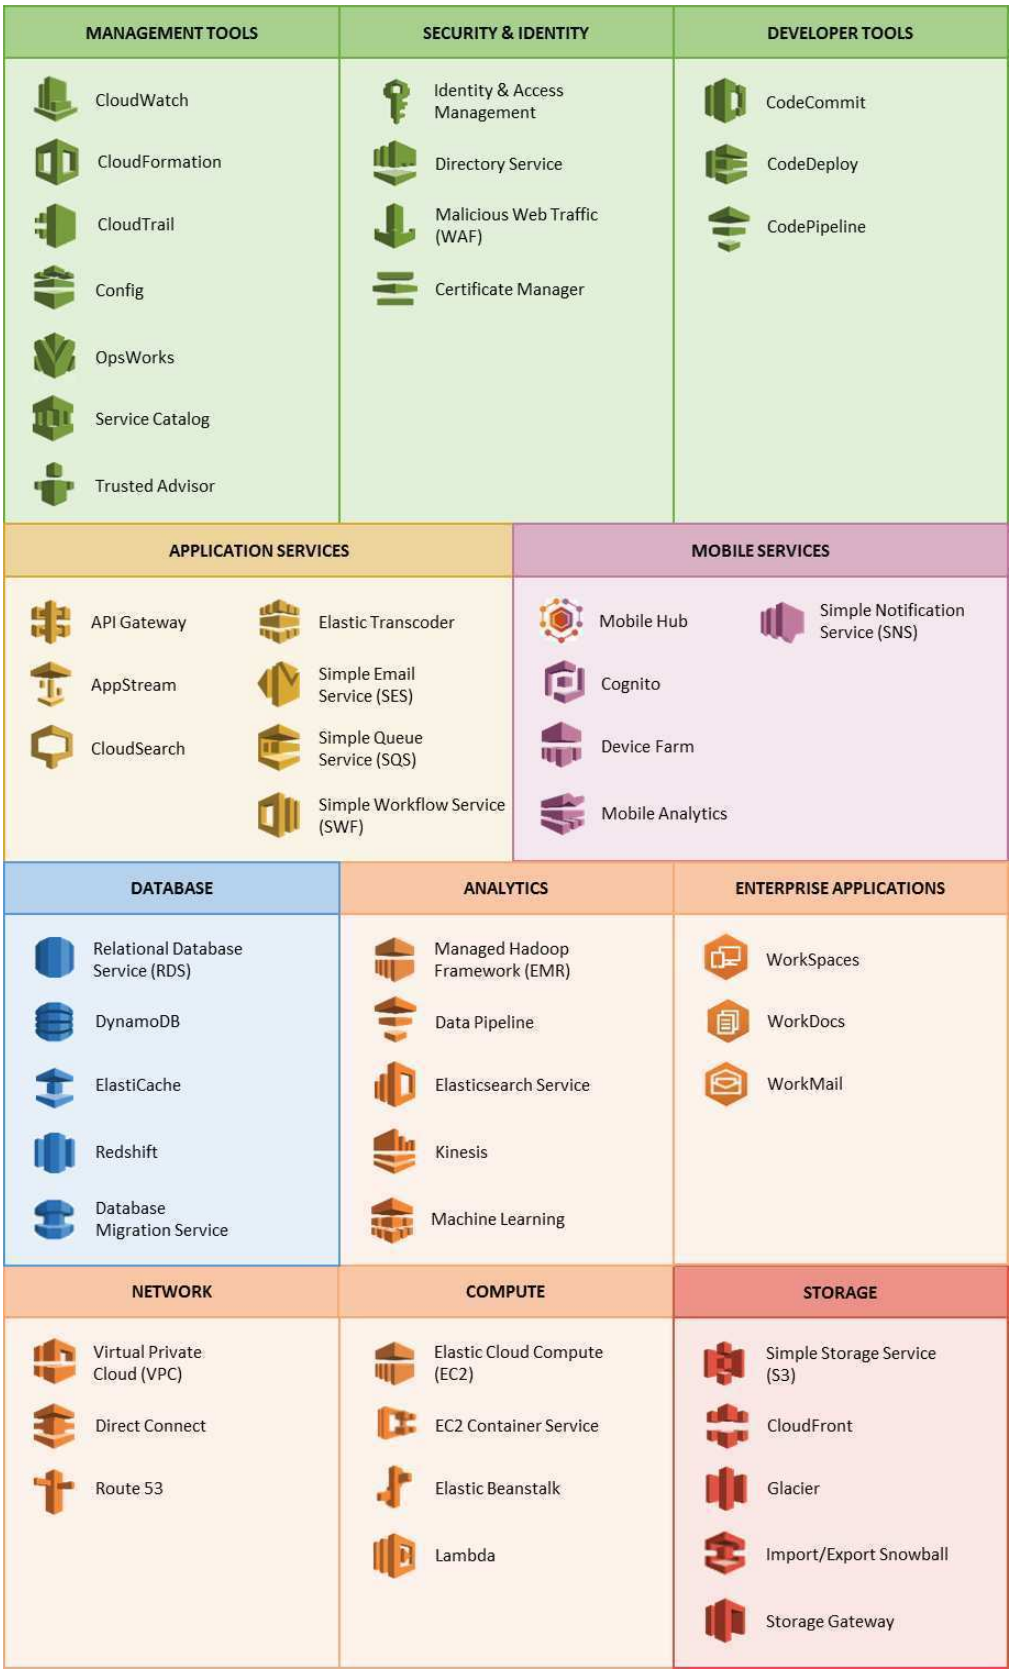

Figure 10. AWS Specific Technologies and Services

6.1.2.2 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

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, while you provision and use what you need via a web application, mobile client, or programmatically through published and well documented APIs.

NIST characteristics of Broad Network Access require that services are available over a network and accessed through standard mechanisms that encourage the use of heterogeneous platforms. Team DLT understands that this independence from geography and deployment of services that are easily network accessible provides NASPO Purchasing Entities with great flexibility in deploying, connecting, and accessing IT resources.

How Team DLT's Cloud Solutions Satisfy NIST Characteristic and Capability: Team DLT affirms that AWS complies with these characteristics as they provide a simple way to access servers, storage, databases, and a broad set of application services over the Internet through the AWS Management Console which provides customers with the ability to provision services needed via a web application, mobile client, command line access through Secure Shell (SSH) or programmatically through published and well-documented APIs. Additionally, AWS owns and maintains the network-connected hardware required for these application services.

AKAMAI

Akamai conforms to the Broad Network Access methodology for authorized users. The broad network access characteristic of the cloud makes it easy to take advantage of high-speed networks, fast DNS services, and CDN services of Akamai to speed up your application. This is especially important for mobile applications, where networks are slower, making every byte or millisecond of latency count. Akamai can help cache content closer to the end users, compress content, and make sure it supports conditional requests. All of these can significantly speed up the customer's website.

GOOGLE

Google has a very mature peering and content delivery policy that ensures customers have Broad Network Access to the Google infrastructure from the closest point of presence. The web-based services can be securely accessed from any browser, work on any mobile device such as Android or iPhone, and integrate with other popular email systems like Microsoft Outlook, Apple Mail, and more.

6.1.2.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

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 2.0 published in October 2010.

The main NIST characteristics of Resource Pooling require that “the provider’s computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.” We understand the significance of resource pooling to NASPO Purchasing Entities as it ensures greater efficiency in IT services and provides for economies of scale in pricing.

How Team DLT’s Cloud Solutions Satisfy NIST Characteristic and Capability: Team DLT is highly experienced in developing and architecting diverse cloud models and can provide NASPO customers with the ability to implement cloud models that provide resource pooling abilities based on the requirements of the customer. Leveraging AWS we can provide NASPO with a virtualized, multi-tenant environment that can be used by NASPO Purchasing Entities to support multiple cloud deployment models with specific resource pooling requirements that meet Public, Community, Hybrid, and Private Cloud requirements. 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.

Additionally, Team DLT has the ability to provide government customers with access to AWS GovCloud which is a community cloud designed to allow US government agencies and customers the ability to move sensitive workloads into the cloud by addressing specific regulatory and compliance requirements. The AWS GovCloud framework adheres to U.S. International Traffic in Arms Regulations (ITAR).

AKAMAI

Akamai CDN platform is the largest distributed platform in the world and conforms to the Resource Pooling methodology of NIST. Akamai’s computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact

location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., region). Examples of resources include storage, processing, memory etc.

GOOGLE

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, while at the same time replicating any data across at least two additional geographically dispersed data centers which also may rotate through the session.

6.1.2.4 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

AWS provides a massive global cloud infrastructure that allows you to quickly innovate, experiment, and iterate. Instead of waiting weeks or months for hardware, you can instantly deploy new applications, instantly scale up as your workload grows, and instantly scale down based on 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 then scale those resources down as demand decreases.

The NIST requirements for Rapid Elasticity state that resources should be provisioned and released elastically and may also require for automated scalability to meet rapid outward and inward demands. We understand that the majority of cloud adopters seek the rapid elasticity afforded through the cloud as it allows organizations to quickly scale up (or down) resources based on demand.

How Team DLT's Cloud Solutions Satisfy NIST Characteristic and Capability: AWS provides a massive global cloud infrastructure that allows for quick innovation, experimentation, and iteration of services through elasticity in services and capabilities. Traditional on-premise IT service models require weeks or months for hardware procurement and deployments that through AWS can be instantly deployed and can instantly scale up and down as workloads ebb and flow. AWS compute services, such as Elastic Cloud Compute (EC2), provides customers with a cloud solution that is both flexible and allows for NASPO Purchasing Entities to scale their applications based on the evolving demand for compute. Elastic Load Balancing (ELB) and Auto Scaling can automatically scale a customer's AWS resources up to meet unexpected demand and then scale those resources down as demand decreases.

We propose the use of ELB to provide automated application load distribution for incoming traffic across multiple Amazon EC2 instances. By using ELB, we can help NASPO Purchasing Entities to achieve a high level of fault tolerance in applications, as the service seamlessly distributes workloads and load capacity in response to incoming application traffic. An additional fault tolerance capability that ELB provides is the ability to detect unhealthy EC2 instances for

automatic rerouting of traffic to healthy instances. This automated rerouting can exist for extended periods of time and provide the ability for the unhealthy instances to be restored. Depending on the availability needs of Purchasing Entities' applications, Team DLT can help to design EC2 instances and ELB architecture to be enabled within a single Availability Zone (AZ) or across multiple zones for even more consistent application performance.

We can work with NASPO Purchasing Entities to enable Auto Scaling services to allow for automated resource scaling of Amazon EC2 instances. As part of our services, we can help to capacity plan EC2 instances and define the conditions that trigger the Auto Scaling services, allowing instances to move up or down in capacity according to conditions that we help Purchasing Entities define. Auto Scaling enables customers to closely follow the demand curve for applications, reducing the need to provision Amazon EC2 capacity in advance.

AKAMAI

The Akamai CDN is designed to enable network operators to address network scalability challenges driven by the rapid increase in all types of Internet traffic. The software can elastically scale up and down network resources to support content spikes resulting from the broadcasts/flash-crowds of large events. The technology is also designed to accelerate service creation for service providers utilizing CDNs to launch innovative multiscreen video services that help drive subscriber acquisition and loyalty.

GOOGLE

Google Apps can elastically scale from one user to tens of thousands of users. The usage of the services included are designed to horizontally scale within the confines of the per user storage allocation. Google Apps can be purchased with a 30GB per user allocation or with unlimited storage.

6.1.2.5 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

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.

NIST requires that cloud systems allow for resource control and optimization through metering capabilities similar to a pay-per-use or charge-per-use basis. We understand the importance of measured services to adequately forecast spend and allow for Purchasing Entities to control use

of IT resources. Team DLT will utilize AWS automated monitoring systems to provide a high level of service performance and availability.

How Team DLT's Cloud Solutions Satisfy NIST Characteristic and Capability: We will leverage proactive monitoring 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.

AWS enables customers the ability to eliminate the need for costly hardware and the administrative struggles that accompany management of on-premise systems. Instead of the weeks and months it takes to plan, budget, procure, set up, deploy, operate, and hire for a new project, we can support NASPO Purchasing Entities' immediate cloud deployment of an AWS environment with unprecedented capabilities from 1, 10, 100, or 1,000 servers. Additionally, Team DLT's proposed implementation of AWS provides NASPO the ability to utilize a combination of purchasing models to immediately procure and implement cloud compute instances. These purchasing models can be manipulated and optimized to ensure increased cost efficiencies for NASPO Purchasing Entities and provide greater insight into measured services. The purchasing models include On-Demand, Reserved, Dedicated, and Spot Instances for true flexibility in terms of cost projections and management:

- On-Demand Instances allow customers to pay a fixed rate by the hour with no commitment
- Reserved Instances provide customers with a capacity reservation and offer a significant discount on the hourly charge for an instance
- Dedicated Instances allow customers to pay for compute instances launched on hardware specifically provisioned for the Purchasing Entity
- Spot Instances enable NASPO Purchasing Entities to bid whatever price they want to pay for instance capacity providing for even greater savings if their applications have flexible start and end times

AKAMAI

Akamai meets this characteristic by providing a basic monthly rate service plan and a metered bursting charge to provide flexibility to the Purchasing Entity based on scalable utilization. Akamai measurement tools will provide both the customer and the provider with an account of what has been used/consumed.

GOOGLE

Google Apps is optimized for unlimited use by all end users. Google ensures high availability, low latency and fault tolerance as a part of the contracted services. Metering is performed at the User level for this service.

6.1.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

All AWS offerings are applicable to each subcategory.

AKAMAI

Akamai service offerings can be divided into the following solution areas: Web Acceleration, Media Delivery, Web Security, and Data Center Security. All Akamai offerings are applicable to each subcategory.

GOOGLE

Google Apps are applicable to each subcategory. Google Apps can be subdivided into:

- SaaS
  - Business Continuity/Disaster Recovery
  - Collaboration
  - E-Discovery
  - Office Productivity
  - Message Filtering
  - Mobile Data Management
  - Messaging or Email as a Service

6.1.4 RFP Reference 8.1.4

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

DLT Response:

As stated in sections 6.1.1. and 6.1.2, Team DLT understands the importance of complying with NIST characteristics and having a diverse service and deployment model to meet the unique requirements of each NASPO Purchasing Entity. We certify that we will provide solutions that comply with the requirements of Attachments C & D.

6.1.5 RFP Reference 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**.

DLT Response:

Please see section 6.1.2.

6.2 SUBCONTRACTORS [RFP REFERENCE 8.2]

As the prime contractor, DLT intends to provide all cloud solutions directly to the NASPO Purchasing Entities as described in this proposal. Furthermore, DLT has teamed with Day1 Solutions as described in sections 8.2.2 and 8.2.3 to create a robust services offering to fully support all cloud solutions which may be needed by Purchasing Entities. As both DLT and Day1 maintain high tier qualifications within the Amazon Web Services Partner Network, the combination of the capabilities of our two companies in the delivery of services will ensure the best possible solution.

DLT is highly selective in choosing our subcontractors. Our criteria for selection and evaluation of subcontractors include: long-term stability, financial strength, history of customer satisfaction, technical excellence, complementary skills, culture, business practices, and demonstrated commitment to customer success. DLT will continuously evaluate potential subcontractors to be responsive to NASPO's existing and evolving needs throughout the term of the contract. DLT maintains a database of pre-screened and approved subcontractors consisting of proven small and large IT products, services and solution based companies. Subcontractors are evaluated based on their ability to provide services described in the NASPO ValuePoint Master Agreement for Cloud Solutions. Criteria for subcontractor selection are designed to optimize the relationship for the long-term and include:

- Cloud domain-specific capabilities to support the existing and future requirements of NASPO
- Demonstrated ability to perform in specific NASPO task areas to include IaaS, SaaS and PaaS
- Strong past performance record, with emphasis on diverse enterprise level customers
- Proven record of service excellence (both technical expertise and quality management) in support of complex, multidisciplinary and geographically dispersed task orders
- Demonstrated access to resources qualified and available to support multiple task order requirements
- Competitive employee benefits to attract and retain proven performers
- Certifications within their core competencies

6.2.1 RFP Reference 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.

DLT Response:

As an AWS Premier Partner, DLT meets the highest standards for partners within the AWS Partner Network (APN). To achieve this status a partner must have invested significantly in their AWS practice, have extensive experience in deploying customer solutions on AWS, have a strong bench of trained and certified technical consultants, have at least one APN Competency,

have expertise in project management, and have a healthy revenue-generating consulting business on AWS.[10] As such, it is DLT's belief that it can meet all of the requirements of this solicitation. That being said, the cloud environment is highly diverse and ever changing, so in order to offer as wide an array of support and experience to the NASPO community, DLT has teamed with Day1 Solutions (Team DLT). Day1 also holds many AWS competencies including Advanced Consulting, Managed Services, Direct Connect, Big Data and more. This teaming relationship ensures that NASPO will have a one-stop solution provider that can handle all needs brought forth by the diverse NASPO community. As the prime contract holder, DLT will manage all joint solutions brought forth to ensure that all deliveries meet the high standards set forth by DLT throughout this solicitation.

As part of Team DLT, Day1 is a current contract holder of WSCA and is currently authorized to provide services to NASPO and Purchasing Entities. Day1 will be leveraged as a subcontractor to provide insight into how to manage and operate NASPO ValuePoint Master Agreement for Cloud Solutions as stated in solicitation CH16012 and accompanying Questions and Answers, Amendments, etc. As a prime contract holder for the existing WSCA vehicle, Day1's extensive knowledge with the contract will prove to be vital to the successful transition and continued operation of the NASPO ValuePoint Master Agreement. Day1's discrete knowledge of existing WSCA customers will allow them to stay engaged as a subcontractor. Additionally, Day1's specific experiences with existing Participating Entities and Purchasing Entities allows for even greater understanding of work at the task order level.

Day1 will also be leveraged across the contract as a part of our services delivery team for value-add services. Day1's various partnership agreements with AWS make them qualified to provide end-to-end AWS specific solutions to NASPO.

6.2.2 RFP Reference 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.

DLT Response:

Team DLT will work on an ongoing basis to determine which companies' skill sets best support the project needs of the individual Purchasing Entity. By working closely in tandem, Team DLT will ensure the best delivery to the customer by utilizing the strengths of both companies. As the prime contract holder, DLT will manage all service delivery engagements and pull in the strengths of our partner as needed to deliver the most complete solution for the Purchasing Entity.

As there are no defined workloads or use cases specified in the solicitation it is difficult to determine, and thus enumerate, each position which will be utilized to support the Purchasing Entities varying needs. Our teams of in-house AWS certified engineers will work closely with Purchasing Entities to determine the scope and requirements of each project. As previously

[10] https://aws.amazon.com/partners/premier/

stated, the combination of skillsets and experience that comprise Team DLT will be sufficient to meet any need that may arise during the term of this contract.

6.2.3 RFP Reference 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.

DLT Response:

Team DLT, comprising DLT's own proven capabilities and Day1 Solutions (a current NASPO WSCA prime and AWS Advanced Consulting Partner), will combine to provide a comprehensive set of capabilities to meet the needs of NASPO Purchasing Entities. Day1's qualifications as described in 8.2 more than qualify them to support the requirements of this RFP and as the prime contract holder, DLT will manage all service delivery engagements and pull in the strengths of our partner as needed to ensure compliance of all requirements. Please refer to section 6.2.1 for additional information on Day1 Solutions' qualifications.

6.3 WORKING WITH PURCHASING ENTITIES [RFP REFERENCE 8.3]

6.3.1 RFP Reference 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.

DLT Response:

If DLT has actual knowledge of a confirmed breach that affects the security of any Purchasing Entity's Content or Customer Data subject to applicable data breach notification law, DLT's Contract Manager will ensure that the appropriate Team DLT staff (a) promptly notify the Purchasing Entity via electronic mail and telephone, (b) take commercially reasonable measures to address the breach in a timely manner to mitigate the effects and minimize any damage resulting from such breach, and (c) take prompt and appropriate corrective action aimed at preventing the reoccurrence of a similar breach in the future.

AKAMAI

Akamai will work with the Purchasing Entities to establish a "run-book" before the contract is signed. This document has information related to the personnel involved, contact information, follow-the-sun support information, etc. With regards to response times, Akamai stands by its SLAs for delivery and cloud security. Please see Appendix 3, section 2.

GOOGLE

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. 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 the Purchasing Entity or, at Google's discretion, by direct communication (e.g., by phone call or an in-person meeting).

6.3.2 RFP Reference 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.

DLT Response:

DLT agrees not to use any information obtained via award or administration of this contract to conduct any of the listed prohibited activities unless prior authorization is obtained from the Participating Entity. DLT agrees that it will not use any information gathered from this contract to create marketing campaigns. No advertising profiles will be created nor any advertising pushed to the services being proposed.

6.3.3 RFP Reference 8.3.3

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

DLT Response:

AMAZON WEB SERVICES

The AWS suite of services allows customers to create custom environments of their own configuration and size at any time. Should the customer wish to create a test/staging environment identical to a production environment, they simply need to create an Amazon Machine Image (AMI) of their production environment and use that image to create an identical test/staging environment. This process can also work in reverse to take a test/staging environment and create a full production environment.

AKAMAI

Akamai's solution does not provide application hosting.

GOOGLE

Google Apps, a SaaS offering, does not include a testing/staging environment.

6.3.4 RFP Reference 8.3.4

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

DLT Response:

AMAZON WEB SERVICES

The AWS Management Console is a single destination for managing all AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. Use the AWS Management Console to perform any number of tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console also enables customers to manage all aspects of their AWS account, including accessing monthly spending by service, managing security credentials, or even setting up new AWS Identity and Access Management (AWS IAM) users. The AWS Management Console supports all AWS regions and lets customers provision resources across multiple regions.

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. Additionally, AWS provides two graphical user interfaces, the AWS Management Console and the AWS ElasticWolf Client Console. The AWS ElasticWolf Client Console has incorporated Section 508 requirements and AWS has prepared a Voluntary Product Accessibility Template (VPAT) for the Console, which outlines the Console's accessibility features.

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 file transfers to and from Amazon Simple Storage Service (Amazon S3).

USE EXISTING MANAGEMENT TOOLS

Many of the tools that organizations use to manage on-premises environments can be integrated with AWS as well. Integrating an AWS environment can provide a simpler and quicker path for cloud adoption, because a customer's operations team does not need to learn new tools or develop completely new processes. For example:

- AWS Management Portal for vCenter enables customers to manage their AWS resources using VMware vCenter. The portal installs as a vCenter plug-in within the existing vCenter environment. Once installed, it enables customers to migrate VMware VMs to Amazon EC2 and manage AWS resources from within vCenter. The AWS resources that customers create using the portal can be located in their AWS account, even though those resources have been 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 customers can use to import their preexisting Virtual Machines (VMs) to Amazon EC2. Using the connector, importing a VM is as simple as selecting a VM 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) into which the VM should be imported. Once the VM has been imported, customers can launch it as an instance from the AWS Management Console and immediately take advantage of all the features of Amazon EC2.

- AWS Management Pack for Microsoft System Center[11] enables customers to view and monitor their AWS resources directly in the Operations Manager console. This way, customers can use a single, familiar console to monitor all of their resources, whether they are on-premises or in the AWS cloud. You get a consolidated view of all 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 in the Operations Manager console.

AWS SERVICES

Appendix 4 contains a list of AWS's cloud services. Links are provided to the service website where you can find all service documentation, pricing information, FAQs, release notes, developer information and relevant case studies.

AKAMAI

The services Akamai provides are accessed through a portal that supports 508-compliant EIT (e.g., third-party browsers, keyboards, monitors, etc.). Akamai represents that its services subject to Section 508 compliance support, either directly or through equivalent facilitation, Section 508-compliant software and technologies, and that any further Akamai EIT accompanying our services are Section 508-compliant to the extent disclosed in Akamai's VPAT (Voluntary Product Accessibility Template™) available to U.S. federal contracting officials and to customers upon request.

GOOGLE

The Google Accessibility Engineering team has designed many features into the Google Apps suite to assist people with disabilities.[12] This includes integration with popular screen readers, speech to text services, and text to speech services to name a few.

6.3.5 RFP Reference 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.

[11] http://aws.amazon.com/windows/products/system-center/
[12] https://www.google.com/accessibility/index.html

DLT Response:

Team DLT asserts that our proposed solutions are readily accessible through multiple methods to include CLI, mobile, and through the web. The table below provides a list of applicable web browsers that NASPO Purchasing Entities may use to access provisioned resources.

Browser compatibility:

Browser	Version	AWS	Akamai	Google
Google Chrome	Latest 3 Versions	All services	All services	All services (current browser version plus 2 previous versions)
Mozilla Firefox	Latest 3 Versions	All services	All services	All services (current browser version plus 2 previous versions)
Microsoft IE	11, 10, 9	All services	All services, Version 10 and higher	All services (current browser version plus 2 previous versions)
Microsoft Edge	12	All services	All services	All services (current browser version plus 2 previous versions)
Apple Safari	9, 8, 7, 6	All services	All services	All services (current browser version plus 2 previous versions)

6.3.6 RFP Reference 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.

DLT Response:

DLT agrees that it will meet with any Purchasing Entity (PE) prior to the execution of any Service Level Agreement in order to determine if any sensitive or personal information will be stored or used in the Service that is subject to any law, rule or regulation providing for specific compliance obligations. Once we have met with the PE and fully understand any specific requirements, we will craft requirements for any of our proposed technologies to create a solution that meets the requirements.

6.3.7 RFP Reference 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.

DLT Response:

DLT utilizes industry standard project management and scheduling ideologies when implementing solutions with our customers. That being said, in most situations cloud solution delivery tends to be fairly custom based on the individual requirements of the customer. DLT will work in a consultative manner, for any technologies that we support, to develop delivery schedules that are mutually acceptable to the customer and DLT.

6.4 CUSTOMER SERVICE [RFP REFERENCE 8.4]

6.4.1 RFP Reference 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).

DLT Response:

Team DLT uses a coordinated system of Quality Assurance (QA) and Quality Control (QC) activities that include work product reviews, health checks, program performance surveillance, and quality monitoring to drive continuous improvement. Our QC approach follows the Project Management Institute's (PMI) quality management processes throughout project execution incorporating QC on all of our projects. Our team uses an internally dedicated QC organization with processes that include a structured review and audit of all activities and deliverables to verify they comply with the applicable procedures, standards, and deadlines.

We work directly with the project delivery team during the early stages of a project to establish plans, standards, metrics, and procedures to ensure that measurable, meaningful, and usable methods are identified for performing reviews and audits throughout the program life cycle. Where necessary we develop a Quality Control Plan (QCP) that comprehensively defines how we manage the quality of work products for each service contract to meet the program and mission requirements. We accomplish this through the continued application of our Quality Management System that incorporates the following:

- QA approaches establish the high-level management framework that is detailed in technical operations.

- QC Checkpoints review products and services at defined points of production to verify conformance through a rigorous inspection process. For inspections, we utilize statistical sampling to determine the number of pieces to be inspected, which saves time and enhances overall accuracy in the process. We create checklists upstream to ensure technical accuracy.

- QA Checkpoints provide independent examination of process outputs and use subject matter expertise to confirm that final outputs meet the input requirements.

Team DLT will provide a multi-tiered approach to issues resolution and escalation that includes an assignment of a Program Manager (PM) to manage task orders and all project resources, deliverables, artifacts, etc., from requirements to delivery.

Our PM will incorporate feedback from NASPO and all Project Managers to ensure that all issues are accurately tracked and escalated and that we have a comprehensive view of all outstanding action items. Team DLT has the processes and procedures in-place to facilitate rapid and quality responses that include centralized team planning and a decentralized execution process. The benefit of this approach to the overall contract is that centralized task order coordination facilitates overall resource tracking leading to the best utilization of resources; better prioritization of tasks; ensures clear communication; enables us to implement a centralized repository for lessons learned; and provide a single focal point. Decentralized task order execution ensures dedication to the task, rapid decision-making capability, and maximum responsiveness for multiple tasks.

Our PM has the authority to make commitments for Team DLT on technical and administrative matters related to this master agreement and will work in tandem with our Contract Manager (see section 5.1). Additionally, the PM will supervise all staff members supporting both the master agreement and all task orders, overseeing Task Order Managers that provide day-to-day direction to Task Leads and staff to ensure NASPO customer's satisfaction.

Our PM has reach-back to additional Team members and corporate SMEs. This element of our organizational structure ensures that high quality support is consistently available and delivered. Additionally, our reach-back capability enables us to respond quickly to new resource requirements and/or surges in workload requirements. As a services organization with Subject Matter Experts, each function within this element is input into a matrix for the NASPO contract and used as required to satisfy the requirements of individual NASPO Purchasing Entities.

Team DLT has compiled an extensive collection of webpages and associated Uniform Resource Locator (URL) specific to AWS services and resources and will make these available to NASPO. Additionally, as an AWS Advanced Consulting partner, we have access to the Amazon Partner Network (APN) with a comprehensive library of white papers, operational guides, implementation guides, best practices, etc., (content exclusive to APN partners) that we will make available to NASPO Purchasing Entities as required.

AKAMAI

Quality Assurance: Akamai ensures regular customer touch points as part of several service management packages. Customers can discuss challenges, deploy new configurations or review existing configurations during these meetings. Akamai also proactively reaches out to its customers after any attack activity and furnishes remedies, next steps, etc.

Escalation Plan: Escalation plans are well communicated between Akamai and the Purchasing Entity; proper priority is an essential component so the appropriate teams can respond to the request within acceptable timelines.

Service Level Agreement can be found in Appendix 3, section 2.

GOOGLE

Google Apps, a SaaS offering, provides Quality Assurance for each release. Service related issues should be escalated to the DLT Service Center for remediation. Service Level Agreement can be found in Appendix 3, section 3.

6.4.2 RFP Reference 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.
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.
c. Customer Service Representative will respond to inquiries within one business day.
d. You must provide design services for the applicable categories.
e. You must provide Installation Services for the applicable categories.

DLT Response:

DLT holds customer service as one of our highest priorities and have dedicated resources to meet all the requirements of this section. Support options are available to all DLT offerings.

a. Each Purchasing Entity will have a dedicated representative to support the needs of the customer on an ongoing basis. As we work closely with our customers, frequent communication will ensure that contact information will always be current.

b. DLT is available via phone and email 24x7x365.

c. DLT will comply with the one (1) business day requirement.

d. DLT has a full complement of technical personnel to work closely with our account representatives to provide design services for applicable categories.

e. DLT Expert Services offers implementation services for all applicable categories.

6.5 SECURITY OF INFORMATION [RFP REFERENCE 8.5]

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

1. Purchasing Entities continue to own their data.

2. Purchasing Entities choose the geographic location(s) in which to store their data—it does not move unless the Purchasing Entity decides to move it.

3. Purchasing Entities can download or delete their data whenever they like.

4. Purchasing Entities 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.

AMAZON WEB SERVICES 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.

6.5.1 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

All data management functions within the AWS environment can be managed with the AWS Management Console.

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. Figure 11 depicts the current AWS regions and Edge Locations, along with new regions that are coming soon.

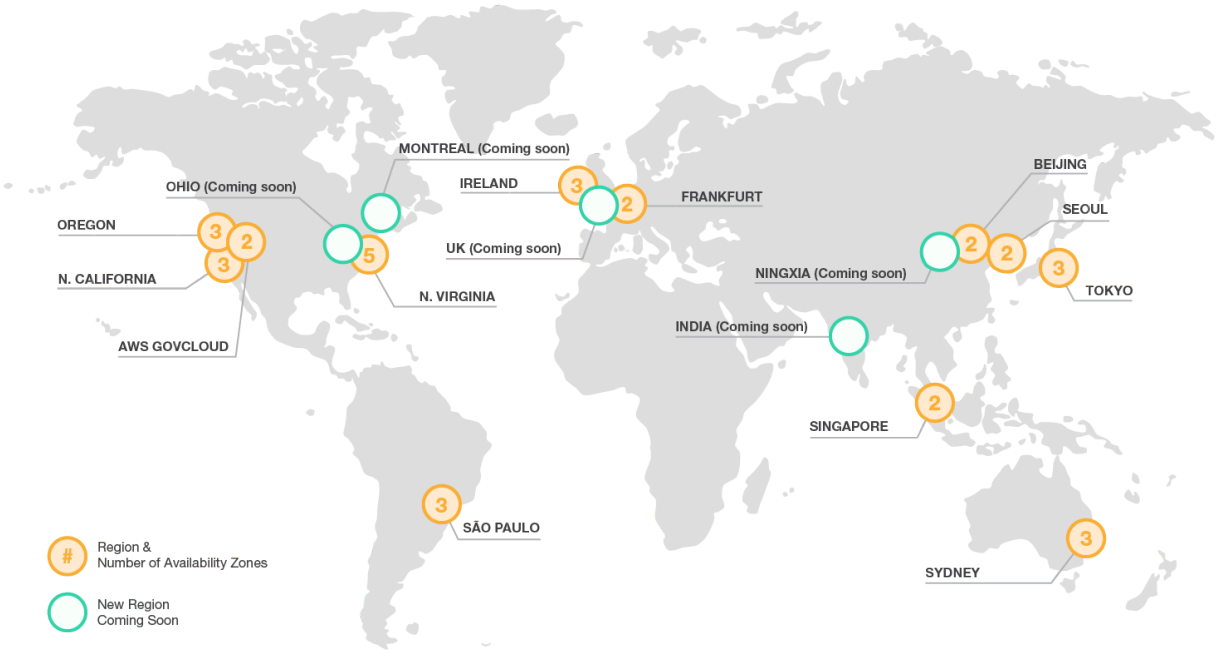

Figure 11. AWS Global Infrastructure[13]

The AWS products and services that are available in each region are listed on Amazon's website.[14]

SHARED RESPONSIBILITY MODEL

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 model (see the NIST Definition of Cloud Computing models). Customers should be clear as to their responsibilities in each 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),

[13] http://aws.amazon.com/about-aws/global-infrastructure/
[14] http://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/

other associated application software, configuration of the AWS-provided security group firewalls, and other security, change management, and logging features.

AWS's shared responsibility model is further explained on the AWS website.[15]

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.

DATA ENCRYPTION

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.

Securing Data 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 solution 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.[16]

Additionally, the *Securing Data at Rest with Encryption* whitepaper[17] provides an overview of the options for encrypting data at rest in AWS cloud services. It describes these options in terms of where encryption keys are stored and how access to those keys is controlled. Both server-side and client-side encryption methods are discussed with examples of how each can be accomplished in various AWS cloud services.

Securing Data in Transit

Protecting data in transit when running applications in the cloud involves protecting network traffic between clients and servers and network traffic between servers.

---

[15] http://aws.amazon.com/compliance/shared-responsibility-model/
[16] http://docs.aws.amazon.com/AmazonS3/latest/dev/s3-dg.pdf
[17] https://aws.amazon.com/whitepapers/encrypting-data-at-rest/

Services from AWS provide support for both Internet Protocol Security (IPSec) and Secure Socket 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.

The AWS *Overview of Security Processes* whitepaper (please see Appendix 5) provides greater detail on how to protect data in transit and at rest in the AWS cloud.

With AWS encryption, technologies such as SSL/TLS for data in transit and AES 256 for data at rest can be utilized to protect data.

DATA DELETION/DISPOSAL

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.

As an AWS authorized partner, Team DLT will leverage the AWS Federal Risk and Authorization Management Program (FedRAMP) certifications in the development of cloud-based solutions for NASPO customers. AWS is a FedRAMP Compliant Cloud Service Provider (CSP). AWS has completed the testing performed by a FedRAMP-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). AWS' compliance with FedRAMP requirements was achieved based on testing performed against the stringent NIST 800-53 Rev. 3 – Moderate baseline requirements, plus additional FedRAMP security controls. The HHS authorization is at the Moderate impact level to store, process, and protect a diverse array of sensitive government data.

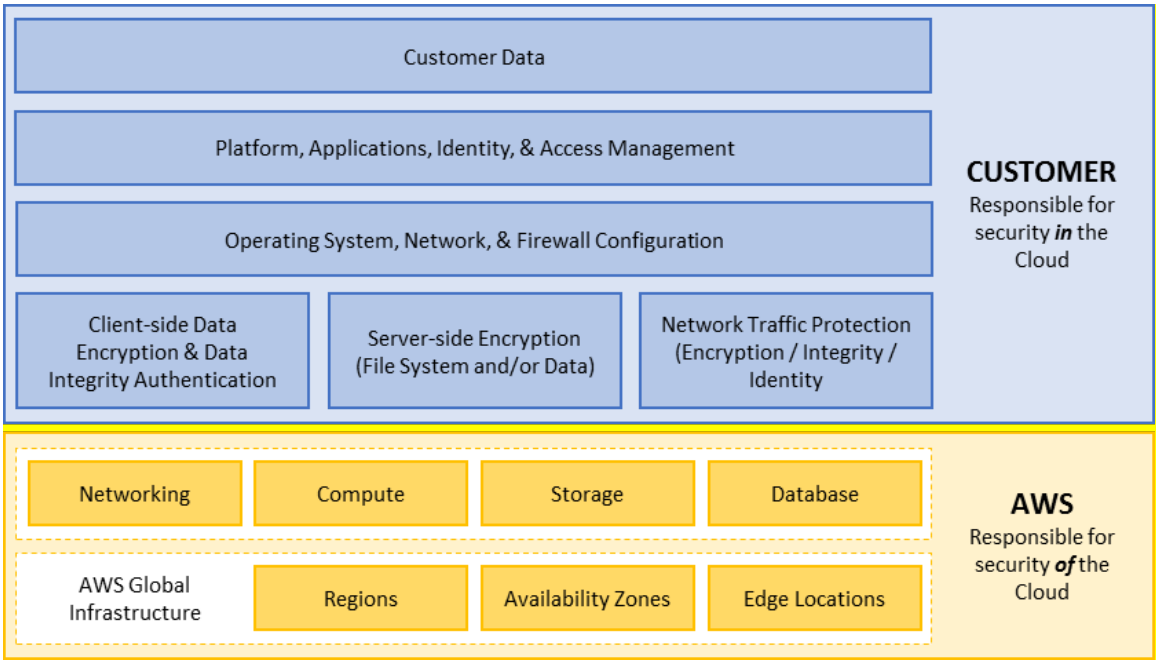

Figure 12. AWS Shared Responsibility Model

In support of the shared responsibility model (Figure 12), Team DLT will help to implement a wide variety of security services for NASPO's cloud based environment to provide greater control over services categorized as "Customer" responsibility in the shared model. These services include network security, access control, monitoring & logging, backup & replication, and data encryption. A brief synopsis of these services that support the security of information follows:

Network Security	Developing virtual and/or logical private cloud segments helps to increase network security and allows for NASPO to control cloud resource communications. We can assist in developing the network architecture and defining security inspection, detection, and monitoring points.
Access Control	Multiple layers of access control allow only Purchasing Entities' authorized users, customers, and applications to access cloud-based resources. Options such as setting up access control policies, individual user accounts, and unique credentials help to secure these controls.

Monitoring & Logging	Cloud-based resources require a much more integrated monitoring tool for keeping track of and monitoring cloud resources. The monitoring and logging capabilities we can help to implement can provide instant visibility into a Purchasing Entity's inventory as well as user and application activity.
Backup & Replication	NASPO's security strategy should include regular backups or snapshots of cloud-based objects and data. In many cases, backups can be set up to occur automatically and, in other cases, a Purchasing Entity can configure snapshots using a variety of backup options.
Data Encryption	Team DLT believes in using encryption wherever possible and enabling customers with multiple encryption methodologies. For instance, encrypting data at rest or in transit may be a requirement or security regulation that Purchasing Entities will look to follow.
Data Archival	With Amazon S3's lifecycle policies, customers can configure their data stores to be archived to Amazon Glacier or deleted after a specific period of time. We can support Purchasing Entities in using policy-driven automation to quickly and easily reduce storage costs, as well as, save time. With this, Purchasing Entities can specify a prefix, a time period, a transition to Amazon Glacier, and/or an expiration. 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.

AKAMAI

Akamai only caches publically available data from origin servers. For data handling, Akamai is fully compliant with FedRAMP Moderate JAB Level P-ATO and is delivering data for almost 90% of the entire US Government's web sites including all 15 Cabinet Level Agencies. Akamai is also PCI-Level-1 compliant and is handling data delivery for major transactional customers and retailers globally. Akamai will also sign up for the BAA of HIPAA and complies with the Designated School Official status in FERPA.

GOOGLE

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.

Google undergoes audits against key Security and Privacy Controls that cover Corporate Security Policies, Operation Security, Access Control, Organization Security, Asset Classification and Control, Personnel Security, Systems Development and Maintenance, Disaster Recovery and

Business Continuity and Regulatory Compliance. The results of these audits have led to SOC2, ISO2001 and FedRAMP Certifications. The audit reports are confidential records that can be reviewed by interested parties under NDA.

These security assertions are also backed by contractual obligations in the Google Terms of Service and Data Processing Amendment (please see Appendix 6) which outlines such things as Customer ownership of their data, confidentiality of all customer data, restricted scope of data use pertaining to the delivery of the contracted services only (meaning no advertising profiles) and the proper handling of data during the contract period and following the contract period.

6.5.2 RFP Reference 8.5.2

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

DLT Response:

AMAZON WEB SERVICES

DLT Solutions as a practice takes a consultative approach with all our customers. As part of this approach we work to educate the customer on the AWS Shared Responsibility Model (described in section 8.5.1). As Purchasing Entities may have disparate legal requirements relating to data privacy and security, this consultation is critical to ensure that both DLT and the Purchasing Entity work together to meet all requirements.

AKAMAI

Akamai is fully compliant with FedRAMP Moderate JAB Level P-ATO. Akamai is also PCI-Level-1 compliant and is handling data delivery for major transactional customers and retailers globally. Akamai will also sign up for the BAA of HIPAA and complies with the Designated School Official status in FERPA. Akamai will work with Purchasing Entities to determine individual requirements in order to meet applicable laws.

GOOGLE

Google commits via Terms of Service (please see Appendix 6) to abide by all applicable laws and to hold customers' data as confidential. The Terms make it the Customer's responsibility to respond to any third-party requests for data and, should Google receive such a request, Google will only respond to legal requests after notifying the customer if permitted by law to do so.

6.5.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Due to the nature of the Amazon Web Services model, neither DLT nor AWS will have any access to any customer data stored in AWS. As such, the customer will have complete control of

any data and it is not possible for DLT or AWS to use customer data for any purpose at any time. Should DLT inadvertently gain access to any customer data, we will fully comply with the provisions of this contract.

We can help to educate and inform NASPO Purchasing Entities on the AWS Shared Responsibilities model as described in section 6.5.1. Through this shared responsibilities model, Team DLT believes that NASPO Purchasing Entities will gain assurance that their user accounts and data remain inaccessible to DLT or AWS staff except in the course of daily operations, response to services and expressed terms of the Master Agreement.

AKAMAI

Akamai does not access customer 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.

GOOGLE

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.

6.6 PRIVACY AND SECURITY [RFP REFERENCE 8.6]

6.6.1 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

AWS's industry-leading security strength benefits you in many ways, one of which is by using a platform that is audited extensively by independent third-party assessors. At times, these audits confirm we can meet new requirements, even as they are issued, and this is the case for the National Institute of Standards and Technology (NIST) guidelines 800-171, which were released in June 2015. 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 we have already been audited under our 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 Appendix D of NIST Special Publication 800-171.[18]

AKAMAI

Akamai complies with NIST on all applicable specifications (e.g., 800-145, 800-53, 800-171, 800-122). Most of the applicable NIST specifications are also covered under the FedRAMP umbrella.

GOOGLE

Google currently maintains a FedRAMP Authorization to Operate. FedRAMP incorporates the relevant NIST SP and FIPS security requirements. Further, Google contractually commits to maintaining SOC2 and ISO27001 certifications.

6.6.2 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

The AWS cloud infrastructure has been designed and is managed in alignment with regulations, standards[19], and best practices, including:

- 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 [SSAE]

[18] http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-171.pdf
[19] For information on all of the security regulations and standards with which AWS complies, visit the AWS Compliance page (https://aws.amazon.com/compliance/).

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)

AKAMAI

PCI, ISO, FISMA, BITS, HIPAA, FedRAMP, NIST 800-53, NIST 800-171, FIPS 200

GOOGLE

Google has a 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 CJI 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.

6.6.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

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 model (see the NIST Definition of Cloud Computing models). Customers should be clear as to their responsibilities in each 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's shared responsibility model is further explained in section 6.5.1.

Please see SLAs in Appendix 3, section 1.

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. Figure 11 on page 66 depicts the current AWS regions and Edge Locations, along with new regions that are coming soon.

AWS NETWORK MONITORING AND PROTECTION

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. 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.

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.

By default, the AWS network provides significant protection against traditional network security issues, and we can implement further protection to ensure a greater security in depth. The following are a few examples:

- Distributed Denial of Service (DDoS) Attacks: AWS API endpoints are hosted on a 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.

- 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 in direct 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 Amazon website.[20] 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

- 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 interfaces can be placed into promiscuous mode, the hypervisor will not deliver any traffic 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 Address Resolution Protocol (ARP) cache poisoning do not work within Amazon EC2 and Amazon Virtual Private Cloud (VPC). While Amazon EC2 does provide ample protection against one customer inadvertently or maliciously attempting to view another's data, as a standard practice all sensitive traffic should be encrypted.

[20] http://aws.amazon.com/contact-us/report-abuse/

Team DLT has access to several security information and event management (SIEM) providers and has the in-house ability to implement these toolsets to develop a holistic view of the security posture of the Purchasing Entities' networks and IT infrastructure. We can develop a SIEM strategy that focuses on capturing and analyzing logs from multiple sources such as user activities, application usage patterns, system logs, and other security related events. Our experiences with SEIM in cloud environments is that they often have several gaps and weaknesses and typically do not replicate the capabilities of an enterprise architecture. An integration of native and third-party products are often used in filling those gaps.

We can leverage a SIEM tool for IaaS and PaaS implementations to collect, index and harness machine data across an infrastructure in real time. The additional benefits of using a SIEM tool include:

- E-Discovery: Search every data source required for E-Discovery from one place. Get instantaneous results across large data sets.
- FISMA: Securely collect, index and store all log and Machine Data along with audit trails to meet NIST requirements.
- HIPAA: Search all machine data to instantly assess reports of EPHI leakage and meet HIPAA's explicit log requirements.
- PCI: Rapid compliance with explicit PCI requirements for log retention/review and change monitoring, comprehensive reporting on all PCI controls such as passwords and firewall policy.
- SOX: Make the ambitious chore of compliance-mandated routine log review easy and straightforward.

AKAMAI

Threats outside the data center are mitigated either by our Kona Site Defender or Prolexic products. Kona Site Defender stops web applications attacks near the source of the malicious activity. Kona Site Defender intercepts threats at the edge of the internet, near the attacker, keeping malicious threats away from your origin. Prolexic mitigates DDoS attacks by diverting the user traffic to a closest scrubbing center where mitigation is applied. For threats within the data center, Akamai follows the guidelines specified by FedRAMP, PCI, and HIPAA.

GOOGLE

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.

Google has automated mechanisms that continuously scan the infrastructure 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 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 uses a proprietary storage and processing mechanism that isolates processing of data in chroot'd jails within a physical server. Access permissions restrict the ability for processes to interact between jails. Additionally, tenant data is striped in chunks across many different drives with each chunk having its own access control list. This helps ensure that data is logically isolated between customers in storage and during processing.

In addition to the inherent processing and storage mechanisms described above, Google's security controls including least privilege rights, logging and auditing have been implemented consistent with FedRAMP requirements.

6.6.4 RFP Reference 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.).

DLT Response:

AMAZON WEB SERVICES

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. Please see response to section 6.6.3 for additional details.

AKAMAI

Akamai adheres to the FedRAMP standard of data confidentiality as described below.

FEDRAMP PRIVACY REQUIREMENTS

Contractor shall be responsible for the following privacy and security safeguards:

1. To the extent required to carry out the FedRAMP assessment and authorization process and FedRAMP continuous monitoring, to safeguard against threats and hazards to the security, integrity, and confidentiality of any non-public Government data collected and stored by the Contractor, the Contractor shall afford the Government access to the Contractor's facilities, installations, technical capabilities, operations, documentation, records, and databases.

2. If new or unanticipated threats or hazards are discovered by either the Government or the Contractor, or if existing safeguards have ceased to function, the discoverer shall immediately bring the situation to the attention of the other party.

3. The contractor shall also comply with any additional FedRAMP privacy requirements.

4. The Government has the right to perform manual or automated audits, scans, reviews, or other inspections of the vendor's IT environment being used to provide or facilitate services for the Government. In accordance with the Federal Acquisitions Regulations (FAR) clause 52.239-1, contractor shall be responsible for the following privacy and security safeguards:

    a) The Contractor shall not publish or disclose in any manner, without the Contracting Officer's written consent, the details of any safeguards either designed or developed by the Contractor under this contract or otherwise provided by the Government. Exception - Disclosure to a Consumer Agency for purposes of C&A verification.

b) To the extent required to carry out a program of inspection to safeguard against threats and hazards to the security, integrity, and confidentiality of Government data, the Contractor shall afford the Government access to the Contractor's facilities, installations, technical capabilities, operations, documentation, records, and databases within XX hours. (FedRAMP recommends 72 hours) The program of inspection shall include, but is not limited to:

- Authenticated and unauthenticated operating system/network vulnerability scans
- Authenticated and unauthenticated web application vulnerability scans Authenticated and unauthenticated database application vulnerability scans Automated scans can be performed by Government personnel, or agents acting on behalf of the Government, using Government operated equipment, and Government specified tools.

c) If new or unanticipated threats or hazards are discovered by either the Government or the Contractor, or if existing safeguards have ceased to function, the discoverer shall immediately bring the situation to the attention of the other party.

GOOGLE

The Google Apps for Business Agreement (see Appendix 7) details the obligations of Google to hold all customer data as confidential and wholly owned by the customer. More details on the policies and procedures to enforce this are presented in section 6.6.3.

6.6.5 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Please refer to section 6.6.2 for information pertaining to the list of third-party attestations, reports, security credentials, and other controls obtained by AWS.

AKAMAI

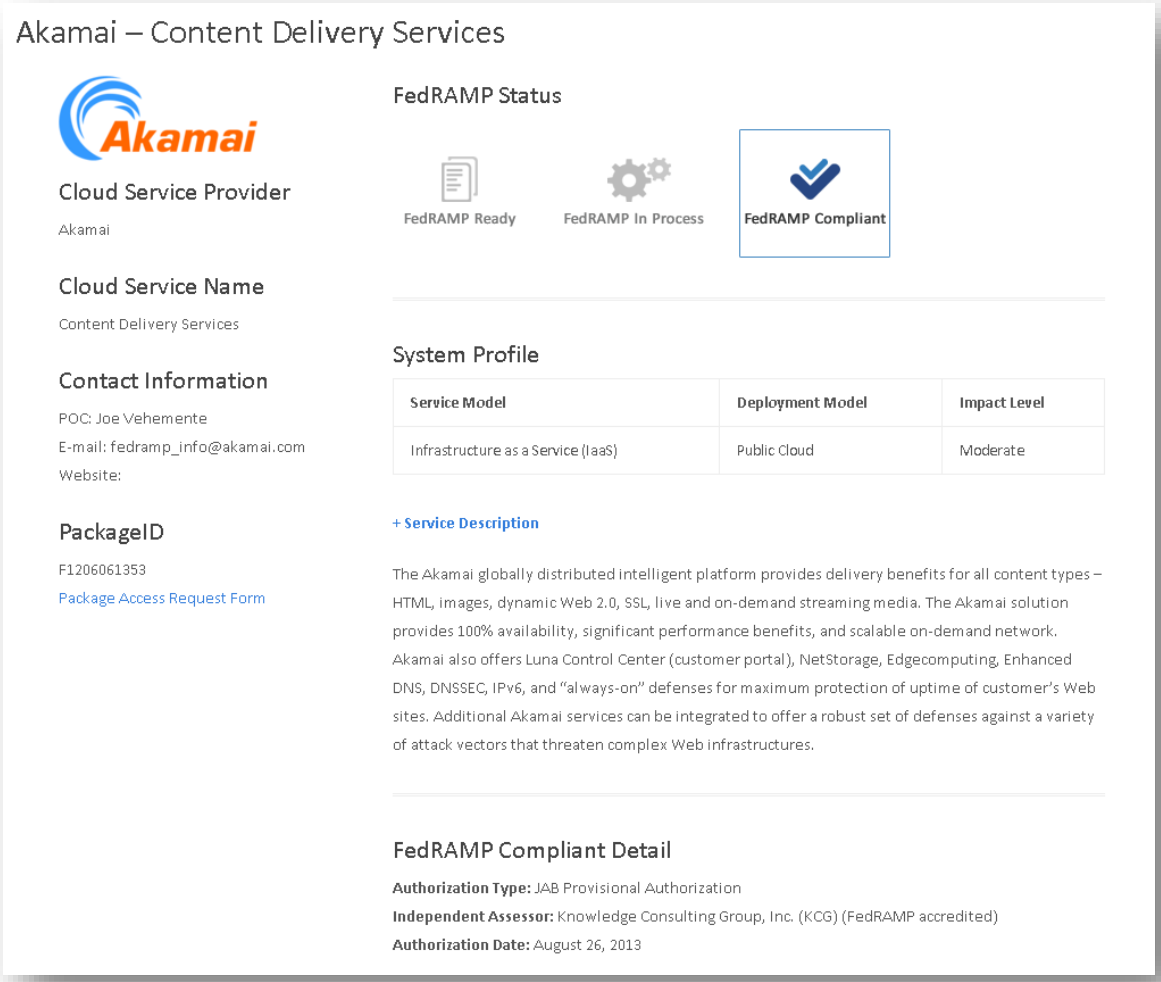

Figure 13. Akamai FedRAMP Status (March 9, 2016)[21]

GOOGLE

- SOC 1 (SSAE 16)
- SOC 2
- SOC 3
- ISO 27001
- ISO 27018
- FedRAMP

[21] Source: https://www.fedramp.gov/marketplace/compliant-systems/akamai-content-delivery-services/

6.6.6 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Team DLT can support NASPO Purchasing Entities' needs for a robust logging process using multiple AWS features and capabilities such as AWS CloudTrail, CloudWatch, and LogAnalyzer (along with third-party tool) to monitor instances, manage, and analyze log files. AWS CloudTrail is a web service that can be set up to record specific Application Program Interface (API) calls to supported AWS services in. These records can be delivered into a log file to an Amazon Simple Storage Service (Amazon S3) bucket. AWS CloudTrail can alleviate common challenges experienced in an on-premise environment, making it easier for Purchasing Entities to enhance security and operational processes while demonstrating compliance with policies or regulatory standards. With AWS CloudTrail, Purchasing Entities can get a history of AWS API calls for their account, including API calls made via 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.

Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications 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 any log files that applications generate. Purchasing Entities can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health, using these insights to react and keep their application running smoothly. We can guide Purchasing Entities to use CloudWatch Logs to monitor and troubleshoot systems and applications using existing system, application, and custom log files. Our team can help set up a strategy for shipping existing system, application, and custom log files to CloudWatch Logs and monitor these logs in near real-time. This helps Purchasing Entities better understand and operate their systems and applications, and they can store their logs using highly durable, low-cost storage for later access.

We can also help to implement LogAnalyzer for AmazonCloudFront for Purchasing Entities who want to analyze their Amazon CloudFront Logs using Amazon Elastic MapReduce (Amazon EMR). Using Amazon EMR and the LogAnalyzer application Purchasing Entities can generate usage reports containing total traffic volume, object popularity, a breakdown of traffic by client IPs, and edge location. LogAnalyzer application produces four sets of reports based on Amazon CloudFront access logs. Reports are formatted as tab delimited text files and delivered to the Amazon S3 bucket that customers specify.

- The Overall Volume Report displays total amount of traffic delivered by CloudFront over the course of a 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.

Amazon CloudFront's Access Logs provide detailed information about requests made for content delivered through Amazon CloudFront, AWS's 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.

AKAMAI

Akamai supports logging in CLF, W3C, or a hybrid of the two formats. Logs are collected at 15-30 minute intervals and can be transferred to Purchasing Entities' local SIEM. Appendix 8 contains the Akamai *Log Delivery Service User Guide*.

GOOGLE

Google has extensive logging from physical access to facilities, hands-on maintenance logs, network latency and perimeter traffic logs and more. The logs that are most relevant to Purchasing Entities are the logging activities of their own end users which includes User Name, Email Address, login data, login type, IP address of login, message sent, message received, type of log in, mobile access, mobile os, mobile model, number of files owned, files shared, viewed, downloaded, edited, deleted, calendar events, changes, RSVPs, number of hangouts, chrome device usage, security reports around 2-step verification, tokens in use, POP use, IMAP use, suspicious login, logins by type, groups logs, email log search, Admin Audit logs, predefined alerts and custom alerts related to Admin Audit, Token, Calendar, Login and Drive Audit Streams.

Google can contractually agree to audits to maintain certification at least once every 18 months.

6.6.7 RFP Reference 8.6.7

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

DLT Response:

AMAZON WEB SERVICES

AWS can restrict visibility of NASPO Purchasing Entities' cloud-hosted data to specific users and groups through role-based access control using the AWS Identity and Access Management (IAM) tool. AWS IAM allows Purchasing Entities to control the level of access they have to AWS resources. AWS IAM is a web service that enables AWS customers to manage users and user permissions in AWS. With AWS IAM, each user can have unique security credentials, eliminating the need for shared passwords or keys and allowing the security best practices of role separation and least privilege. 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 this tool, Purchasing Entities can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access. We can guide Purchasing Entities in setting proper permissions to specify who has access to AWS resources and which actions they can perform on those resources.

AKAMAI

Akamai only caches publically available data from origin servers. Restricting visibility will occur at the origin server and not at the caching locations.

GOOGLE

Google's obligations, as described in section 6.6.4, are to ensure there is no unauthorized access of Customer data. This is largely the responsibility of the customer. Purchasing Entities are responsible for ensuring their end users use the service according to acceptable use policies. During the onboard process, Purchasing Entities will be assisted in the configuration controls that are included to enforce acceptable use and security policies and trained in how to maintain oversight on the ongoing usage of the services in scope.

6.6.8 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

If DLT has actual knowledge of a confirmed breach that affects the security of any Purchasing Entity Content or Customer Data subject to applicable data breach notification law, DLT's Contract Manager will ensure that the appropriate Team DLT staff (a) promptly notify the Purchasing Entity via electronic mail and telephone, (b) take commercially reasonable measures to address the breach in a timely manner to mitigate the effects and minimize any damage resulting from such breach, and (c) take prompt and appropriate corrective action aimed at preventing the reoccurrence of a similar breach in the future.

The "AWS Security Center" can be leveraged by Purchasing Entities for security and compliance details about AWS. Purchasing Entities can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

AKAMAI

As all data cached by the Akamai service is intended to be public facing and thus readily available to all, the concept of a data breach in this case is not relevant.

GOOGLE

If Google becomes aware of a Data Incident, Google will promptly notify the Purchasing Entity 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 the Purchasing Entity or, at Google's discretion, by direct communication (e.g., by phone call or an in-person meeting). A Purchasing Entity 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.

6.6.9 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

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 Virtual Private Cloud. For example, you can create a public-facing subnet for your webservers 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. The design, engineering, and implementation of Amazon VPC provides the NASPO Purchasing Entities with isolated hosted server deployments. Amazon VPC includes the following characteristics:

- Provision logically isolated computing/processing resources to host NASPO customer applications/systems
- Isolated network that can be user defined to host the required NASPO customer applications/systems
- Controlled access to VPC that can be established through VPN connections

- Additional tools and services that can be configured to provide encryption in-transit and at-rest capabilities

- Dedicated instances to provide greater isolation and private cloud-like performance with computing instances on single-tenant compute hardware. This will ensure that sensitive data and compute has dedicated processing that only serve designated Purchasing Entities' resources.

AKAMAI

As all data cached by the Akamai service is intended to be public facing and thus readily available to all, security controls for physical and virtual Zones of Control Architectures (ZOCA) are not relevant.

GOOGLE

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 set foot in one of our data centers.

Google operates a global, multi-tenant 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 lifecycle at Google from arrival to final destruction. These disks are component parts 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 chroot'd jails within a physical server. Access permissions restrict the ability for processes to interact between jails. Additionally, tenant data is striped in chunks across many different drives with each chunk having its own access control list. This helps ensure that data is logically isolated between customers in storage and during processing.

In addition to the inherent processing and storage mechanisms described above, Google's security controls, including least privilege rights, logging and auditing, have been implemented consistent with FedRAMP requirements.

6.6.10 RFP Reference 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)

DLT Response:

AMAZON WEB SERVICES

Amazon's shared security model clearly delineates the security responsibility that falls under the customer's purview, but does not provide detailed guidance on building secure systems in accordance with FedRAMP guidelines. The AWS Security Architecture (Figure 14) is designed to provide an additional layer of guidance that can help organizations "right-size" the security approach so they can migrate faster while reducing compliance related security gaps in their system. The AWS Security Architecture is not meant to be prescriptive or comprehensive, but instead discusses best practices and illustrates how customers can configure and implement AWS services to make security and compliance easier. Regardless of the initial size or scope of the workload, the AWS Security Architecture will provide foundational guidance and direction so that Purchasing Entities can design their AWS infrastructures to be scalable, secure, manageable, and compliant.

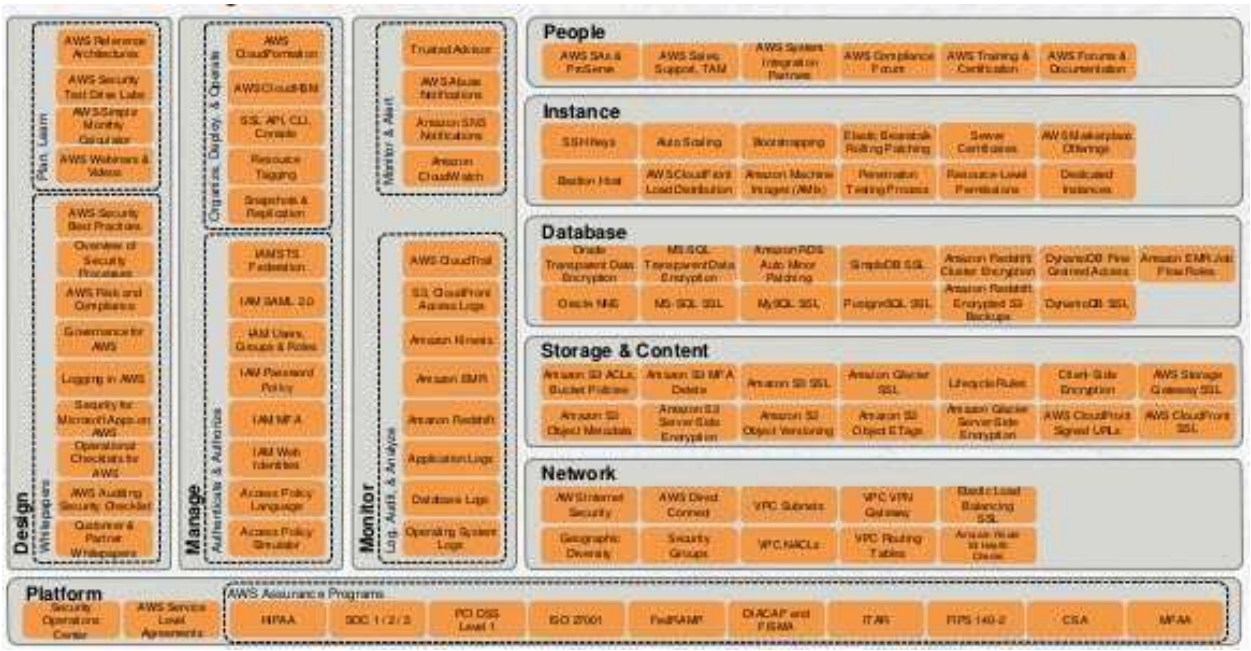

Figure 14. AWS Security Architecture

AKAMAI

Akamai's secure technical architecture fully complies with FedRAMP.

GOOGLE

Google's secure technical architecture fully complies with FedRAMP. The Google Apps Security Whitepaper (please refer to Appendix 9) explains the Security controls that support this service.

6.6.11 RFP Reference 8.6.11

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

DLT Response:

AMAZON WEB SERVICES

Due to the nature of the Amazon Web Services model, neither DLT nor AWS will have any access to any Purchasing Entity's data stored in AWS. As such, the Purchasing Entity will have complete control of any data and it is not possible for DLT or AWS to use customer data for any purpose at any time. Should DLT inadvertently gain access to any customer data, we will fully comply with the provisions of this contract.

AKAMAI

Akamai does not host any sensitive data on its platform.

GOOGLE

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

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's 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 filling 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 Google's 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. Additionally, the privacy team conducts research providing thought leadership on privacy best practices for Google emerging technologies.

6.6.12 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

For details for data at rest and in transit, please see section 6.9. Additionally, please see the below list of security and compliance certifications AWS possesses.

Published Certifications:

- AWS ISM Letter of Compliance
- AWS ASD Letter of Certification
- AWS ISO 9001:2008 Certification
- AWS ISO 27001:2013 Certification
- AWS ISO 27017:2015 Certification
- AWS ISO 27018:2014 Certification
- Multi-Tier Cloud Security Standard Level-3 (CSP) Certification
- AWS SOC 3 Report

Additional AWS Assurance Programs are shown in Figure 15.

AKAMAI

Akamai Compliance: ISO 27001, HIPAA, PCI DSS, and FedRAMP

GOOGLE

Google Apps Compliance: ISO 27001, SSAE 16 / ISAE 3402 / SAS 70 Audit, and FedRAMP

6.6.13 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Due to the nature of the Amazon Web Services model, neither DLT nor AWS will have any access to any customer data stored in AWS. As such, Purchasing Entities will have complete control of any data and it is not possible for DLT or AWS to use customer data for any purpose at any time. Should DLT inadvertently gain access to any customer data, we will take the measures described in section 6.3.1.

AKAMAI

As all data cached by the Akamai service is intended to be public facing and thus readily available to all, the concept of a data breach in this case is not relevant.

Figure 15. Additional AWS Assurance Programs

GOOGLE

Google will take and implement appropriate technical and organizational measures to protect Purchasing Entities' data against accidental or unlawful destruction or accidental loss or alteration or unauthorized disclosure or access or other unauthorized processing. If Google becomes aware of a Data Incident, Google will promptly notify the Purchasing Entity as permitted by law 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 the Purchasing Entity or, at Google's discretion, by direct communication (e.g., by phone call or an in-person meeting).

6.7 MIGRATION AND REDEPLOYMENT PLAN [RFP REFERENCE 8.7]

6.7.1 RFP Reference 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 de-provisioning 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.

DLT Response:

AMAZON WEB SERVICES

DLT is supplying the Purchasing Entity access to the full complement of AWS services. All AWS services are for the use of the customer at their discretion and under the full control of the Purchasing Entity. Is it completely at the Purchasing Entity's discretion as to when to initiate end-of-life activities to close down a service (application running on AWS). The Shared Responsibility Model (described in section 6.5.1) remains in place during end of life activities. Purchasing Entities can opt to maintain existing functionality within the AWS platform during migration for redundancy purposes and have complete control of data at all times.

AKAMAI

Akamai provides more than a one-year notice before it marks a product as EOL. Customers of that particular product are sent regular and timely information related to EOL and transition plans. Account teams also work hand in hand with the customers to make this transition as smooth as possible.

GOOGLE

We are constantly changing and improving our Services. We may add or remove functionalities or features, and we may suspend or stop a Service altogether.

You can stop using our Services at any time. Google may also stop providing Services to you, or add or create new limits to our Services at any time.

We believe that you own your data and preserving your access to such data is important. If we discontinue a Service, where reasonably possible, we will give you reasonable advance notice and a chance to get information out of that Service.

6.7.2 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Under the AWS business model, the customer has complete control of their data at all times. As long as the Purchasing Entity maintains an active AWS customer, they may at any time initiate

data retrieval from AWS. Per the AWS Shared Responsibility Model, there is no specific SLA for return of customer data as this completely within the Purchasing Entity's control.

AKAMAI

As the source data resides at the customer's origin, it is not anticipated that any data would need to be retrieved.

GOOGLE

A Purchasing Entity may choose the data migration and deletion options that work best for them. Google does not have a customer data access/export process or procedure. Each Purchasing Entities' exit strategy will vary to some degree based on their use of the services and the platform to which they are migrating.

As a general rule, the termination clause in the Google master terms provides customers with the ability to provide advance notice of their intent to not renew. It would be the expectation that the Purchasing Entity will have already engaged with a new platform provider in advance of the end date and work was already in progress to migrate data to the new platform. Google provides customers with a variety of migration tools under the Google TakeOut program as well there are third party independent software vendors that offer migration services and tools using Google APIs.

Google's obligations to protect the customer data and keep the services accessible and secure will persist until the Purchasing Entity's chosen end date. Data migration activities do not impact the ongoing use of the services by the end users so no downtime needs to be factored in for this activity. The Purchasing Entity would be able to coordinate with Google on the project plan for the exit strategy. Google agrees that once the Purchasing Entity has moved all their data, deleted all the records and end user accounts, all data will be expunged from all servers and backup takes within 180 days. Keep in mind that this means the individual chunks of encrypted data. It should not be assumed that at some later date within that 180 days that any data recovery can be accomplished. Once data is deleted by the customer and there are no legal holds on the data, pointers to that data are deleted making recovery virtually impossible.

6.8 SERVICE OR DATA RECOVERY [RFP REFERENCE 8.8]

6.8.1 RFP Reference 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).

DLT Response:

AMAZON WEB SERVICES

AVAILABILITY AND FAULT-TOLERANT DESIGN

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.

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 as described in section 6.5.1. 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 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.

AWS GovCloud[22] (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.

[22] http://aws.amazon.com/govcloud-us/

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.

Under the Shared Responsibility Model, Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are the responsibility of the Purchasing Entity. DLT will work with the Purchasing Entity to properly architect their solution to meet their requirements.

AKAMAI

Akamai only caches publically available data from origin servers. These caches do not need to be backed up. In the event of a data loss for caches, the origin servers will be sourced for repopulating cached content.

GOOGLE

a. Due the redundant nature of the Google infrastructure there is no expectation of any extended downtime. If a data center in use by any of your end users suffers a catastrophic failure, Google's dynamic health monitoring would just reroute the session to a different data center.

b. Google performs real time data replication to avoid unrecoverable data loss. Customers that configure the retention policies in Google Apps Vault can ensure that end users cannot inadvertently or maliciously delete email in advance of any defined retention schedules.

c. Google's corporate network is separate from its production infrastructure. Google is a global organization that has support personnel located at key office installations and tests our business continuity programs annually. These tests are to validate that if corporate headquarters is off the grid, critical business functions can be picked up by other staff.

d. Google's RTO is zero.

e. RPO/RTO objectives are zero.

6.8.2 RFP Reference 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 within the continental United States.

DLT Response:

AMAZON WEB SERVICES

As NASPO Purchasing Entities begin to adopt the AWS platform, Team DLT will guide them through different methods for backup and restore methodologies such as methods for data

backups and server image backups. 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 via code rather than being hardware based.
- Capacity is available at incremental cost rather than up-front 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 RATHER THAN SERVERS

The Amazon Elastic Compute Cloud (Amazon EC2) service enables the backup and recovery of a standard server, such as a web server or application server, so that customers can focus on protecting their configuration and the state of data rather than 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, temporary files, and so on. 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 on an Amazon Machine Image (AMI) and can also 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 it the identity of the web server and any security credentials required for initial setup. The instance is based on a pre-built 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 file(s).

- 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 also available for a number of other Linux distributions.

Figure 16 depicts a traditional backup approach and Figure 17 depicts an Amazon EC2 backup approach.

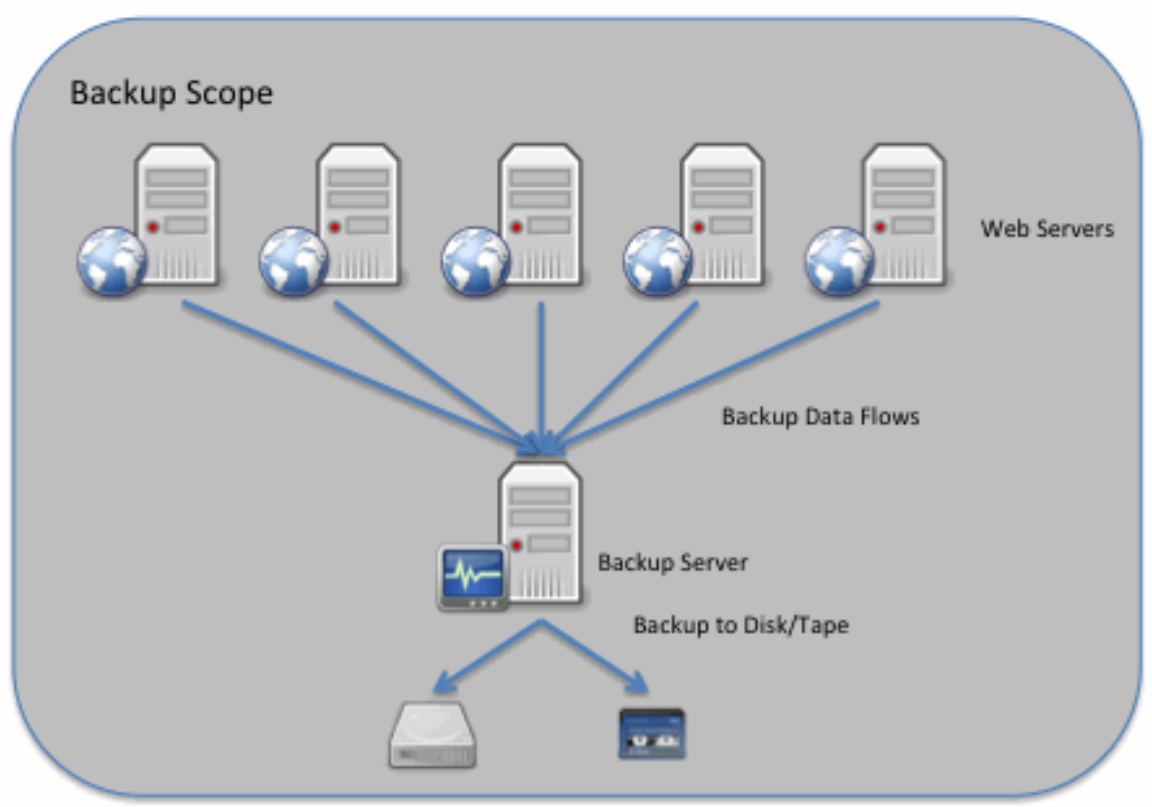

Figure 16. Traditional Backup Approach

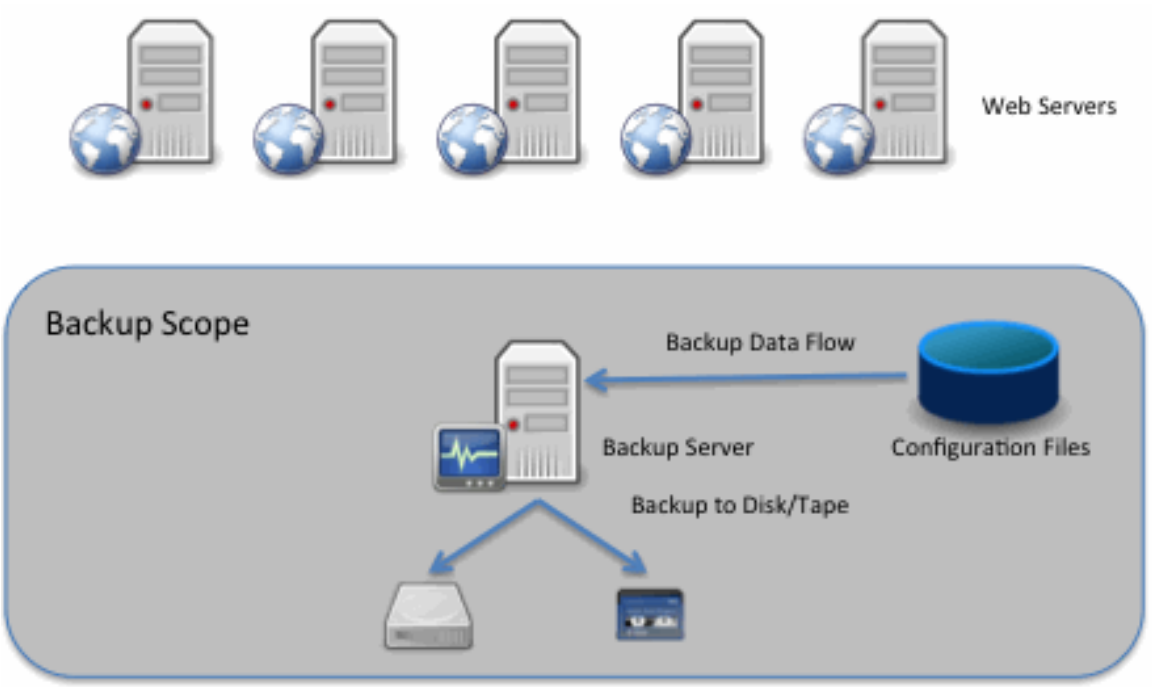

Figure 17. Amazon EC2 Backup Approach

In this case, there is no need to back up the server itself. The relevant configuration is contained in the combination of the AMI and the configuration file(s). So, the only components requiring backup and recovery are the AMI and configuration file(s).

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, Purchasing Entities can create totally independent copies of the AMI by:

- Sharing the original AMI to another specified AWS account controlled by the Purchasing Entity
- Starting a new instance based on 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. Of course, Purchasing Entities can also create multiple copies of the AMI within the same account.

CONFIGURATION FILES

Customers use a variety of version management approaches for configuration files, and they can follow the same regime for the files used to configure their Amazon EC2 instances. For example, a Purchasing Entity could store different versions of configuration files in designated locations and securely control them like any other code. That Purchasing Entity could then back up these code repositories using the appropriate backup cycle (e.g., daily, weekly, monthly) and snapshots to protected locations. Further, Purchasing Entities can use Amazon S3 to store their configuration files, taking advantage of the durability of the service in addition to backing up the files to an alternate location on a regular basis.

DATABASE AND FILE SERVERS

Backing up data for database and file servers differs from the web and application layers. In general, database and file servers contain larger amounts of business data (tens of GB to multiple TB) that must be retained and protected at all times. In these cases, Purchasing Entities 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 less than 1 TB), an alternative backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. While 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 (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. With data centers in 12 regions around the world (4 in the United States), AWS provides a set of cloud-based DR services that enable rapid recovery of IT infrastructure and data.

GENERAL DISASTER RECOVERY/COOP AND BACKUP REQUIREMENTS AND ISSUES

Some of the minimum needs and requirements in a traditional DR approach are:

- Facilities to house additional infrastructure, including power and cooling
- Security to ensure 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 all mission-critical services, including storage appliances for the supporting data, and servers to run applications and back-end services such as user authentication, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), monitoring, and alerting.

AWS CAPABILITIES FOR DR/COOP/BACKUP SOLUTIONS

With AWS, customers can eliminate the need for additional physical infrastructure, off-site 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 Purchasing Entities 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 DR capabilities:

- Fast Performance: Fast, disk-based storage and retrieval of files.
- No Tape: Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software.
- Compliance: Minimize downtime to avoid breaching Service Level Agreements (SLAs).
- Elasticity: Add any amount of data quickly. Easily expire and delete without handling media.

- Security: Secure and durable cloud DR 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 multiple DR strategies. Figure 18 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 DR event.

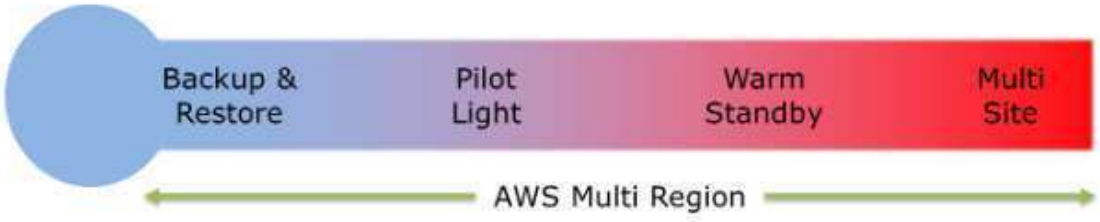

Figure 18. Spectrum of DR Options

Each DR option is discussed in more detail below:

Backup and Restore	This DR method is usually the easiest to adopt as it mimics a traditional IT environment. In most traditional environments, data is backed up to tape and sent off-site regularly. To support this type of DR requirement with AWS, Day1 may leverage backups to an Amazon Simple Storage Service (S3) bucket. Using Amazon S3 is ideal for backup data, as it is designed to provide 99.999999999% durability of objects over a given year. Transferring data to and from Amazon S3 is typically done via the network, and it is therefore accessible from any location. Additionally, we can support the use of AWS Storage Gateway, where NASPO customers can automatically back up on-premises data to Amazon S3.
Pilot Light for Simple Recovery into AWS Warm Standby Solution	In this scenario we will guide the NASPO customers in understanding the most critical components of their system. These core components will be pre-configured and waiting in sedentary state in the cloud. This scenario is similar to the Backup and Restore scenario This scenario is analogous to a backup and restore scenario; however, NASPO customers must ensure 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, NASPO would execute the DR plan which rapidly provisions a full-scale production environment around these components.
Warm Standby Solution in AWS	In a Warm Standby DR scenario, a scaled-down version of the fully functional environment is always running in the cloud. It decreases recovery time as some services are always running. By identifying business-critical

	systems, customers could fully duplicate these systems on AWS and have them always on.
Multi-site Solution Deployed on AWS and On-Site	A multi-site solution will run in the cloud as well as on existing on-premise infrastructure in an active-active configuration. During a disaster situation, traffic would simply be re-routed to the cloud based instances, which can scale to handle the full production load.

DR RESOURCES

There are multiple resources to help organizations start using AWS for a DR/COOP and backup solution:

- Read the AWS whitepaper *Using AWS for Disaster Recovery* [23]
- Read the Forrester whitepaper *File Storage Costs Less in the Cloud than In-House* [24]
- Review a sample AWS DR architecture[25]

AKAMAI

Akamai only caches publically available data from origin servers. These caches do not need to be backed up. In the event of a data loss for caches, the origin servers will be sourced for repopulating cached content.

GOOGLE

a. Online tape backups are performed on each server daily

b. Standard tape backup routines

c. Stored on tape at data center location in locked vault. Some tapes may be stored offsite with an archiving vendor.

d. As described elsewhere in this response, Google replicates customer data in real time to all data centers located in the continental United States, there is no single server, nor a single data center that stores your data. In this manner, service availability is assured even in the face of a server failure or issues with an entire data center.

6.9 DATA PROTECTION [RFP REFERENCE 8.9]

6.9.1 RFP Reference 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.

[23] http://media.amazonwebservices.com/AWS_Disaster_Recovery.pdf
[24] http://media.amazonwebservices.com/Forrester_File_Storage_Costs_Less_In_The_Cloud.pdf
[25] http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_disasterrecovery_07.pdf

DLT Response:

AMAZON WEB SERVICES

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. Purchasing Entities 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.

SECURING DATA 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*.[26]

Additionally, the *Securing Data at Rest with Encryption* whitepaper[27] provides an overview of the options for encrypting data at rest in AWS cloud services. It describes these options in terms of where encryption keys are stored and how access to those keys is controlled. Both server-side and client-side encryption methods are discussed with examples of how each can be accomplished in various AWS cloud services.

SECURING DATA IN TRANSIT

Protecting data in transit when running applications in the cloud involves protecting network traffic between clients and servers and network traffic between servers.

Services from AWS provide support for both Internet Protocol Security (IPSec) and Secure Socket 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.

The *AWS Security Best Practices* whitepaper[28] provides greater detail on how to protect data in transit and at rest in the AWS cloud.

With AWS, encryption technologies such as SSL/TLS for data in transit and AES 256 for data at rest can be utilized to protect data.

AKAMAI

Akamai does not host any sensitive data on its platform.

---

[26] http://docs.aws.amazon.com/AmazonS3/latest/dev/Welcome.html
[27] http://media.amazonwebservices.com/AWS_Securing_Data_at_Rest_with_Encryption.pdf
[28] http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

GOOGLE

All connections from customer end point devices to Google's Front End Servers are encrypted with enforced HTTPS sessions using Forward Secrecy. Google websites and properties use robust public key technologies: 2048-bit RSA or P-256 ECDSA SSL certificates issued by a trusted authority (currently the Google Internet Authority G2).

All data in transit between Google's Data Centers traverses across Google's private fiber network using a customized, proprietary encryption technology.

Google hard drives leverage technologies like FDE (full disk encryption) and drive locking to protect data at rest.

These methods of encryption are fully managed by Google and Google's Key Management Servers based on 128-bit or stronger Advanced Encryption Standard (AES).

ENCRYPTION KEYS AND CIPHERS SUPPORTED BY GOOGLE

- Protocols
  - TLS 1.2
  - TLS 1.1
  - TLS 1.0
  - SSL 3.04
  - QUIC
- Cipher suites
  - ECDHE_RSA with AES
  - ECDHE_RSA with 3DES
  - ECDHE_ECDSA

6.9.2 RFP Reference 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.

DLT Response:

On a case-by-case basis, DLT will be willing to sign agreements such as a Business Associate Agreement (BAA) with a Purchasing Entity.

6.9.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Due to the nature of the Amazon Web Services model, neither DLT nor AWS will have any access to any Purchasing Entity's data stored in AWS. As such, the Purchasing Entity will have complete control of any data and it is not possible for DLT or AWS to use that data for any purpose at any time. Should DLT inadvertently gain access to any customer data, we will take the measures described in section 6.3.1.

Team DLT's engagement with NASPO Purchasing Entities will outline access levels and describe account privileges. We use multiple levels of control including using elevated privileges only when required, using role based authentication, and AWS IAM to control access to data for the intended purpose of each engagement.

AKAMAI

Akamai will not use the customer data or customer-related data it caches for any other purpose including but not limited to data mining, and shall not resell nor otherwise redistribute information gained from its access to the data received as a result of this RFP.

GOOGLE

Google's Master Terms of Service and Data Processing Amendment (please see Appendix 6) outline such things as Customer ownership of their data, confidentiality of all customer data, restricted scope of data use pertaining to the delivery of the contracted services only (meaning no advertising profiles) and the proper handling of data during the contract period and following the contract period.

6.10 SERVICE LEVEL AGREEMENTS [RFP REFERENCE 8.10]

6.10.1 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

AWS SLAs are standard and non-negotiable. However, under the DLT Managed Services offerings there is the flexibility for enhanced SLAs on a case-by-case basis.

AKAMAI

Our service is designed to support a wide variety of state, local, federal, and commercial customers. Given such breadth, consistency across operations is of paramount importance. Custom SLAs imply exceptions to the standard processes, thus impacting consistency in service delivery. Consequently, it is not feasible to customize SLAs.

GOOGLE

Google's SLA is not negotiable. In order to deliver SaaS services at scale any vendor would need to have a standard SLA to avoid any conflicts or disputes in the operational performance of the service.

6.10.2 RFP Reference 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 within which the infrastructure must operate to meet IT System and Purchasing Entity's requirements.

DLT Response:

Service Level Agreements (SLAs) for all proposed products and services may be found in Appendix 3.

6.11 DATA DISPOSAL [RFP REFERENCE 8.11]

Specify your data disposal procedures and policies and destruction confirmation process.

DLT Response:

AMAZON WEB SERVICES

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

1. Purchasing Entities continue to own their data.
2. Purchasing Entities choose the geographic location(s) in which to store their data—it does not move unless the Purchasing Entity decides to move it.
3. Purchasing Entities can download or delete their data whenever they like.
4. Purchasing Entities 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

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.

AKAMAI

Akamai does not host any sensitive data on its platform, therefore data destruction is not required.

GOOGLE

There are two dimensions to data disposal / destruction.

First, Google manages the lifecycle of all servers in all data centers. If a server needs to be retired, it is first erased, degaussed and then crushed or shredded on site. Any data that was on that server that is an active component of any customer's data set would have been migrated to a different healthy server. No one server is dedicated to a single customer.

Second, customers are responsible for their own data retention policies and enforcement using the tools provided with Google Apps Vault. Customers are able to set retention policies on certain data types with options to have that data expunged at a set time, and customers have an expectation that at the conclusion of a service agreement Google would ensure their data is deleted. With Google Apps Vault, customer data disposal requirements can be met by the customer setting a retention policy with an expunge rule for certain data types within a certain time span. This is configured at the discretion of the customer. At the conclusion of the service agreement, the Customer has access to a number of data migration tools to move their data to another vendor, delete the data and the user accounts from Google and Google will ensure that the data is removed within 180 days from all disks and all backup tapes.

6.12 PERFORMANCE MEASURES AND REPORTING [RFP REFERENCE 8.12]

6.12.1 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

AWS's overall availability SLA is 99.95%. For additional product specific SLAs, please refer to Appendix 3, section 1.

AKAMAI

Akamai provides a 100% platform availability backed by a written SLA. Please refer to Appendix 3, section 2.

GOOGLE

Google's SLA is 99.9% and Google has been operating at 99.98% and higher for the past three years. This is due to the dynamic nature of resource allocation, data center and server assignment and real time replication of data. Please refer to Appendix 3, section 3.

6.12.2 RFP Reference 8.12.2

Provide your standard uptime service and related Service Level Agreement (SLA) criteria.

DLT Response:

AMAZON WEB SERVICES

As described in section 4.4.1 (page 17), Team DLT can pass along AWS specific SLAs to NASPO Purchasing Entities. AWS currently provides SLAs on a service-by-service basis, and not all services currently have fully defined SLAs. The current list of AWS products with SLAs include EC2, S3, CloudFront, Route53, and RDS. Figure 19 provides a list of AWS services, the outlined monthly uptime, and consequences to AWS if specified criteria are not achieved.

Service	Monthly Uptime Percentage	Service Credit
EC2	Less than 99.95% but equal to or greater than 99.0%	10%
	Less than 99.0%	30%
S3	Equal to or greater than 99.0% but less than 99.9%	10%
	Less than 99.0%	25%
	Equal to or greater than 98.0% but less than 99.0% (for S3 Standard - Infrequent Access)	10%
	Less than 98.0% (for S3 Standard - Infrequent Access)	25%
CloudFront	Equal to or greater than 99% but less than 99.9%	10%
	Less than 99%	25%
Route53	Route 53 was not 100% Available: 5-30 minutes	1 Day Service Credit
	Route 53 was not 100% Available: 31 minutes - 4 hours	7 Days Service Credit
	Route 53 was not 100% Available: More than 4 hours	30 Days Service Credit
RDS	Less than 99.95% but equal to or greater than 99.0%	10%
	Less than 99.0%	30%

Figure 19. AWS SLAs and Consequences

AKAMAI

Please refer to Appendix 3, section 2.

GOOGLE

Please refer to Appendix 3, section 3.

6.12.3 RFP Reference 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.

DLT Response:

The DLT Service Center, established in 2005, delivers U.S.-citizen, U.S.-soil technical support, 24x7x365. A public sector-focused support offering, the DLT Service Center includes AWS certified technical support engineers who are always available to provide expertise. The DLT Service Center is a sanctioned AWS support offering with direct-line access to AWS resources, technical support tools, and if needed, AWS backline engineers to support any AWS troubleshooting requirements Purchasing Entities may need. Upon procurement of solutions from DLT, the Purchasing Entity will be provided a DLT Service Center Welcome Package, which includes an 800 number, email address, and customer portal to use for logging cases for support. Please refer to Appendix 10.

AKAMAI

Akamai Customer Care provides the 24x7 single point of contact for all customer support issues for customers live on the platform. All escalations are logged into the central ticketing database that is accessible thru the LUNA Control Center.

At all times, the Purchasing Entity will have access to an on-call team providing best-of-class support. The team is composed of a globally distributed staff with significant professional experience and advanced degrees. Designed to be customer-centric, our technical support process ensures that a Purchasing Entity always has access to a knowledgeable consultant who can solve any problem. Leveraging state-of-the-art support technologies, our consultants address technical needs efficiently and quickly. There are three ways to contact Akamai Customer Care:

- control.akamai.com – Luna Control Center
- E-mail (ccare@akamai.com)
- Telephone 1.877.4.AKATEC [1.877.425.2832]

Customers are granted up to fifteen (15) support requests per year. While the Akamai platform is proven to be very reliable, Purchasing Entities have unlimited access to technical support for any Akamai-related issues that impact the performance or availability of Akamai services. Support for critical P1 service issues must be requested by phone to ensure a timely response.

RESPONSE TIMES

- Critical Impact (P1) 2 Hours
- High Impact (P2) 4 Hours
- Low Impact (P3) 2 Business Days

LIVE SUPPORT AVAILABILITY

- Critical Impact (P1) 24 x 7
- High Impact (P2) Business Hours
- Low Impact (P3) Business Hours

Business Hours are, by geography, as follows, Monday through Friday, excluding local holidays:

- North America (GMT -05:00) - 9:00 AM to 9:00 PM ET
- Europe (GMT) - 9:00 AM to 6:00 PM
- Asia-India (GMT +05:30) - 9:00 AM to 6:00 PM
- Asia-Japan/Singapore (GMT +08:00) - 9:00 AM to 6:00 PM

GOOGLE

The DLT Service Center, established in 2005, delivers U.S.-citizen, U.S.-soil technical support, 24x7x365. A public sector-focused support offering, the DLT Service Center includes Google certified technical support engineers who are always available to provide expertise. The DLT Service Center is a sanctioned Google support offering with direct-line access to Google resources, technical support tools, and if needed, Google backline engineers to support any Google troubleshooting requirements NASPO Purchasing Entities may need. Upon procurement of solutions from DLT, the Purchasing Entity will be provided a DLT Service Center Welcome Package, which includes an 800 number, email address, and customer portal to use for logging cases for support.

6.12.4 RFP Reference 8.12.4

Describe the consequences/SLA remedies if the Respondent fails to meet incident response time and incident fix time.

DLT Response:

The DLT Service Center will maintain our industry leading SLAs for the NASPO Master Agreement. Our SLAs are outlined below and can be tailored to a more aggressive response time if required by a Purchasing Entity. DLT's SLA remediation process includes escalating any and all cases that are in jeopardy of missing an SLA to the highest severity level (S1) until the case is resolved. To date DLT Service Center has never missed an SLA. DLT's vendor-focused support and managed service offering follows our vendor guidelines on SLA requirements, and at this time does not include any consequential contract language.

Severity Level	Classification	Impact	Initial Response
S1	Emergency	Service not available (all users and functions unavailable).	Live transfer to next available engineer. If one is not available call back in 30 minutes
S2	Critical	Significant degradation of Service	1 hour call back
S3	Serious	Limited degradation of service, or urgent new capabilities and/or services	4 business hour call back
S4	Minimal	Small issue, tuning, or non-urgent new capabilities and/or services	Next Business Day

AMAZON WEB SERVICES

Please refer to section 6.12.2. Figure 19 on page 107 details the AWS SLAs and Consequences.

AKAMAI

Akamai provides a 100% uptime SLA. In the event, Akamai does not meet SLA, credits are issued. Based on product purchased, measurable performance metrics and credits are outlined in the product specific SLA documents (please refer to Appendix 3, section 2.)

GOOGLE

Customers are eligible for service credits if Google fails to meet the SLA (please refer to Appendix 3, section 3).

6.12.5 RFP Reference 8.12.5

Describe the firm's procedures and schedules for any planned downtime.

DLT Response:

AMAZON WEB SERVICES

AWS can schedule events for your instances, such as a reboot, stop/start, or retirement. These events do not occur frequently. If a Purchasing Entity's instance will be affected by a scheduled event, AWS sends an email to the email address associated with the Purchasing Entity's AWS account prior to the scheduled event, with details about the event, including the start and end date. Depending on the event, the Purchasing Entity might be able to take action to control the timing of the event.

AKAMAI

Not applicable because Akamai's network is available 24x7x365 with no planned downtime.

GOOGLE

Under normal circumstances, all hardware in any data center—including Google—will eventually need either maintenance or replacement. There are also unforeseen circumstances that can happen to hardware that can cause it to fail in any number of ways. When these events happen at Google, the relevant data is transparently moved from the affected hardware to hardware that is working normally. This is done without any interaction from the customer and no downtime.

6.12.6 RFP Reference 8.12.6

Describe the consequences/SLA remedies if disaster recovery metrics are not met.

DLT Response:

See section 6.12.4.

6.12.7 RFP Reference 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.

DLT Response:

Figure 20 shows an example of real-time performance reporting that will be available to NASPO Purchasing Entities via the DLT Service Center Portal.

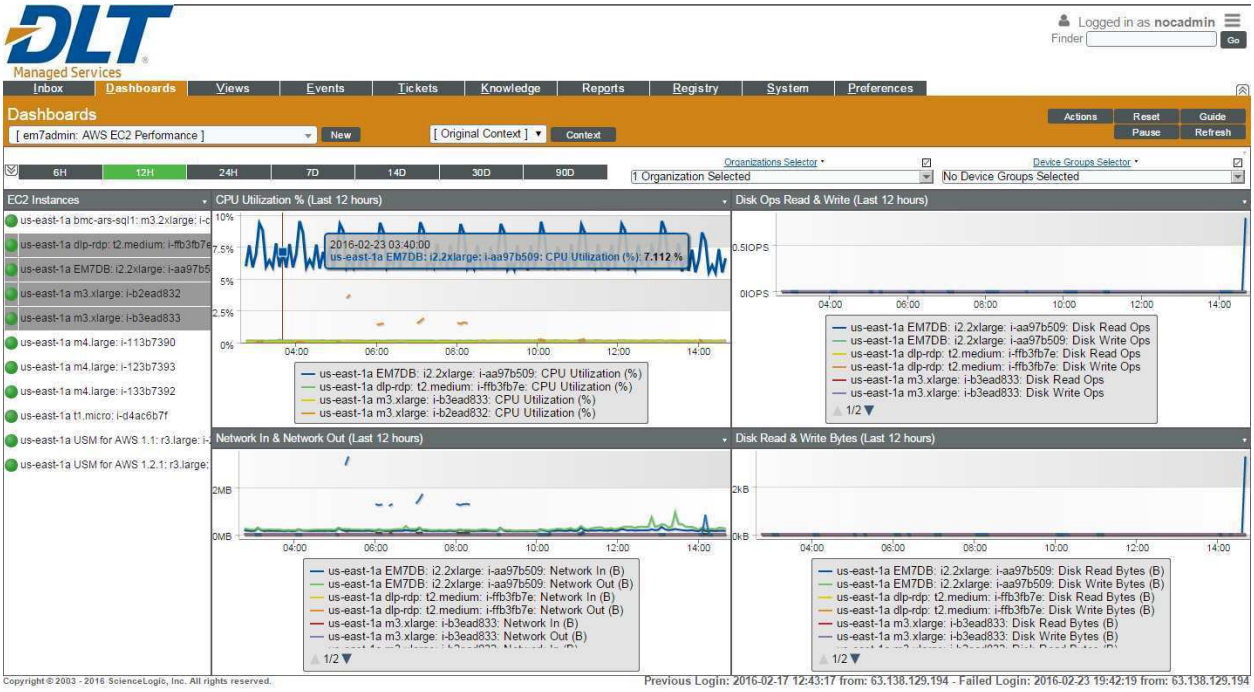

Figure 20. DLT Service Center Portal

AKAMAI

There are several ways to access performance reports:

1. Luna Control Center, our internet portal, provides reports in two formats, HTML or CSV.
2. Configure reports to be automatically sent by email on a recurring basis in either HTML or CSV format.
3. Web Services using SOAP interface instead of a browser.
4. Luna Control Center MIB provides real time statistics and events directly to an end user's Enterprise Network Management System. SNMP MIDB download is available on Windows and Linus systems.
5. Log Delivery Service provides server logs for various services.

Please see Appendix 11 for sample reports.

GOOGLE

Google provides access to the real time Apps Status Dashboard which is freely available over the web.[29]

6.12.8 RFP Reference 8.12.8

Ability to print historical, statistical, and usage reports locally.

DLT Response:

See Figure 21 for sample usage reports. Reports can be customized and tailored to Purchasing Entities' needs. Customers can print locally.

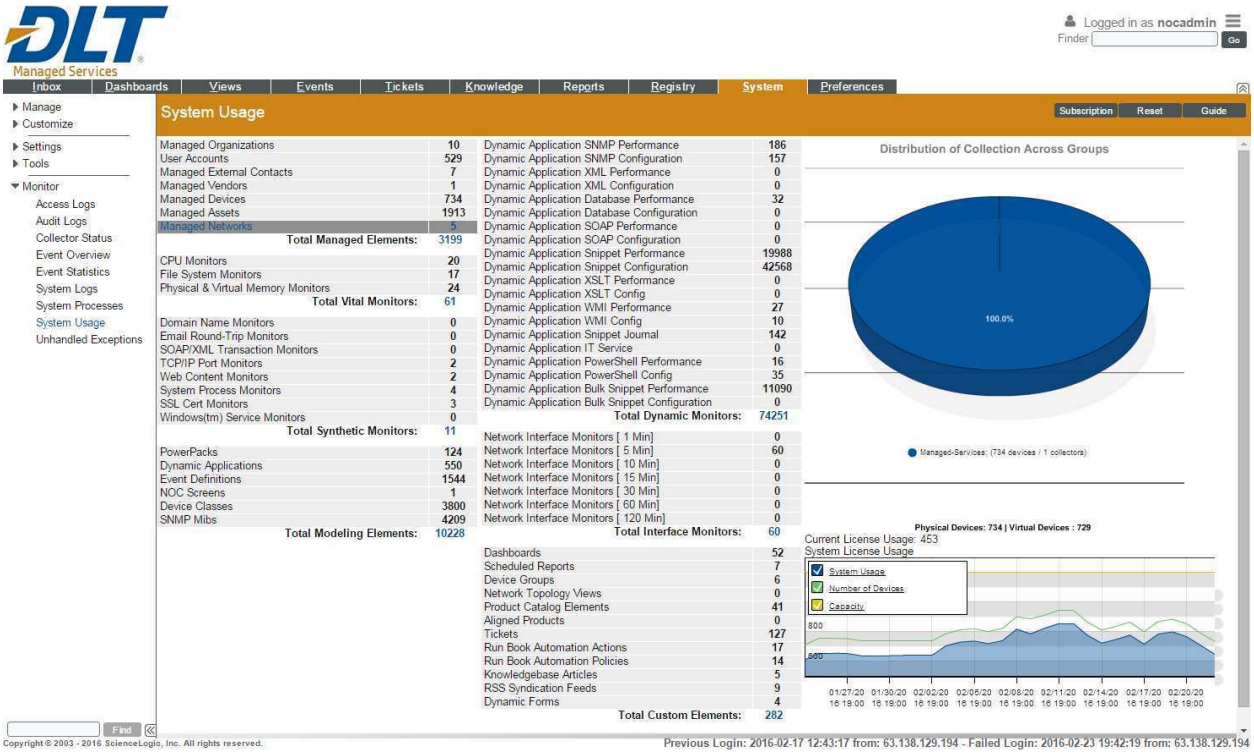

Figure 21. Sample Usage Report

AKAMAI

Printing is done through the Luna Control Center.

GOOGLE

Purchasing Entities' IT Administrators have access to the Google Apps Control Panel which provides on-demand access to a range of reports that can be printed locally.

[29] http://www.google.com/appsstatus#hl=en&v=status

6.12.9 RFP Reference 8.12.9

Offeror must describe whether or not its on-demand deployment is supported 24x365.

DLT Response:

All on-demand services offered by Team DLT in our response to this RFP are supported 24x7x365.

6.12.10 RFP Reference 8.12.10

Offeror must describe its scale-up and scale-down, and whether it is available 24x365.

DLT Response:

AMAZON WEB SERVICES

Elastic Load Balancing and Auto Scaling can automatically scale your AWS cloud-based resources up to meet unexpected demand, and then scale those resources down as demand decreases.

AUTO SCALING

Auto Scaling allows customers to automatically scale their Amazon EC2 capacity up or down according to conditions that they define. Auto Scaling is well suited for applications that experience hourly, daily, or weekly variability in usage. Customers can automatically scale their Amazon EC2 fleet or maintain their Amazon EC2 fleet at a set size.

Auto Scaling enables customers to closely follow the demand curve for their applications, reducing the need to provision Amazon EC2 capacity in advance. For example, Purchasing Entities can set a condition to add new Amazon EC2 instances in increments of three instances to the Auto Scaling Group when the average CPU utilization of the Amazon EC2 fleet goes above 70%. Similarly, Purchasing Entities can set a condition to remove Amazon EC2 instances in the same increments when CPU utilization falls below 10%.

Often, customers may want more time to allow their fleet to stabilize before Auto Scaling adds or removes more Amazon EC2 instances. Purchasing Entities can configure a cooldown period for their Auto Scaling Group, which tells Auto Scaling to wait for some time after taking an action before it evaluates the conditions again. Auto Scaling enables Purchasing Entities to run their Amazon EC2 fleet at optimal utilization.

ELASTIC LOAD BALANCING

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances. It enables customers to achieve even greater fault tolerance in their applications, seamlessly providing the amount of load balancing capacity needed in response to incoming application traffic. Elastic Load Balancing detects unhealthy instances and automatically reroutes traffic to healthy instances until the unhealthy instances have been restored. Purchasing Entities can enable Elastic Load Balancing within a single Availability Zone or across multiple zones for even more consistent application performance.

AMAZON CLOUDWATCH

Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications that customers run on AWS. Purchasing Entities 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 database instances, as well as custom metrics generated by applications and services and any log files your applications generate. Purchasing Entities can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health, using these insights to react and keep applications running smoothly.

Amazon CloudWatch's metrics and alarms can work together with Auto Scaling and ELB to dynamically deploy new instances on-demand, as depicted in Figure 22.

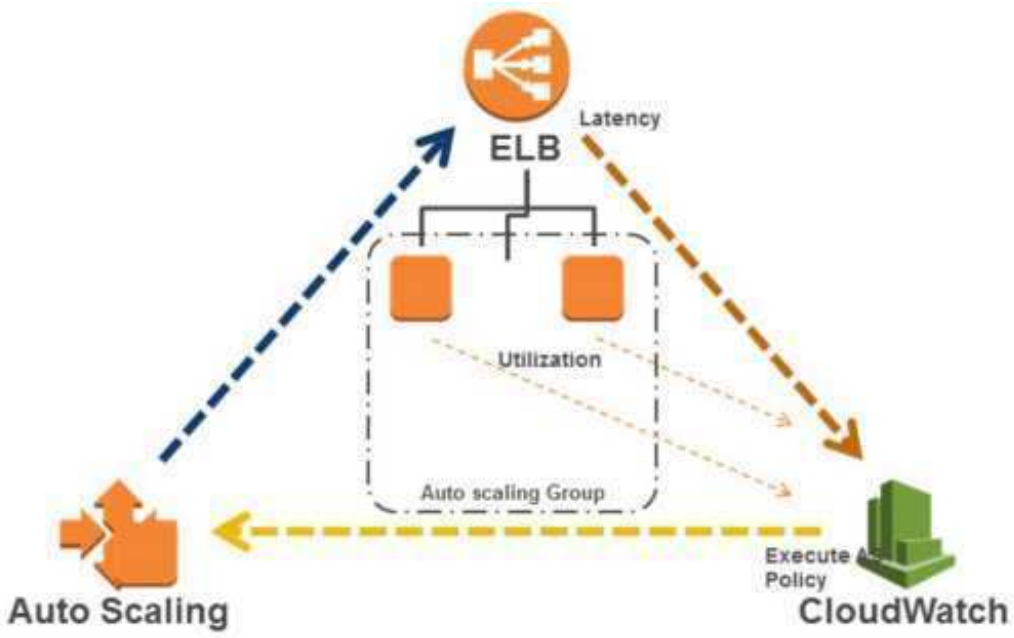

Figure 22. Auto Scaling with ELB and CloudWatch alarms

Auto Scaling is available 24x7x365 and instances can be added or turned down in minutes rather than hours.

AKAMAI

Akamai's network is available 24x7x365 and will scale as needed.

GOOGLE

Purchasing Entities would purchase a set of user licenses and can add new users within the boundaries of those licenses. Scale-up orders can be placed at any time if the Purchasing Entity exceeds or is close to exceeding those licenses. Scale-down orders can be placed at the time of renewals and after the user accounts to be excluded have been deleted or marked as former employee.

6.13 CLOUD SECURITY ALLIANCE [RFP REFERENCE 8.13]

Describe your level disclosure of disclosure with CSA Star Registry for each Solution offered.

a. Completion of a CSA STAR Self-Assessment, as described in Section 5.5.3
b. Completion of Exhibits 1 **and** 2 to Attachment B.
c. Completion of a CSA STAR Attestation, Certification, or Assessment.
d. Completion CSA STAR Continuous Monitoring.

DLT Response:

AMAZON WEB SERVICES

AWS is compliant with Level 1 CSA STAR Registry Self-Assessment. Please refer to AWS' self-assessment found within our Risk and Compliance Whitepaper, page 25-61, Appendix 2. This is the latest CAIQ (v3) released by the CSA.

Per the CSA definitions, AWS aligns with Level 2 via the determinations in our third party audits for SOC and ISO:

- Level 2 Attestation is based on SOC2, which can be requested under NDA.[30] The SOC 2 report audit attests that AWS has been validated by a third party auditor 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 certification is available on the AWS website.[31]

All of the AWS self-assessed assertions within the CSA STAR Registry Self-Assessment are backed by independent, third party audits across multiple compliance programs. We continue to assert we raise the bar on CSA's "attestation" and "certification" program.

AKAMAI

Please see Akamai's completed CAIQ in the file:

- DLT File 04 – Attachment B Exhibit 1 – CAIQ – Akamai.xlsx

GOOGLE

Google completed the Self Assessment and provided completed Exhibit 1 to Attachment B and confirms that Exhibit 2 to Attachment B maps to the security controls as described in Exhibit 1.

Please see Google's completed CAIQ in the file:

- DLT File 05 – Attachment B Exhibit 1 – CAIQ – Google.xlsx

[30] http://aws.amazon.com/compliance/contact/
[31] http://d0.awsstatic.com/certifications/iso_27001_global_certification.pdf

6.14 SERVICE PROVISIONING [RFP REFERENCE 8.14]

6.14.1 RFP Reference 8.14.1

Describe in detail how your firm processes emergency or rush services implementation requests by a Purchasing Entity.

DLT Response:

AMAZON WEB SERVICES

With AWS, the customer has full access to turn up any AWS service at any time with near-zero lag time.

AKAMAI

An Emergency Integration fee is assessed when a customer requires an Akamai service integrated on a faster than normal timeline (less than 10 business days). The Purchasing Entity would contact Akamai Sales /Professional Services Team to initiate emergency request.

GOOGLE

Once Google receives a valid purchase order, the order will be transmitted via the normal channels to Google's Order Management. A request for expedite with justification from the Purchasing Entity will be submitted directly to Google Order Management and the Google Account Management team.

6.14.2 RFP Reference 8.14.2

Describe in detail the standard lead-time for provisioning your Solutions.

DLT Response:

AMAZON WEB SERVICES

None. See section 6.14.1.

AKAMAI

The standard lead time for integration onto the Akamai platform is typically 10-15 business days depending on customer availability, number of services ordered, and complexity of solution.

GOOGLE

Standard lead time is three business days.

6.15 BACK UP AND DISASTER PLAN [RFP REFERENCE 8.15]

6.15.1 RFP Reference 8.15.1

Ability to apply legal retention periods and disposition by agency per purchasing entity policy and/or legal requirements.

DLT Response:

AMAZON WEB SERVICES

AVAILABILITY AND FAULT-TOLERANT DESIGN

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.

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 as described in section 6.5.1. 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 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.

AWS GovCloud[32] (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.

[32] http://aws.amazon.com/govcloud-us/

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.

AKAMAI

Akamai does not store customer's data. In the event of an outage, our caches are restored from the customer's origin website.

GOOGLE

Google Apps Vault in combination with the configuration of the service so that agency users are in distinct Organizational Units (OUs) allows for retention policies to be set by OU.

6.15.2 RFP Reference 8.15.2

Describe any known inherent disaster recovery risks and provide potential mitigation strategies.

DLT Response:

AMAZON WEB SERVICES

See section 6.15.1.

AKAMAI

There are no known disaster recovery risks.

GOOGLE

There are no known disaster recovery risks. Google's data storage design ensures that customer data is replicated across several servers and multiple data centers. If the customer declares a local disaster based on their own facilities where primary access to the Internet is cut off, a good mitigation plan would be to have mobile data services or wireless services as a backup.

6.15.3 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

See section 6.15.1.

Additionally, the AWS cloud supports many popular DR architectures, described in section 6.8.1.

Team DLT understands that there are inherent risks associated with Disaster Recovery (DR). These risks include additional costs, capacity planning, replication failure, loss of data, management of multiple environments, etc. With AWS, NASPO Purchasing Entities can eliminate the majority of these risks such as the need for additional physical infrastructure, off-site 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. The following table provides an overview of the advantages of using AWS for DR as opposed to traditional on-premise DR approaches:

Traditional DR/COOP and Backup Requirements and Issues	AWS Capabilities for DR/COOP/Backup Solutions
• Facilities to house additional infrastructure, including power and cooling. • Security to ensure 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 all mission-critical services, including storage appliances for the supporting data, and servers to run applications and back-end services such as user authentication, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), monitoring, and alerting.	• Fast Performance: Fast, disk-based storage and retrieval of files. • No Tape: Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software. • Compliance: Minimize downtime to avoid breaching Service Level Agreements (SLAs). • Elasticity: Add any amount of data, quickly. Easily expire and delete without handling media. • Security: Secure and durable cloud DR platform with industry-recognized certifications and audits. • Partners: AWS solution providers and system integration partners to help with deployments.

AKAMAI

Akamai has no hosting environment.

GOOGLE

Please refer to section 6.8.1.

6.16 SOLUTION ADMINISTRATION [RFP REFERENCE 8.16]

6.16.1 RFP Reference 8.16.1

Ability of the Purchasing Entity to fully manage identity and user accounts.

DLT Response:

AMAZON WEB SERVICES

NASPO Purchasing Entities can use the AWS Management console to access and manage all AWS resources through a simple and intuitive web-based user interface secured through Secure Socket Layers (SSL) encryption. Team DLT recommends the use of AWS Multi-Factor Authentication (MFA) to secure access to a Purchasing Entity's AWS resources. AWS MFA is a simple best practice that adds an extra layer of protection on top of a typical user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for a user name and password (the first factor—"what you know"), as well as for an authentication code from an AWS MFA device (the second factor—"what you have").

Additionally, Team DLT will help to implement Identity and Access Management (IAM) to control access to a Purchasing Entity's AWS specific services and resources. This will ensure that only users approved by the Purchasing Entity have the ability to access the internal IT resources. Please see section 6.6.7 for more information on IAM.

Team DLT will provide full root credential access to the Purchasing Entities, allowing them to fully manage identity and user accounts. We will work consultatively with the Purchasing Entities to educate them on all of the AWS native tools to achieve this requirement.

AKAMAI

Managing users and roles can be performed through the Luna Control Center.

GOOGLE

User administration is the responsibility of the Purchasing Entity.

6.16.2 RFP Reference 8.16.2

Ability to provide anti-virus protection, for data stores.

DLT Response:

DLT's Managed Services offering provides best of industry anti-virus protection for data stores and provides this offering to the NASPO Purchasing Entities.

AKAMAI

Akamai does not host any data on its platform and anti-virus protection is not needed.

GOOGLE

Google scans all servers and all data stores for viruses.

6.16.3 RFP Reference 8.16.3

Ability to migrate all Purchasing Entity data, metadata, and usage data to a successor Cloud Hosting solution provider.

DLT Response:

AMAZON WEB SERVICES

As the customer owns and controls all of their data within the AWS infrastructure and has complete control of the same, they may migrate the data to any location at their discretion.

DLT is supplying NASPO Purchasing Entities access to the full complement of AWS services. All AWS services are for the use of the Purchasing Entity at their discretion and under the full control of the Purchasing Entity. It is completely at the Purchasing Entity's discretion as to when to initiate end-of-life activities to close down a service (application running on AWS). The Shared Responsibility Model (described in section 6.5.1) remains in place during end-of-life activities. A Purchasing Entity can opt to maintain existing functionality within the AWS platform during migration for redundancy purposes and has complete control of data at all times.

Team DLT recommends the AWS Import/Export as a data migration tool to move large amounts of data into and out of AWS. AWS Import/Export allows for movement of large volumes of data using portable storage devices for transport. We can help Purchasing Entities transfer data to AWS directly using storage devices and Amazon's high-speed internal network, bypassing the Internet. For significant data sets, AWS Import/Export is often faster than Internet transfer and more cost effective than upgrading existing connectivity. With Import/Export, encryption is mandatory, and AWS will encrypt data using the password the Purchasing Entity specified and transfer it onto the device.

AKAMAI

Akamai does not host origin content. All content is placed by the consumer of cloud services on what itself feels is cacheable in storage. As such, Akamai does not rely on origin infrastructure. In fact, many clients switching from one cloud hosting provider to another leverage Akamai to make sure their static content remains up during a migration period. Akamai acts as a "swing site" for those migrating sites keeping the non-transactional elements operating smoothly during a live cutover from one cloud host to another. Moreover, as Akamai is a C Name change, a customer could easily turn Akamai off if it no longer wants the service without interruption. We would simply point them back to the origin for the content after the Akamai contract expired.

GOOGLE

Google offers a number of data export tools that allow customers to pull their data out of the services in use in order to port that data to a successor Cloud hosting solution provider. There are also third party migration tools. Most hosting solution providers have their own data import tools that can read the Google data stores with the appropriate admin credentials provided by the Purchasing Entity.

6.16.4 RFP Reference 8.16.4

Ability to administer the solution in a distributed manner to different participating entities.

DLT Response:

DLT currently supports hundreds of customers with multiple of hundreds of accounts across virtually all areas of the public sector. We are fully capable of administering AWS solutions across multiple Participating Entities.

AKAMAI

Akamai uses Content Provider (CP) Codes for tracking and reporting of delivered content. All CP Codes are tied to one or more service, which are then tracked and reported under that CP Code. The Luna Control Center Traffic, URLs, (Unique) Visitors, Response, Offload, etc. all display traffic information by CP code. CP Codes can be used to separate content or services for tracking and reporting purposes.

GOOGLE

Google Apps allows a Purchasing Entity to organize their users into an organizational unit (OU) hierarchy by agency and/or location. The Admin Roles that the Purchasing Entity chooses to use can be assigned to specific users and constrained to specific OUs. Google Apps is currently implanted across many large scale entities in all sectors public and private. The solution can scale at any level, and DLT is fully capable of administering this solution at scale.

6.16.5 RFP Reference 8.16.5

Ability to apply a participating entity's defined administration polices in managing a solution.

DLT Response:

DLT Solutions supports a large variety of customers, each of which has particular administration policies for managing AWS. DLT has process and backend business systems to fully support our customer's requirements. By working closely with customers, we gain a full understanding of policy and requirements to meet their individual needs.

AKAMAI

Luna Control Center offers customer self-service User Management capabilities. This tool allows customers and partners to create new users, define profiles and set access rights and privileges to deliver the right level of information, often matching the organization's internal business structure.

GOOGLE

With Google Apps, the Admin Panel is accessed by the Purchasing Entity's IT Admin staff which allows them to set Service configuration policies for sharing, email security and compliance filtering and access permissions.

With Google Cloud Platform, the Developer's Console provides tools to control the services and scope as needed.

6.17 HOSTING AND PROVISIONING [RFP REFERENCE 8.17]

6.17.1 RFP Reference 8.17.1

Documented cloud hosting provisioning processes, and your defined/standard cloud provisioning stack.

DLT Response:

AMAZON WEB SERVICES

The AWS Management Console is a single destination for managing all AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. A Purchasing Entity can use the AWS Management Console to perform any number of tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console also enables customers to manage all aspects of their AWS account, including accessing monthly spending by service, managing security credentials, or even setting up new AWS Identity and Access Management (AWS IAM) users. The AWS Management Console supports all AWS regions and lets Purchasing Entities 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, Purchasing Entities 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 file transfers to and from Amazon Simple Storage Service (Amazon S3).

USE EXISTING MANAGEMENT TOOLS

Many of the tools that organizations use to manage on-premises environments can be integrated with AWS as well. Integrating an AWS environment can provide a simpler and quicker path for cloud adoption, because a customer's operations team does not need to learn new tools or develop completely new processes. For example:

- AWS Management Portal for vCenter enables customers to manage their AWS resources using VMware vCenter. The portal installs as a vCenter plug-in within the existing vCenter environment. Once installed, it enables customers to migrate VMware VMs to Amazon EC2 and manage AWS resources from within vCenter. The AWS resources that customers create using the portal can be located in their AWS account, even though those resources have been 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 (VMs) to Amazon EC2. Using the connector, importing a VM is as simple as selecting a VM 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) into which the VM should be imported. Once the VM has been imported, customers can launch it as an instance from the AWS Management Console and immediately take advantage of all the features of Amazon EC2.

- AWS Management Pack for Microsoft System Center[33] enables customers to view and monitor their AWS resources directly in the Operations Manager console. This way, customers can use a single, familiar console to monitor all of their resources, whether they are on-premises or in the AWS cloud. You get a consolidated view of all 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 in the Operations Manager console.

ADDITIONAL AWS RESOURCES

Section 508

In 1998, the Congress of the United States of America amended the Rehabilitation Act to require Federal agencies to make their electronic and information technology accessible to people with disabilities. Inaccessible technology interferes with an individual's ability to obtain and use information quickly and easily. Section 508 was enacted to eliminate barriers in information technology, to make available new opportunities for people with disabilities, and to encourage development of technologies that will help achieve these goals.

The law applies to all Federal agencies when they develop, procure, maintain, or use electronic and information technology. Under Section 508 (29 U.S.C. ' 794d), agencies must give disabled employees and members of the public access to information that is comparable to the access available to others. AWS offers the Voluntary Product Accessibility Template (VPAT) upon request.

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. Additionally, AWS provides two graphical user interfaces, the AWS Management Console and the AWS ElasticWolf Client Console. The AWS ElasticWolf Client Console has incorporated Section 508 requirements and AWS has prepared a Voluntary Product Accessibility Template (VPAT) for the Console, which outlines the Console's accessibility features.

AWS Services

Appendix 4 contains a list of AWS's cloud services. Links are provided to the service website where you can find all service documentation, pricing information, FAQs, release notes, developer information and relevant case studies.

[33] http://aws.amazon.com/windows/products/system-center/

AKAMAI

Not applicable.

GOOGLE

Not applicable to Google Apps Software-as-a-Service offering.

6.17.2 RFP Reference 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).

DLT Response:

AMAZON WEB SERVICES

AWS will provide Purchasing Entities with access to multiple tools and methods for hosting and provisioning AWS resources to include the AWS Management Console, AWS Command Line Interface (CLI), and even existing on-premise management tools. Purchasing Entities can use the AWS Management Console as a single destination for managing all AWS resources to perform a number of tasks, from deploying new applications to monitoring the health of applications. The AWS Management Console can also provide Purchasing Entities the ability to manage all aspects of their AWS account.

Additionally, from the AWS Management Console (Figure 23), Purchasing Entities can use AWS CLI as a unified tool for managing AWS cloud services. With just one tool to download and configure, Purchasing Entities can control multiple AWS resources from the command line and automate them through custom developed scripts.

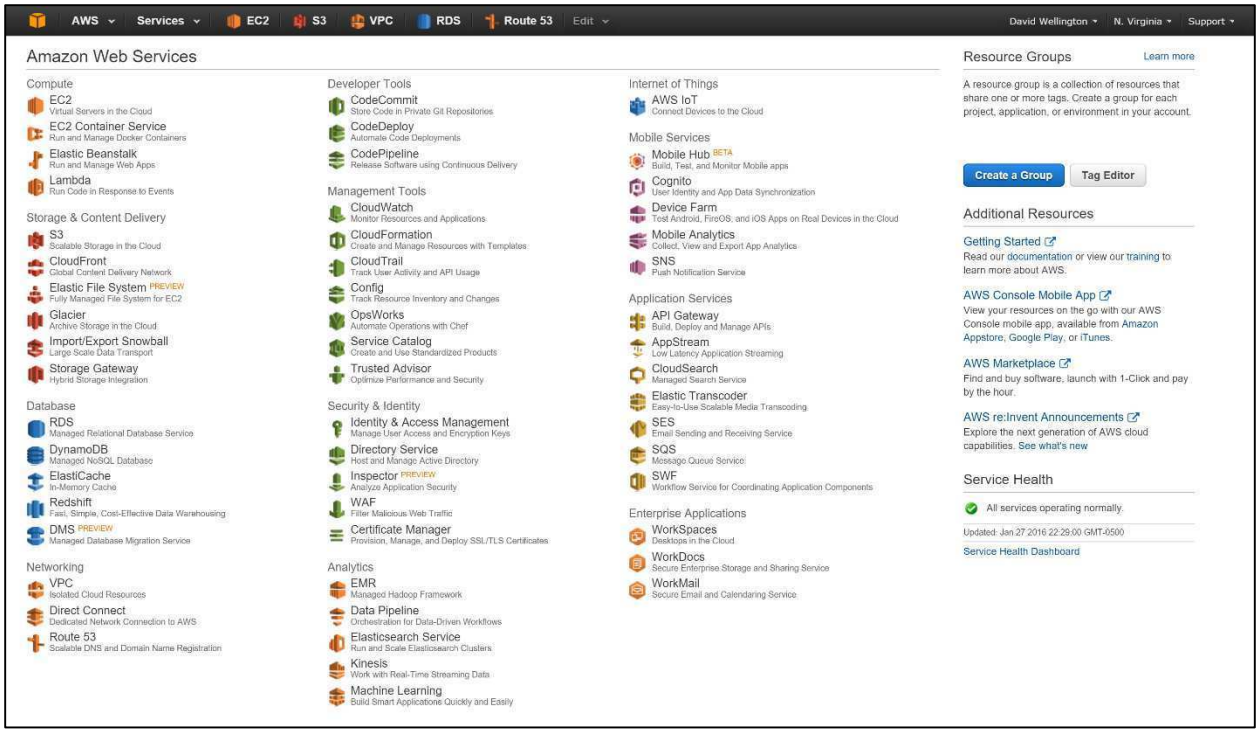

Figure 23. The AWS Management Console

Many of the tools that organizations use to manage on-premises environments can be integrated with AWS as well. Team DLT understands that integrating an AWS environment can provide a simpler and quicker path for cloud adoption, reducing the need for operations teams to learn new tools or develop completely new processes. The examples below provide context as to how we can help integrate existing on-premise tools with AWS:

- AWS Management Portal for VMWare vCenter enables customers to manage their AWS resources using VMware vCenter. The portal installs as a vCenter plug-in within the existing vCenter environment. Once installed, it enables customers to migrate VMware VMs to Amazon EC2 and manage AWS resources from within vCenter. The AWS resources that customers create using the portal can be located in their AWS account, even though those resources have been 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 (VMs) to Amazon EC2. Using the connector, importing a VM is as simple as selecting a VM 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) into which the VM should be imported. Once the VM has been imported, customers can launch it as an instance from

the AWS Management Console and immediately take advantage of all the features of Amazon EC2.

- AWS Management Pack for Microsoft System Center[34] enables customers to view and monitor their AWS resources directly in the Operations Manager console. This way, customers can use a single, familiar console to monitor all of their resources, whether they are on-premises or in the AWS cloud. You get a consolidated view of all 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 in the Operations Manager console.

- ScienceLogic: You can deploy the ScienceLogic platform in under 30 minutes, receiving actionable information from your entire IT stack across public, private and hybrid clouds, under a single "pane of glass." We take this "pane of glass" one step further with our built-in, Smart™ Action enhanced, run book automation enabling you to automatically take actions for quick issue resolution. We use the industry's most advanced discovery engine to intelligently discover your entire IT infrastructure, from your datacenter to your Amazon Web Services resources. Once discovered, our platform automatically and intelligently provides a visual representation of your infrastructure as well as a comprehensive view of your IT and business services in one dashboard. By providing alerting, ticketing, monitoring, asset management and more in a single solution we completely eliminate the need to glue and stitch disparate tools. Our all-in-one solution is "turbo-boosted" by over 1,000 dynamic Apps and protected by the world's only True™ Multi-Tenant system.

    - Hybrid cloud – Now you can see your infrastructure wherever it is. Whether on-premises or in the cloud, reap the benefits of cloud without sacrificing management visibility.
    - See your servers, storage, and networks in the public cloud
    - Monitor over 300 AWS parameters
    - Visibility into Azure cloud services
    - Map dependencies between components with CloudMapper
    - Get to the cloud faster without losing visibility

AKAMAI

- Kona Site Defender (KSD): Defends against all types of DDoS, web application and direct-to-origin attacks.

- Client Reputation: Complements KSD with an additional layer of protection against malicious actors.

- Kona DDoS Defender (KDD): DDoS defense to protect websites against the increasing scale, frequency and sophistication of attacks.

[34] http://aws.amazon.com/windows/products/system-center/

- Prolexic Routed: DDoS defense for protecting data center infrastructure against large, complex attacks.
- Prolexic Connect: DDoS defense for protecting data center infrastructure against large, complex attacks.

GOOGLE

Each Google Apps for Work user can store up to 30 GB of content for free (compared to 15 GB with Google Apps free edition or individual consumer accounts). This storage is shared between Gmail, Google Drive, and Google Photos. If you're approaching the limit or if you've run out of storage, you can purchase additional storage or reduce your storage usage.

The Google Apps Admin console tool provides reports and logs making it easy to examine potential security risks, measure user collaboration, track who signs in and when, analyze administrator activity, and much more. The reports have interactive graphics and tables that show broad, domain-level data alongside granular, user-level details.

6.18 TRIAL AND TESTING PERIODS (PRE- AND POST-PURCHASE) [RFP REFERENCE 8.18]

6.18.1 RFP Reference 8.18.1

Describe your testing and training periods that your offer for your service offerings.

DLT Response:

AMAZON WEB SERVICES

AWS offers a free tier that a customer can use outside of this contract. The AWS free tier is designed to enable you to get hands-on experience with AWS Cloud Services. Customers may access this tier with an account direct with AWS. However, under the AWS Partner Model, the free tier is not applicable to accounts managed by DLT Solutions.

AKAMAI

Training options include regional classroom training, on-line training, and custom on-site training. Testing specific to a customer implementation is provided as a standard part of the available installation/integration services. The integration and on-boarding is a fully managed project including testing periods (i.e., alert mode vs deny mode for attack mitigation) and tuning prior to production to ensure the service is fully operational.

GOOGLE

You can use Google Apps for 30 days free with no commitment. Sign up in minutes and create accounts for up to 10 pilot users in your organization. To get the most out of your trial, we recommend identifying at least three important tasks you want to explore, such as collaborating on documents, launching a video meeting in Hangouts or checking Gmail from your mobile device.

6.18.2 RFP Reference 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.

DLT Response:

AMAZON WEB SERVICES

Customers can stand up a test and/or proof-of-concept environment for acceptance of specific workloads on AWS. Team DLT will consult with the Purchasing Entity to help them establish this environment.

AKAMAI

Akamai will work with Purchasing Entities to determine proof-of-concept (POC) objectives, success criteria, solution scope, configuration scope and implementation plan. Once the solution is on boarded to the Akamai platform, the end user will have access to the online portal, Luna Control Center, for reporting, altering, configuration management, training, helpdesk and documentation. A typical POC is between 30-60 days.

GOOGLE

Google Apps has been commercially available since 2006 and over the years a large number of customers in public and private sector have adopted Google Apps in whole or in part. These customers are the best source of proof in the form of case studies, success stores and direct reference calls. Should a Purchasing Entity require a proof of concept, the success criteria and timeline for the evaluation would first be documented and then, based on the agreed upon scope, a specified number of user accounts can be provided for testing.

6.18.3 RFP Reference 8.18.3

Offeror must describe what training and support it provides at no additional cost.

DLT Response:

AMAZON WEB SERVICES

AWS offers free instructional videos and labs to gain foundational knowledge about key AWS services. Instructional videos explain key concepts and terminology and present a step-by-step console demonstration of an AWS service. Additionally, the DLT Presales Engineering staff will support you in properly architecting your cloud-based systems at no cost.

AKAMAI

Akamai provides training modules available through the Luna Control Center portal. Each module includes in-depth information and software simulations, including a basic overview of Akamai, technical how-to, troubleshooting, and administrations.

Akamai Customer Care provides the 24x7 single point of contact for all customer support issues for customers live on the platform. All escalations are logged into the central ticketing database that is accessible through the Luna Control Center.

GOOGLE

Google provides free online training for Google Apps.[35]

6.19 INTEGRATION AND CUSTOMIZATION [RFP REFERENCE 8.19]

6.19.1 RFP Reference 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.

DLT Response:

Each AWS service can be integrated via AWS RESTful APIs. AWS APIs allow you to programmatically interface resources via tools such as a command-line interface or writing your own application with an SDK.

Team DLT can also leverage our Consulting Methodology to ensure a repeatable process and to provide consistency in service integration for NASPO. The Consulting Methodology includes five main phases: Strategize, Assess, Design, Deploy, and Maintain. Figure 24 and the paragraphs below provide an overview of how the Consulting Methodology enable Team DLT to deliver the complementary service integration of NASPO Purchasing Entities' specific solutions.

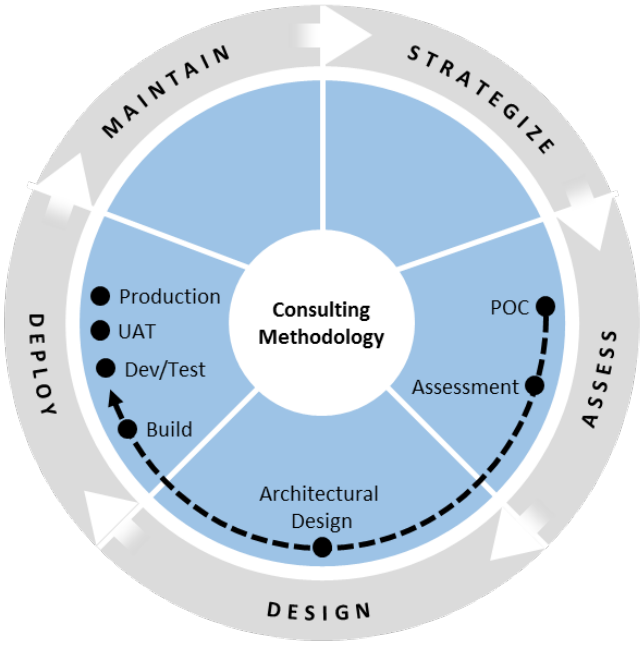

Figure 24. Consulting Methodology

We have categorized activities associated with service integration into the Consulting Methodology with five phases described below:

[35] https://apps.google.com/learning-center/

- Strategize: Team DLT understands that the successful integration of any new technology or service must follow an approved, repeatable, and well-qualified process. Additionally, even the most methodical processes can sometimes be delayed because of organizational hesitation to adopt due to uncertainty or misunderstanding of the technology. A strong and comprehensive strategy can help to diminish these delays and can facilitate organizational buy-in of the adoption. We will help to develop specific strategies for each service integration to include lift-and-shift, code refactoring, etc.

- Assess: The Assess phase is where the team will perform an overall assessment of integration points and provide insight into possible alternatives. During the Assess phase, the team will work with the Purchasing Entity's stakeholders to understand the functionality of services and outline existing capabilities. We will also work with these stakeholders to further understand the needs and technical requirements of existing systems in the cloud. This is where we develop our in-depth knowledge of the As-Is architecture and start to define the expectations of environments. As part of our consultation services, our team will use exploration tools to document and develop a holistic map of the existing environment as well as actionable recommendations. Our recommendations may result in alternative configurations, introduction of new technologies, methodologies, etc. The alternatives assessment may include activities such as unbiased market research, market analysis, and vendor analysis.

  The team may also perform an Independent Verification and Validation (IV&V) of any ongoing cloud solutions. Our approach to performing IV&V analysis will help determine deficiencies in a system or processes and will allow the Purchasing Entity to understand possible new approaches to rectifying issues.

- Design: Throughout this phase, the team will help to develop the necessary schedule, design elements, and recommendations to be used in Purchasing Entities' cloud solution environments. In this phase, the To-Be architectures are developed and socialized with the Purchasing Entity. Additionally, technology considerations will be determined and integration plans developed to ensure proper future implementation. This phase sometimes includes the development of documents such as Communications Plan, Design Documents, To-Be Architecture, and Implementation Plans. We can also help with business case documentation, cost estimates, troubleshooting, and best practices.

- Deploy: The Deploy phase takes all design elements into consideration and provides a path forward into the development of Purchasing Entities' cloud solution environments. The fundamental requirement of this phase is to assist Purchasing Entity staff to achieve the proposed architecture. The Deploy phase is completed when environments are baselined and are ready for daily operation. The Deploy phase also includes providing all aggregated information on each environment and ultimately transitioning the environments system to the Maintain Phase.

- Maintain: The Maintain phase is achieved after transitioning from the Deploy phase. The deployed architecture is turned over to the Purchasing Entity and is fully operational. Team DLT's vendor-authorized support and optional Managed Services are available to the Purchasing Entity to troubleshoot any issues that may arise during this phase.

Additionally, our Subject Matter Experts (SMEs) can provide considerable expertise in the areas of storage, networking and other areas as well as training and education on Cloud platforms to help Purchasing Entities maximize their Cloud investment and reduce risks associated with transitioning from a legacy IT model to the Cloud.

Team DLT understands that repeatable processes and a proven methodology are critical to the successful modernization and migration of any legacy IT system. Our experience with cloud migration has enabled us to develop a multi-step cloud migration methodology that our Professional Services engineering team will leverage to support the migration planning of a Purchasing Entity's platform. As the migration of a Purchasing Entity's solution may differ by platform, service, or application, our team will leverage our consulting methodology to determine the best path for migration planning based on each specific scenario.

AKAMAI

Akamai's CDN service perfectly integrates with any standard SIEM for logging. Also, APIs/interfaces are provided for handling cache purges, DNS zone transfers, etc.

GOOGLE

Google offers open source access to a rich API library. The APIs that are relevant to Google Apps have been collected under the Google Apps Admin SDK. These APIs make it possible to integrate with any SMTP-based solution, file storage or document management system, calendar based solutions, LDAP directories, identity management services and more. Google offers a MAPI plug in for Microsoft Outlook and a Google Drive plug-in for Microsoft Office.

6.19.2 RFP Reference 8.19.2

Describe the ways to customize and personalize the Solutions you provide to meet the needs of specific Purchasing Entities.

DLT Response:

AMAZON WEB SERVICES

Team DLT understands that, as part of the Master Agreement, we will be required to provide Purchasing Entities with highly customized and personalized services to meet their individual requirements. To ensure complete customization of our solutions, we leverage a 4-step Assessment Lifecycle to help scope and customize IT-based engagements and provide consistency in delivery. This thorough lifecycle also allows for organization buy-in from all Purchasing Entities' stakeholders prior to implementation. Our Assessment lifecycle has been

established through our understanding of industry best practice, lessons learned, and strong subject matter expertise in cloud based IT implementations.

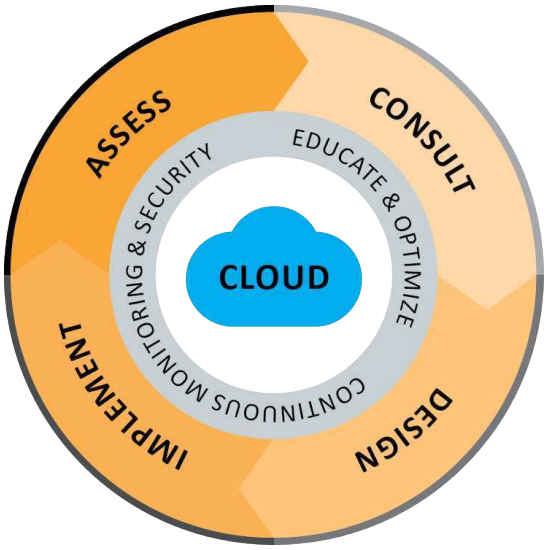

Figure 25. Team DLT's Assessment Lifecycle

Team DLT's 4-step assessment lifecycle is described below:

- Assess: During the Assess phase, we will work with Purchasing Entities' stakeholders to understand and define their specific environment and organizational needs. Additionally, we will work with these stakeholders to define their business and technical requirements. We look for technical specifications and requirements such as application dependencies, and server and storage sizing. We will work directly with Purchasing Entities to outline the specific scope of the assessment. During this phase our team will collect critical project and system information such as application refactoring, networking configuration, connectivity requirements, server dependencies, etc. We leverage existing templates from our document library to ensure repeatability of assessments if required by a Purchasing Entity. Examples of these templates include user questionnaires, server sizing worksheets, etc.

- Consult: During the Consult phase, we leverage the in-depth analysis provided during the Assess phase to provide advisory and consultation services on best path forward. We provide suggestions on technical considerations to include changes in server, storage, and network requirements. Our team will work with Purchasing Entities to understand implications of modifying specific IT components and define a path of least resistance. Our team works to become the trusted advisor through our unique skillsets in the cloud space and our strong experiences of similar previous engagements. Examples of work products that come from the Consult phase may include As-Is architecture diagrams, system inventory, etc.

- Design: We will design and customize a solution specific to the Purchasing Entities through the understanding we have developed in the Assess and Consult phases. Design

considerations include the use of existing IT systems, integrating the Purchasing Entity's environments, etc. Throughout this phase, we will leverage all information we obtained in the Assess and Consult phases to develop any additional architecture and design elements required for successful service delivery of the IT systems. Examples of these design elements may include To-Be architecture diagrams and AWS machine image configurations, server specifications, etc.

- Implement: The Implementation phase takes all outputs from the previous phases into consideration and provides a final path forward into the implementation of the project. The fundamental requirement of this phase is to implement the agreed-upon design. The end result is a highly customized and repeatable process that can be used by other Purchasing Entities.

AKAMAI

Based on the specific requirements, Akamai Professional Services can design and develop solutions that can be deployed at the Akamai edge—making the development cycle faster and cost effective. Our experts use proven methodologies to help move faster, unleashing the power of the Akamai Platform for your business. Akamai experts have helped some of the most dynamic, interactive and engaging Web properties on the Internet keep up with the changing online environment and stay lightning fast at the same time. Some of the capabilities include performance testing and tuning, automated workflow, online event management, project management, third party load testing, end user validation, failover solutions, custom reporting and analytics, geo-targeting, mobile device identification, A/B testing, user prioritized access, and private content delivery networks. Security capabilities include traffic filtering, configuration assistance, end user validation, geo-targeting, proactive security event monitoring, security event after action reporting, security incident management, threat update reviews, table top attack drills and customized vulnerability scanning and analysis.

GOOGLE

Google Apps can be customized by a Purchasing Entity in terms of branding but overall the means to deliver a SaaS solution at scale is to offer a single version for all customers. Training can be customized to meet specific business use cases.

6.20 MARKETING PLAN [RFP REFERENCE 8.20]

Describe how you intend to market your Solutions to NASPO ValuePoint and Participating Entities.

DLT Response:

NASPO ValuePoint and Participating Entities will benefit significantly from DLT's marketing expertise. DLT distinguishes itself in the marketplace through its government-focused, uniquely insightful marketing practice, which includes outreach and education. We believe that our high volume, cost-efficient direct marketing format allows us to maintain a pricing advantage over many other product resellers and integrators. With the high level of vendor partnerships that DLT maintains, we can offer NASPO Purchasing Entities increased savings through product specials

and promotions as well as increased exposure through sales contests targeted at expanding contract awareness.

DLT has a long history of marketing to the state, local and education markets. DLT leverages more than 190,000 SLED contacts to conduct state and local year-end campaigns with the specific goal of helping them spend their IT budgets.

Upon award of the NASPO contract, DLT will seek approval from NASPO to implement a Marketing Plan that can include such activities as:

- Public Relations: With NASPO's approval, upon award of the contract, DLT will create and distribute a press release highlighting the addition of the NASPO contract to our contract portfolio to all of DLT's media contacts. NASPO may approve content prior to distribution, if required. DLT will also create and distribute press releases highlighting significant wins ($1M+) that come from the NASPO contract. Lastly, DLT will use NASPO contract wins to identify unique stories that can be turned into case studies and subsequent press opportunities.

- NASPO Contract page on DLT website: The Contracts section of DLT's website lists our current contract vehicles (see Figure 26). Upon award, DLT will create a page for the NASPO contract that will include DLT's point of contact, contract details, ordering instructions, and a link to NASPO ValuePoint eMarket Center. DLT will first obtain NASPO's approval to include contract information on the DLT website.

- Trade Shows: Participation and/or sponsorship in national trade shows and conferences and targeted state and local events, including pre-and post-event communications and follow-up.

- Tech Days: In addition to participating in events hosted by third-party organizations, DLT will develop direct events that highlight the NASPO contract and NASPO-hosted events.

- Webinars/Virtual Event: Over the years, DLT has leveraged webinars to highlight and promote customer issues and the products/solutions that solve those issues. The webinars (one-off or series) are live and then hosted on-demand on our website.

- Email Campaigns: DLT has invested heavily in Eloqua, an email marketing automation tool by Oracle. Using this tool, DLT will create periodic email campaigns highlighting the NASPO contract, the products and services that are included on the contract as well as information about events, tech days and any other pertinent news.

- Advertising: DLT will work with media outlets focused on state, local and education to advertise the NASPO contract and the solutions on the contract. DLT will also work with these outlets to leverage their lead generation programs.

- Social Media/Web Marketing: DLT is the leading vendor in the public sector in leveraging social media and pay-per-click (PPC) marketing to attract new business. DLT will create a social media campaign leveraging more than 1,000 followers on LinkedIn and Twitter to promote the NASPO contract and the products and services included on the contract. DLT

will also create PPC campaigns on Google, Bing and Yahoo using targeted keywords to promote the NASPO contract and the products and services included on the contract.

- Sales Collateral: DLT will create a state, local and education one-page information sheet that highlights the NASPO contract and the products and services included on the contract. The information sheet will be also be available electronically and can be used as a follow-up to a sales call or any in-person events DLT Solutions hosts or attends.

- Call Blitzes: The DLT Sales team will conduct periodic call blitzes highlighting the NASPO contract and the products and services included on the contract. The call blitzes will be conducted in unison with the email campaigns and will include 35-50 calls by each sales representative. On occasion, DLT will also be able to leverage non-state and local sales representatives for call blitzes.

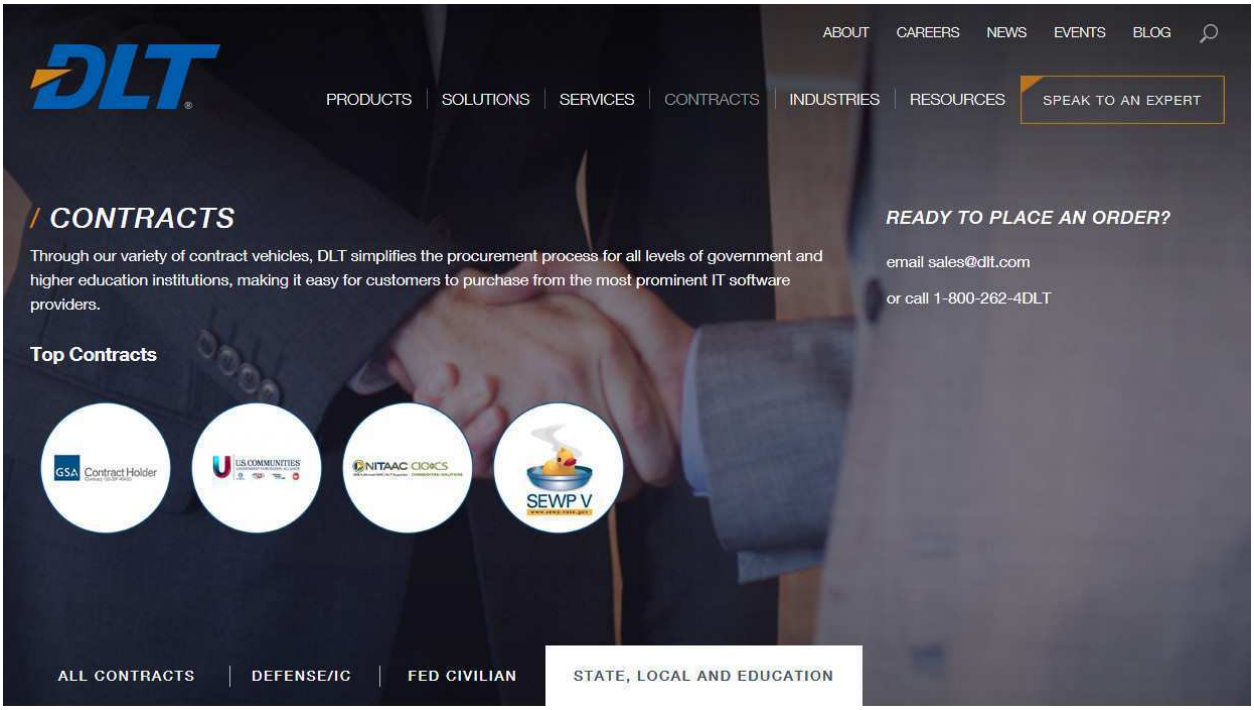

Figure 26. The Contracts page of DLT's website[36]

[36] http://www.dlt.com/government-it-contracts

6.21 RELATED VALUE-ADDED SERVICES TO CLOUD SOLUTIONS [RFP REFERENCE 8.21]

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.

DLT Response:

DLT provides the following value-added services for all proposed solutions.

DLT SALES ENGINEERING

This group is the sales team's technical encyclopedia, representing how our solutions solve your specific needs. This group serves five primary functions:

- Sales Consultancy: Help you understand, compare, and contrast the solutions that are available for purchase and troubleshoot problems with existing implementations to ensure that the solutions work successfully.
- Product and Service Engineering: Modify, improve, or introduce advanced technology to customers who may have an application but who have not yet acquired knowledge of the material or technique in question.
- Education and Adoption: Advise, teach and mentor on how best to use a product or service, develop pilots and proof of concepts, and provide demos on the product or solution.
- Tailored Solutions: Collaborate with vendor engineering to determine how a product or services could be implemented to suit your needs.
- Solution Engineering: Address specific needs, problems and requirements through the design of specific information systems or application using documentation that describes the structure and behavior of a solution as well as the process to deliver it.

Purchasing Entities may engage Sales Engineering whenever they are interested in new or enhanced capabilities. This team is staffed with highly experienced engineers and architects who possess deep product knowledge and significant certifications.

SERVICE OPERATIONS SUPPORT OVERVIEW

The DLT Operations Center, which is a state of the art facility geared towards Service Center and Managed Services operations, will be used to fulfill NASPO's value-added services requirements in this RFP. The DLT Operations Center includes both the DLT Service Center and DLT Managed Services operations.

The DLT Service Center, established in 2005, delivers U.S.-citizen, U.S.-soil technical support, 24x7x365. A public sector-focused support offering, the Service Center includes certified technical support engineers who are always available to provide expertise. The DLT Service Center is a sanctioned vendor support offering with direct-line access to vendor resources,

technical support tools, and if needed, vendor backline engineers to support any troubleshooting requirements.

The DLT Managed Services offering comprises the proven DLT management methodology and suite of tools, coupled with DLT certified engineering oversight, giving NASPO Purchasing Entities the ability to have DLT monitor, manage, and provision their environments to meet the requirements in this RFP via a safe, secure and methodical means.

EXPERT SERVICES

This team provides product consulting, implementation services and project delivery around core technology competencies. The Expert Services team leverages comprehensive product knowledge, delivery best practices, and strategic partnerships to migrate and integrate your capabilities into the cloud.

Our Cloud offerings are designed to mentor customers through increasing capacity and adding capabilities on the fly without investing in new infrastructure, training new personnel, or licensing new software.

This team includes many of the technology experts that will help you get value from our technology solutions. DLT offers staff expertise such as Project Managers, Solution and Technology Architects, Subject Matter Experts, Business Analysts and other implementation skills.

6.22 SUPPORTING INFRASTRUCTURE [RFP REFERENCE 8.22]

6.22.1 RFP Reference 8.22.1

Describe what infrastructure is required by the Purchasing Entity to support your Solutions or deployment models.

DLT Response:

AMAZON WEB SERVICES

AWS does not require any up-front hardware to compute in the cloud. Cloud computing lets you focus on your own users, rather than on the heavy lifting of racking, stacking, and powering servers. Purchasing Entities can enhance the AWS offering in areas such as VPC connectivity by utilizing their on-premise routers to add private communication between on-premise and cloud networks.

AKAMAI

No infrastructure is required for this service.

GOOGLE

No infrastructure is required other than access to the Internet and the use of a modern browser. Optional services that integrate with LDAP services and identity management services and can handle data migration would call for locally sourced and installed Windows or Linux-based servers.

6.22.2 RFP Reference 8.22.2

If required, who will be responsible for installation of new infrastructure and who will incur those costs?

DLT Response:

AMAZON WEB SERVICES

AWS-provided services do not require additional hardware installation costs.

AKAMAI

As the service does not require any customer provided infrastructure, there is no installation required.

GOOGLE

All infrastructure is managed by Google. Customer costs are only incurred based on subscription use of the services.

6.23 ALIGNMENT OF CLOUD COMPUTING REFERENCE ARCHITECTURE [RFP REFERENCE 8.23]

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).

DLT Response:

All DLT's proposed technologies comply with the NIST Cloud Computing Reference Architecture. Please see Appendix 1 Products/Service Model Matrix.

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

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/

For a more detailed description of AWS Compliance please see
AWS Compliance page: 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**.

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 FedRAMP[sm] compliance please see the AWS FedRAMP[sm] 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://d0.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 assessmen**t 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://d0.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://d0.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℠ 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 FedRAMPsm 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 FedRAMPsm 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 99.999999999% 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 *Application Security*	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 *Customer Access Requirements*	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 *Data Integrity*	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 *Data Security / Integrity*	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 *Audit Planning*	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 *Independent Audits*	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. 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 - 02.5	Do you conduct external audits regularly as prescribed by industry best practices and guidance?	
	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 *Information System Regulatory Mapping*	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 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.
	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 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 *Business Continuity Planning*	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 *Business Continuity Testing*	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 *Power / Telecommunications*	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 *Environmental Risks*	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 *Equipment Location*	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 *Equipment Maintenance*	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 *Equipment Power Failures*	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 *Impact Analysis*	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 *Policy*	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 *Retention Policy*	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/.
	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?	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 -11.5	Do you test your backup or redundancy mechanisms at least annually?	
Change Control & Configuration Management *New Development / Acquisition*	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 *Outsourced Development*	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. 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.
	CCC-02.2	Do you have controls in place to detect source code security defects for any outsourced software development activities?	
Change Control & Configuration *Management Quality Testing*	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/. 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 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. 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.
	CCC-03.2	Is documentation describing known issues with certain products/services available?	
	CCC-03.3	Are there policies and procedures in place to triage and remedy reported bugs and security vulnerabilities for product and service offerings?	
	CCC-03.4	Are mechanisms in place to ensure that all debugging and test code elements are removed from released software versions?	
Change Control & Configuration Management *Unauthorized Software Installations*	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 *Production Changes*	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 *Classification*	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 *Data Inventory / Flows*	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 *eCommerce Transactions*	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. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
	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)?	
Data Security & Information Lifecycle Management *Handling / Labeling / Security Policy*	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 *Nonproduction Data*	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 *Ownership / Stewardship*	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 *Secure Disposal*	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 *Asset Management*	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 *Controlled Access Points*	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 *Equipment Identification*	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 *Offsite Authorization*	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 *Offsite equipment*	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 *Policy*	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 *Secure Area Authorization*	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 *Unauthorized Persons Entry*	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. AWS Physical Security Mechanisms are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Datacenter Security *User Access*	DCS -09.1	Do you restrict physical access to information assets and functions by users and support personnel?	
Encryption & Key Management *Entitlement*	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 *Key Generation*	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/). 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. 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 FedRAMP.
	EKM -02.2	Do you have a capability to manage encryption keys on behalf of tenants?	
	EKM -02.3	Do you maintain key management procedures?	
	EKM -02.4	Do you have documented ownership for each stage of the lifecycle of encryption keys?	
	EKM -02.5	Do you utilize any third party/open source/proprietary frameworks to manage encryption keys?	
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. 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 FedRAMP.
	EKM-04.2	Are your encryption keys maintained by the cloud consumer or a trusted key management provider?	
	EKM-04.3	Do you store encryption keys in the cloud?	
	EKM-04.4	Do you have separate key management and key usage duties?	
Governance and Risk Management *Baseline Requirements*	GRM-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.
	GRM-01.2	Do you have a capability to continuously monitor and report the compliance of your infrastructure against your information security baselines?	
	GRM-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 *Risk Assessments*	GRM-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.
	GRM-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 *Management Oversight*	GRM-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 *Management Program*	GRM-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/.
	GRM-04.2	Do you review your Information Security Management Program (ISMP) least once a year?	
Governance and Risk Management *Management Support / Involvement*	GRM-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). AWS manages third-party relationships in alignment with ISO 27001 standards. AWS Third Party requirements are reviewed by independent external
Governance and Risk Management *Policy*	GRM-06.1	Do your information security and privacy policies align with industry standards (ISO-27001, ISO-22307, CoBIT, etc.)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	GRM-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/.
	GRM-06.3	Can you provide evidence of due diligence mapping of your controls, architecture and processes to regulations and/or standards?	
	GRM-06.4	Do you disclose which controls, standards, certifications and/or regulations you comply with?	
Governance and Risk Management *Policy Enforcement*	GRM-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. 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.
	GRM-07.2	Are employees made aware of what actions could be taken in the event of a violation via their policies and procedures?	
Governance and Risk Management *Business / Policy Change Impacts*	GRM-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 *Policy Reviews*	GRM-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.
	GRM-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 *Assessments*	GRM-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
	GRM-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*	GRM-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. 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.
	GRM-11.2	Do you make available documentation of your organization-wide risk management program?	
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. 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-03.2	Do you document employee acknowledgment of training they have completed?	
	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 *Employment Termination*	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 *Portable / Mobile Devices*	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 *Nondisclosure Agreements*	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 *Roles / Responsibilities*	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 area available at: http://aws.amazon.com/security and http://aws.amazon.com/compliance.

Control Group	CID	Consensus Assessment Questions	AWS Response
Human Resources *Acceptable Use*	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 *Training / Awareness*	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 *User Responsibility*	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 *Workspace*	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 *Audit Tools Access*	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 *User Access Policy*	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 *Diagnostic / Configuration Ports Access*	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 *Policies and Procedures*	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 *Segregation of Duties*	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 *Source Code Access Restriction*	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 *Third Party Access*	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 *User Access Restriction / Authorization*	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 *User Access Authorization*	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 *User Access Reviews*	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 *User Access Revocation*	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 *User ID Credentials*	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?	
	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?	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 -12.9	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 *Utility Programs Access*	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 *Audit Logging / Intrusion Detection*	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. 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.
	IVS-01.5	Are audit logs reviewed on a regular basis for security events (e.g., with automated tools)?	
Infrastructure & Virtualization Security *Change Detection*	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 *Clock Synchronization*	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.
Infrastructure & Virtualization Security *Capacity / Resource Planning*	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 *Management - Vulnerability Management*	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 *Network Security*	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. 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. 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-06.4	Are all firewall access control lists documented with business justification?	
Infrastructure & Virtualization Security *OS Hardening and Base Controls*	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?	
Infrastructure & Virtualization Security *Production / Nonproduction Environments*	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/.
	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 *Segmentation*	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 *VM Security - vMotion Data Protection*	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 *VMM Security - Hypervisor Hardening*	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 *Wireless Security*	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 *Network Architecture*	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. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security.
Interoperability & Portability *Data Request*	IPY-02	Is unstructured customer data available on request in an industry-standard format (e.g., .doc, .xls, or .pdf)?	
Interoperability & Portability *Policy & Legal*	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 *Standardized Network Protocols*	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 *Virtualization*	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 *Anti-Malware*	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 *Application Stores*	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 *Approved Applications*	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 *Approved Software for BYOD*	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 *Awareness and Training*	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 *Cloud Based Services*	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 *Compatibility*	MOS -07	Do you have a documented application validation process for testing device, operating system and application compatibility issues?	
Mobile Security *Device Eligibility*	MOS -08	Do you have a BYOD policy that defines the device(s) and eligibility requirements allowed for BYOD usage?	
Mobile Security *Device Inventory*	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 *Device Management*	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 *Encryption*	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 *Jailbreaking and Rooting*	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 *Legal*	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 *Lockout Screen*	MOS -14	Do you require and enforce via technical controls an automatic lockout screen for BYOD and company owned devices?	
Mobile Security *Operating Systems*	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 *Passwords*	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 *Policy*	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 *Remote Wipe*	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 *Security Patches*	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 *Users*	MOS -20.1	Does your BYOD policy clarify the systems and servers allowed for use or access on the BYOD-enabled device?	
	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 *Contact / Authority Maintenance*	SEF-01.1	Do you maintain liaisons and points of contact with local authorities in accordance with contracts and appropriate regulations?	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.
Security Incident Management, E-Discovery & Cloud Forensics *Incident Management*	SEF-02.1	Do you have a documented security incident response plan?	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.
	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 *Incident Reporting*	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 *Incident Response Legal Preparation*	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 *Incident Response Metrics*	SEF-05.1	Do you monitor and quantify the types, volumes and impacts on all information security incidents?	AWS Security Metrics are monitored and analyzed in accordance with ISO 27001 standard. Refer to ISO 27001 Annex A, domain 16 for further details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	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 *Data Quality and Integrity*	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 *Incident Reporting*	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 *Network / Infrastructure Services*	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 *Provider Internal Assessments*	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 *Third Party Agreements*	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 *Supply Chain Governance Reviews*	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 *Supply Chain Metrics*	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 *Third Party Assessment*	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 *Third Party Audits*	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 *Antivirus / Malicious Software*	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 *Vulnerability / Patch Management*	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 *Mobile Code*	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 an 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 FedRAMP[sm].
	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 99.999999999% 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 webservers 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 FedRAMP[sm]. 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 FedRAMP[sm].
	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.

FedRAMP[sm]: The Federal Risk and Authorization Management Program (FedRAMP[sm]) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMP[sm] 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 uses 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: FedRAMP[sm], 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

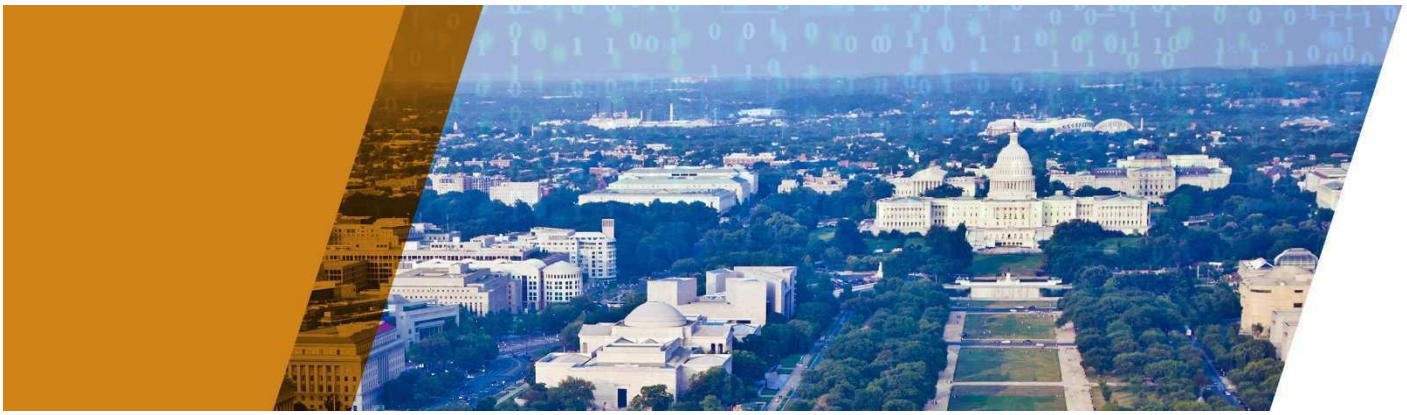

THE STATE OF UTAH
DIVISION OF PURCHASING

NASPO ValuePoint Master Agreement for Cloud Solutions

Utah Solicitation Number CH16012

Technical Proposal – Appendix 3

PRESENTED BY

March 10, 2016

DLT Solutions / 2411 Dulles Corner Park, Suite 800 / Herndon VA 20171 / 703.709.7172 / www.dlt.com
TIN 54-1599882 / DUNS # 78-6468199 / CAGE Code 0S0H9 / GSA Schedule 70 GS-35F-4543G

CONTENTS

1.0 AMAZON WEB SERVICES (AWS) SERVICE LEVEL AGREEMENTS

NOTE: Due to the rapidly evolving nature of AWS's product offerings, the most current SLAs are best reviewed directly on the AWS website via the links below:

- Amazon EC2 SLA: http://aws.amazon.com/ec2-sla/
- Amazon S3 SLA: http://aws.amazon.com/s3-sla
- Amazon CloudFront SLA: http://aws.amazon.com/cloudfront/sla
- Amazon Route 53 SLA: http://aws.amazon.com/route53/sla
- Amazon RDS SLA: http://aws.amazon.com/rds-sla

1.1 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:
    - For Amazon EC2, when all of your running instances have no external connectivity.
    - 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 by the end of the second billing cycle after which the incident occurred and must include:

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.

1.2 AMAZON S3 SERVICE LEVEL AGREEMENT

Last Updated September 16, 2015

This Amazon S3 Service Level Agreement ("SLA") is a policy governing the use of Amazon Simple Storage Service ("Amazon S3") 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 S3. 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 S3 available with the applicable Monthly Uptime Percentage (as defined below) during any monthly billing cycle (the "Service Commitment"). In the event Amazon S3 does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Error Rate" means: (i) the total number of internal server errors returned by Amazon S3 as error status "InternalError" or "ServiceUnavailable" divided by (ii) the total number of requests for the applicable request type during that five minute period. We will calculate the Error Rate for each Amazon S3 account as a percentage for each five minute period in the monthly billing cycle. The calculation of the number of internal server errors will not include errors that arise directly or indirectly as a result of any of the Amazon S3 SLA Exclusions (as defined below).

- "Monthly Uptime Percentage" is calculated by subtracting from 100% the average of the Error Rates from each five minute period in the monthly billing cycle.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon S3 account.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for Amazon S3 for the billing cycle in which the error occurred in accordance with the schedule below.

For all requests not otherwise specified below:

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99.0% but less than 99.9%	10%
Less than 99.0%	25%

For requests to Amazon S3 Standard – Infrequent Access (Standard-IA):

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 98.0% but less than 99.0%	10%
Less than 98.0%	25%

We will apply any Service Credits only against future Amazon S3 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 error 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 S3 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 by the end of the second billing cycle after which the incident occurred and must include:

1. the words “SLA Credit Request” in the subject line;

2. the dates and times of each incident of non-zero Error Rates that you are claiming; and

3. 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 applicable to the month of such request is confirmed by us and is less than the applicable 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 S3 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon S3, or any other Amazon S3 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 S3; (iii) that result from any actions or inactions of you or any third party; (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); or (v) arising from our suspension and termination of your right to use Amazon S3 in accordance with the AWS Agreement (collectively, the “Amazon S3 SLA Exclusions”). If availability is impacted by factors other than those used in our calculation of the Error Rate, then we may issue a Service Credit considering such factors at our discretion.

1.3 AMAZON CLOUDFRONT SERVICE LEVEL AGREEMENT

Last Updated June 1, 2013

This Amazon CloudFront Service Level Agreement ("SLA") is a policy governing the use of Amazon CloudFront 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 CloudFront. 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 CloudFront available with a Monthly Uptime Percentage (defined below) of at least 99.9% during any monthly billing cycle (the "Service Commitment"). In the event Amazon CloudFront does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Error Rate" means: (i) the total number of internal server errors returned by Amazon CloudFront divided by (ii) the total number of requests during that five minute period. We will calculate the Error Rate for each Amazon CloudFront account as a percentage for each five minute period in the monthly billing cycle. The calculation of the number of internal server errors will not include errors that arise directly or indirectly as a result of any of the Amazon CloudFront SLA Exclusions (as defined below).

- "Monthly Uptime Percentage" is calculated by subtracting from 100% the average of the Error Rates from each five minute period in the monthly billing cycle.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for Amazon CloudFront for the billing cycle in which the error occurred in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99% but less than 99.9%	10%
less than 99%	25%

We will apply any Service Credits only against future Amazon CloudFront 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 error 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 CloudFront 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 by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each incident of non-zero Error Rates that you are claiming; and

iii. 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 applicable to the month of such request is confirmed by us and is less than 99.9%, 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 CloudFront SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon CloudFront, or any other Amazon CloudFront 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 CloudFront; (iii) that result from any actions or inactions of you or any third party; (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) arising from our suspension and termination of your right to use Amazon CloudFront in accordance with the AWS Agreement; (vi) that result from exceeding usage limits stated in the Amazon CloudFront documentation; or (vii) that result from use of an origin server other than Amazon S3 (collectively, the **“Amazon CloudFront SLA Exclusions”**). If availability is impacted by factors other than those used in our calculation of the Error Rate, then we may issue a Service Credit considering such factors at our discretion.

1.4 AMAZON ROUTE 53 SERVICE LEVEL AGREEMENT

Last Updated May 15, 2015

This Amazon Route 53 Service Level Agreement ("SLA") is a policy governing the use of Amazon Route 53 (including Private DNS) 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 Route 53. 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 Route 53 100% Available (defined below). In the event Amazon Route 53 does not meet the foregoing commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "100% Available" means that Amazon Route 53 did not fail to respond to your DNS queries during a monthly billing cycle.
- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon Route 53 account.

Service Credits

Service Credits are calculated based on 1 day of Service Credit, which is equal to your average daily Amazon Route 53 query charges for the monthly billing cycle preceding the monthly billing cycle in which the period that Amazon Route 53 was not 100% Available occurred, and are available as follows:

Duration Amazon Route 53 was not 100% Available	Service Credit
5 - 30 minutes	1 day Service Credit
31 minutes - 4 hours	7 days Service Credit
More than 4 hours	30 days Service Credit

We will apply any Service Credits only against future Amazon Route 53 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 error 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 Route 53 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 by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each period that Amazon Route 53 was not 100% Available that you are claiming; and

iii. 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 period that Amazon Route 53 was not 100% Available is confirmed by us, 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 Route 53 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon Route 53, or any other Amazon Route 53 performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of Amazon Route 53’s reasonable control, including any force majeure event, denial-of-service attack or Internet access or related problems beyond the demarcation point of Amazon Route 53; (iii) that result from any actions or inactions of you or any third party; (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) arising from our suspension and termination of your right to use Amazon Route 53 in accordance with the AWS Agreement; (vi) that result from you exceeding usage limits stated in the Amazon Route 53 documentation (collectively, the “Amazon Route 53 SLA Exclusions”); or (vii) that, with respect to public DNS only, result during a period that you were not using all four virtual name servers (for example, ns123.awsdns.com, ns123.awsdns.net, ns123.awsdns.co.uk and ns123.awsdns.org) assigned to your “hosted zone”. If availability is impacted by factors other than those used in our calculation of 100% Available, then we may issue a Service Credit considering such factors at our discretion.

1.5 AMAZON RDS SERVICE LEVEL AGREEMENT

Last Updated June 1, 2013

This Amazon RDS Service Level Agreement ("SLA") is a policy governing the use of the Amazon Relational Database Service ("Amazon RDS") under the terms of the AWS 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 RDS. 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 Multi-AZ instances available with a Monthly Uptime Percentage (defined below) of at least 99.95% during any monthly billing cycle (the "Service Commitment"). In the event Amazon RDS does not meet the Monthly Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Monthly Uptime Percentage" for a given Multi-AZ instance is calculated by subtracting from 100% the percentage of 1 minute periods during the monthly billing cycle in which the Multi-AZ instance was "Unavailable". If you have been running that Multi-AZ instance for only part of the month, your Multi-AZ instance is assumed to be 100% available for the portion of the month that it was not running. Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon RDS SLA Exclusion (defined below).

- "Multi-AZ instance" means an Amazon RDS for MySQL,Oracle or PostgreSQL database instance with the Multi-AZ parameter set to true.

- "Unavailable" means that all connection requests to the running Multi-AZ instance fail during a 1 minute period.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Credits

Service Credits are calculated as a percentage of the charges paid by you for the Multi-AZ instances that did not meet the Monthly Uptime Percentage commitment in a billing cycle 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%	25%

We will apply any Service Credits only against future Amazon RDS 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 or non-performance or other failure by us to provide Amazon RDS 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 will need to submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each Unavailability incident you are claiming;

iii. the DB Instance IDs and the AWS Regions of the affected Multi-AZ instances; and

iv. 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 the request occurred. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon RDS SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon RDS, or any other Amazon RDS 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 RDS;

(iii) that result from any voluntary actions or inactions from you or any third party (e.g., rebooting a database instance, scaling compute capacity, not scaling storage when the storage is full, misconfiguring security groups, VPC configurations or credential settings, disabling encryption keys or making the encryption keys inaccessible, etc.);

(iv) that result from instances belonging to the Micro DB instance class or other instance classes which have similar CPU and memory resource limitations;

(v) that result from you not following the basic operational guidelines described in the Amazon RDS User Guide (e.g., overloading a database instance to the point it is inoperable, creating excessively large number of tables that significantly increase the recovery time etc.);

(vi) caused by underlying database engine software that lead to repeated database crashes or an inoperable database instance;

(vii) that result in long recovery time due to insufficient IO capacity for your database workload;

(viii) 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); or

(ix) that result from any maintenance as provided for pursuant to the AWS Agreement; or

(x) arising from our suspension and termination of your right to use Amazon RDS in accordance with the AWS Agreement (collectively, the “Amazon RDS SLA Exclusions”).

If availability is impacted by factors other than those explicitly used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

2.0 AKAMAI SERVICE LEVEL AGREEMENTS

2.1 AKAMAI CLOUD SECURITY SOLUTIONS SERVICE LEVEL AGREEMENTS (FEBRUARY 2016)

Definitions

"Attack Monitoring Services - Failure to Notify Event" is an event in which Akamai fails to take the defined steps to notify Customer within a period of 15 minutes from the time that Akamai's Security Operations Center (SOC) receives a Critical alert (applicable only to Prolexic Application-Based Monitoring and Prolexic Flow-Based Monitoring Services deployed at the Customer site).

"**Akamai Network**" means the distributed network owned and operated by Akamai.

"**Akamai Prolexic Network**" means the distributed network of specialized network of scrubbing centers owned and operated by Akamai.

"Availability Outage" (applicable only to Kona Site Defender and Kona DDoS Defender) is defined as a period of at least two consecutive failed attempts six (6) minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the contracted security service for the day in which the failure occurred, not to exceed 30 days of fees.

"Emergency Maintenance" means any activity that Akamai, in its sole discretion, deems necessary to correct an immediate threat to the ongoing availability and quality of Akamai's Service offerings

"Managed Kona Site Defender Response Service SLA Violation" – the inability of the Akamai support team to meet the Response time or Live Support Availability as defined.

"**Managed Kona Site Defender Service Initial Response Time"** (applicable only to Support Requests filed against the Kona Site Defender product under Managed Kona Site Defender Service) is the time it takes Customer to get a response on the reported issue from an Akamai technical support representative.

The measurement of the initial response time is the elapsed time from the start of the Security Incident Management process, to the response to Customer by an appropriate service resource to acknowledge the request, respond with a service request number and begin working the issue. This includes time until a response is received in the form of a call back or e-mail or any other customer facing communication. The degree of urgency can vary based upon the issue's priority level.

For Security Events identified through Akamai Managed Kona Site Defender Service Security Event monitoring, Security Incident Management begins once a Security Event has been observed and that event cannot be classified as a false positive and the issue is escalated to the customer.

For Security Events identified by the customer, the Security Incident Management process begins from the time the event is reported by the customer to Akamai Customer Care.

All Support Requests reported via e-mail will be considered as Severity 3.

"Security Severity Levels"

- **"Severity 1"** -- Critical Impact: This class exhibits: a) loss or outage on any portion of a protected property, b) data breach (exfiltration or infiltration) confirmed in progress, or c) defacement of a protected property.

- **"Severity 2"** – Major Impact: This class exhibits: a) degradation in performance on any portion of a protected property, b) suspected data breach, or c) excessive bot activity that may lead to intellectual property compromise.

- **"Severity 3"** – Low Impact: This class exhibits: a) signs of a potential small-scale security incident (log event evidence of malicious traffic that does not impact the origin and may be false positive, b) is a proactive action; "heightened attention" in response to a public threat, for instance, c) includes a possible fraud investigation without immediate evidence of data breach, or d) low-level site scraping activity.

"**Service Outage"** (applicable only to Prolexic Routed, Prolexic Connect, and Prolexic Proxy) means that Akamai's Prolexic Network did not respond to DNS or HTTP queries or the forwarding of IP traffic for more than sixty (60) consecutive seconds.

Akamai's Time To Mitigate and Consistency of Mitigation Service Levels (applicable only to Prolexic Routed, Prolexic Connect, and Prolexic Proxy)

With respect to Customers subscribing to Prolexic Routed, Prolexic Connect, and Prolexic Proxy, Akamai offers a service level ("Service Level") committing to the length of time that it will take Akamai to effectively deploy mitigation.

The Service Level begins at the time that a critical alert is generated by Akamai for Customers subscribed to the standard Always-On mitigation service or for Customers who are otherwise permitted to be running traffic through Akamai's Prolexic Network when a DDoS attack is identified. The time of the critical alert will be determined by relevant ticket correspondence and/or critical alerts in the Akamai portal.

The Service Level for Customers subscribed to an On-Demand mitigation service, if not currently routed through the Akamai Prolexic Network, begins after a Customer notifies Akamai and properly routes traffic through Akamai's Prolexic Network during a DDoS attack. The Time to Mitigate ("TTM") value for these On-Demand Customers depends upon the length of time for the Customer to properly route traffic through Akamai's Prolexic Network, and the length of time it takes for routes to propagate to the Internet at large.

Akamai's Service Level for te following attack types is available exclusively to Prolexic Routed, Prolexic Connect and Prolexic Proxy Services Customer. At a minimum, Service Validation is required once annually. Akamai commits to the following TTM, for each DDoS attack type, as categorized per following:

Attack Type	TTM – Time to Mitigate (typical)	TTM – Time to Mitigate – Guaranteed (Service Level)
* UDP/ICMP Floods	1 minute or less	5 minutes
* SYN Floods	1 minute or less	5 minutes
* TCP Flag Abuses	1 minute or less	5 minutes
* GET/POST Floods	10 minutes or less*	20 minutes
* DNS Reflection	5 minutes or less**	10 minutes
* DNS Attack	5 minutes or less**	10 minutes

** Mitigation requiring traffic analysis and custom signature deployment*

Akamai's Time To Mitigate and Consistency of Mitigation Service Levels (applicable only to Kona DDoS Defender and Managed Kona Site Defender Service)

With respect to Customers subscribing to Kona DDoS Defender and Managed Kona Site Defender Service, Akamai offers a service level ("Service Level") committing to the length of time that it will take Akamai to effectively deploy mitigation, meaning:

- Initial mitigations have been deployed
- They have been effective at mitigating the impact of the immediate attack.
- The benefits of the mitigation were evident within the time window of the SLA.

The Service Level begins at the time that a critical alert is generated by Akamai for Customers subscribed and integrated to the standard Always-On mitigation service when a DDoS attack is identified. The time of the critical alert will be determined by relevant ticket correspondence and/or critical alerts in the Akamai portal.

Akamai's Service Level only for the following attack types is available exclusively to Kona DDoS Defender and

Managed Kona Site Defender Service Customers. At a minimum, a Table Top Drill for Kona DDoS Defender, or a Threat Update Review and a Table Top Drill for Managed Kona Site Defender Service is required once annually and Akamai's Security Specialist recommendations must have been applied to the configuration. Akamai commits to the following TTM, for each DDoS attack type, as categorized per following:

Attack Type	TTM – Time to Mitigate (typical)	TTM – Time to Mitigate – Guaranteed (Service Level)
* UDP/ICMP Floods	1 minute or less	5 minutes
* SYN Floods	1 minute or less	5 minutes
* TCP Flag Abuses	1 minute or less	5 minutes
* GET/POST Floods	10 minutes or less*	20 minutes

Attack Type	TTM – Time to Mitigate (typical)	TTM – Time to Mitigate – Guaranteed (Service Level)
* DNS Reflection	5 minutes or less**	10 minutes
* DNS Attack	5 minutes or less**	10 minutes

** Mitigation requiring traffic analysis and custom signature deployment*
*** Applies to DNS attacks targeting Akamai IP addresses*

Akamai's Consistency of Mitigation Service Level (applicable only to Prolexic Routed, Prolexic Connect, Prolexic Proxy, Kona DDoS Defender and Managed Kona Site Defender Service)

Akamai offers a 95% Consistency of Mitigation Service Level. Consistency of Mitigation is measured by analyzing the ratio of clean traffic to attack traffic that is forwarded to the Customer. Measurement of the Consistency of Mitigation parameter begins after the committed TTM has elapsed. Claims against the Consistency of Mitigation Service Level must be submitted with a packet capture of at least one hour in duration, identifying the total amount of attack traffic forwarded during the event envelope. The event envelope is defined as all or part of the period between the TTM Service Level period and the end of the attack. Evidence of forwarding of attack traffic in excess of 5% of the total traffic volume qualifies for a credit under this Service Level clause.

Remedy for Time to Mitigate and Consistency of Mitigation Service Levels

The TTM is based from the time that traffic is properly routed through Akamai's Prolexic Network or Akamai Network for On-Demand Customers or from the time that a critical alert is generated for services that are AlwaysOn or already routed through Akamai's Prolexic Network or Akamai Network. The TTM is measured based upon the Consistency of Mitigation Service Level terms. During any given calendar month, if Akamai fails to meet the TTM Service Level as measured by the Consistency of Mitigation parameters set forth above, the following credits will be issued:

- Single event – in the event that the TTM Service Level is exceeded – with mitigation not meeting the Consistency of Mitigation Service Level, Akamai will credit Customer's account for such month for the pro-rated charges as follows:
    - Less than one hour: for (1) day of Monthly Service Fees due in respect of the affected Network Protection Services
    - For one hour or more, and less than 6 hours: two (2) days of Monthly Service Fees payable in respect of the affected DDoS Mitigation Services; and
- Multiple Events or Single Event lasting more than 6 hours – in the event that the Time to Mitigate Service Level is exceeded – with mitigation not meeting the Consistency of Mitigation Service Level for a period of six (6) hours or more, or during four (4) or more events within a calendar month, Customer will be credited with seven (7) days of Monthly Service Fees payable in respect of the affected DDoS Mitigation, or Managed Kona Site Defender Services.

All Customers must have successfully completed a Table Top Drill, with any prefix(es) affected, within the previous twelve months in order to qualify for remedy credit under the Time to Mitigate and Consistency of Mitigation Service Levels.

Akamai's Service Availability Service Level (applicable to Prolexic Connect, Prolexic Routed, and Prolexic Proxy Service Outage)

Akamai will provide any credits to Customer (or to Reseller for transfer to Customer, if Reseller is the contracting entity) per the following: Provided Customer reports a Service Outage to Akamai promptly following the occurrence of an event of interruption in Service that Customer believes is a Service Outage, but in any event no later than five (5) days after the event took place, Customer shall be entitled to receive a service credit for Customer's benefit in accordance with the schedule below. Whether an interruption in Services constitutes a Service Outage shall be determined solely by Akamai in its sole good faith discretion supported by records, data and other evidence. If a Service Outage has taken place and Customer notifies Akamai as provided in this Section, Akamai shall provide a credit to Customer as follows:

(i) If a particular Service Outage reported by Customer lasted for more than one minute but less than four (4) consecutive hours during a calendar month, Akamai will credit Customer for such month, the pro-rated charges for one (1) day of Monthly Service Fees of the amount of revenue Akamai receives from Customer with respect to the affected DDoS Mitigation Service(s); or

(ii) If a particular Service Outage reported by Customer lasted for four (4) or more consecutive hours during a calendar month, a credit equal to two (2) days of the Monthly Service Fees payable of the amount of revenue Akamai receives from Customer with respect to the affected DDoS Mitigation Service(s).

The above provision sets forth Customer's sole and exclusive remedy for Service Outages and any other interruptions or failures of Akamai's Managed DDoS Mitigation Service.

Remedy – Akamai Attack Monitoring Services (applicable only to Application-Based and Flow-Based Monitoring Services)

A Customer subscribing to the Akamai Application-Based Monitoring or Flow-Based Monitoring Service is entitled to remedy credit in accordance with this subsection should an Attack Monitoring Services - Failure to Notify Event occur, provided Customer reports the incident to Akamai promptly following the occurrence of an event that Customer believes is an Attack Monitoring Services - Failure to Notify Event, but in any event no later than five (5) calendar days after the event. Whether an incident constitutes an Attack Monitoring Services - Failure to Notify Event shall be determined by Akamai in its sole good faith discretion supported by records, data and other evidence.

(i) If an Attack Monitoring Services - Failure to Notify Event occurs once or more times during a calendar month, Akamai will credit Customer's account for the pro-rated charges for one (1) day's Monthly Service Fees due for each incident, in respect of the affected site(s)' Services; and

(ii) In addition to Customer being entitled to the above credits, in the event that three or more Attack Monitoring Services - Failure to Notify Events occur during a calendar month, Customer shall

have the right, for 30 days following the start of such incident, to terminate the affected Service, without liability.

Akamai's Availability and Performance Service Level (applicable only to Kona Site Defender, Kona DDoS Defender, Web Application Firewall, Bot Manager and Site Shield)

- **Availability SLA:** Akamai offers a service level ("Service Level") committing to 100% availability of the contracted security service.

  The Service Level begins at the time the customer has successfully completed integration and Service Validation for the protected properties, or the contract Billing Effective Date, whichever is later.

- **Performance SLA:** Akamai offers a service level ("Service Level") committing that the security service will not impede origin performance in any period that the protected digital property is not under attack.

  The Service Level begins at the time the customer has successfully completed integration and Service Validation for the protected properties, or the contract Billing Effective Date, whichever is later.

Activation of the Availability and Performance Service Level Agreements occurs once the Customer has successfully completed the following: Customer must enter and indicate the location of two valid test files for the same object (as described in II(c and (d above into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal . Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and, for customers using the optional Site Shield solution, Akamai Professional Services. Customers using Remote Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect five (5 business days after the Customer enters valid test files into the SLA Activation Tool.

Remedy for Availability and Performance Service Levels (applicable only to Kona Site Defender, Kona DDoS Defender, Web Application Firewall, Bot Manager and Site Shield)

If the Service fails to meet the defined service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly security service fee for the day for the protected origin(s) in which the failure occurs, not to exceed 30 days of fees.

The following methodology will be employed to measure the availability and performance of the security service:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's protected origin servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations using a test file on the customer's protected origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the protected origin server (ie, http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Kona Site Defender or the optional Site Shield solution.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the protected origin server (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the protected Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's protected production origin (measured directly, or via a Site Shield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the protected property, for that day in which the failure occurs, not to exceed 30 days of fees.

Akamai's Managed Kona Site Defender Service Response Service Level Agreement (applicable only to Managed Kona Site Defender Service)

Akamai agrees to provide a level of service to Customer's purchasing Managed Kona Site Defender Service as follows:

- Response Time
    - Severity 1 ≤ 30 minutes (must be opened via phone)
    - Severity 2 ≤ 1 hour
    - Severity 3 ≤ 1 business day
- Live Support Availability: An Akamai representative will be available live on the phone to respond to Severity 1 (Critical Impact) and Severity 2 (Major Impact) Service issues 24 hours a day, 7 days a week and 365 days a year. Live Support Availability for severity 3 (Low Impact) cases will be available during normal business hours, Monday through Friday, excluding local holidays, in the following geographies as follows:
    - North America (GMT – 05:00): 9:00 am to 9:00 pm ET
    - Europe (GMT): 08:00 am to 5:00 pm
    - Asia-India (GMT + 05:30): 9:00 am to 6:00 pm
    - Asia-Japan/Singapore (GMT + 08:30): 9:00 am to 6:00 pm

Remedy for Managed Kona Site Defender Service Response SLA Violation

In the event of a Managed Kona Site Defender Service Response SLA Violation, Customer must submit a written request for a credit (email request acceptable) to Customer's applicable Akamai relationship manager within seven days of the alleged SLA Violation. For acknowledged SLA Violations, Customer will receive (as its sole remedy) a credit equal to Customer's monthly Managed Kona Site Defender Service fee for the day in which the failure occurs, not to exceed 30 days of fees per month.

Remedy Terms – General

In order for Akamai to issue a credit in accordance with this SLA, Customer must have an account that is current with payments and in good standing with Akamai, and must be able to confirm that Customer has completed the Integration process for the applicable Service and, if applicable, all Competing Mitigation Hardware and Techniques have been disabled or removed during any mitigation services.

Credits shall only apply for Services provided pursuant to the Monthly Service Fee and/or Monthly Service Overage Fee, and will not apply to any other Service. Customers with subscriptions for more than one DDoS Mitigation Service will only receive credits for affected portion of DDoS Mitigation Service(s). The aggregate credits to be provided in any calendar month shall not exceed 25% of the Monthly Service Fee in respect of the affected Service(s).

2.2 DYNAMIC SITE DELIVERY SERVICE LEVEL AGREEMENT (07/18/2007)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

(a) 100% Uptime: The Service will serve content 100% of the time.

(b) Measurable Performance Enhancement: The Service will deliver content measurably faster than the Customer's web site.

(c) Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject), or via an Akamai SiteShield region if the customer is using the SiteShield feature.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a SiteShield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a SiteShield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a SiteShield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate the Dynamic Site Delivery Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and, for customers using SiteShield, Akamai Professional Services. Customers using Remote SiteShield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the SiteShield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

2.3 FASTDNS SERVICE SERVICE LEVEL AGREEMENT

I. Service Level and Credits

Akamai agrees to provide a level of service to the FastDNS customer providing:

(a) 100 % Uptime: For customer owned DNS zones configured to use the FastDNS service, Akamai will serve DNS resolutions 100% of the time.

(b) Service Credits: If the FastDNS service experiences an outage as defined below as a result of a failure of any component of the FastDNS service operated, managed or maintained by Akamai, the FastDNS customer will receive a Service Credit. A Service Credit is defined as equal to the pro rata FastDNS service fees for the continuous duration of the outage.

In the event that the FastDNS service SLA has not been met, the FastDNS customer shall request a credit via email or in writing to the customer's Account Manager at Akamai within 7 days of the outage.

II. FastDNS Outage Definition

(a) An outage is defined as a 12-minute period of consecutive failed attempts from multiple networks to resolve a DNS record configured for DNS resolution by the FastDNS service.

(b) Testing methodology for the FastDNS service must meet the following requirements:

- A minimum of five (5) testing agents in diverse geographic and network topological locations.
- A sample of at least two (2) data points per testing agent (i.e. a minimum polling interval less than 12 minutes on each testing agent).
- Test results must demonstrate that each testing agent was successful performing a DNS query for the DNS zone to the root DNS servers (.com, .net, .gov, etc.), but unsuccessful performing a DNS query for the DNS zone to the FastDNS service.

2.4 GLOBAL TRAFFIC MANAGEMENT SERVICES SERVICE LEVEL AGREEMENT (04/12/2006)

I. Service Levels and Credits

Akamai agrees to provide a level of service demonstrating:

a. 100% Uptime: Every client request through the Akamai Global Traffic Management service will be directed to a live Customer server 100% of the time.

b. Credits: If the Akamai Global Traffic Management service fails to meet the above service level, the Customer will receive (as its sole remedy) a credit equal to fees for the day in which the failure occurs.

II. SLA Monitoring Methodology

The following methodology will be employed to measure Global Traffic Management service availability:

Agents and Polling Frequency

a. Akamai will perform a customer-specified liveness test against each of the Customer's production servers from at least five (5) geographically and network-diverse locations in major metropolitan areas. Liveness testing occurs at the customer-specified frequency.

b. Akamai will also perform one DNS resolution via the Akamai Global Traffic Management service. Resolutions will take place at least once every 15 minutes.

c. There will be a separate test for each property (sub-domain) registered with Global Traffic Management.

III. Outages

An "Outage" is defined as a failed attempt by the Akamai Global Traffic Management service to resolve the name of a registered property to a live Customer server when there is at least one (1) live Customer web server. If an Outage is identified by this method, the customer will receive (as its sole remedy) a credit equivalent to the fees for the day in which the outage occurred.

2.5 AKAMAI ION PREMIER NORTH AMERICA DOMESTIC SERVICE LEVEL AGREEMENT (MAY 2014)

I. Service Levels

Akamai agrees to provide a level of service for customers with origin infrastructure hosted in North America demonstrating:

(a) Improvement Over Origin: The daily average page delivery time using the Ion Premier service will be at least 30% faster than the daily average delivery time for the same page delivered from the customer's origin server (defined as the daily average page delivery time from origin being at least 1.3 times the daily average page delivery time using Akamai's Ion Premier service).

(b) 100% Availability: The Service will serve content 100% of the time.

II. SLA Monitoring Methodology

A. Performance SLA Monitoring Methodology

The following methodology will be employed to measure the performance improvement provided by the Service:

(1) A single static unauthenticated page (including HTML and associated embedded content), provided by Akamai, will be tested throughout the term of the Ion Premier Order Form for purposes of this SLA.

(2) Delivery times will be tested using Akamai's Site Analyzer testing services to measure the daily average page delivery time computed from the complete global set of available Site Analyzer measurement agents. The tests will be configured to take place once an hour from the measurement agents.

(3) This SLA assumes that there will be no material changes to the test content including, without limitation, metadata applied to the content, agent network used for testing, origin settings and origin infrastructure; any such changes will nullify a deficient test result.

(4) This SLA shall not apply if there is limited or no performance improvement due to causes originating from customer's infrastructure or a third party's infrastructure outside of Akamai's control, including the customer's DNS that provides the CNAME into the Akamai network.

Akamai reserves the right to remove DNS times from the speedup computation if deemed necessary.

B. Availability SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability:

Agents and Polling Frequency

(1) From at least six (6) geographically and network-diverse locations in major metropolitan areas,

(2) Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(3) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com). One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject). The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(4) The Akamaized test content must use a TTL of 2 hours or greater.

(5) The test content will be a file of approximately 10 KB in size.

(6) Polling will occur at approximately 6-minute intervals.

(7) Based on the http GET operations described in II-B above, the response times received from the two sources, (a) the Customer server (directly), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outage

An availability outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly).

IV. SLA Activation

A. Performance SLA Activation

To activate this SLA, Customer and Akamai will designate the applicable page as contemplated by II-A (1) above. Customer shall be responsible for configuring such page to enable testing by the measurement agents. Customer shall be responsible for configuring the Site Analyzer service for measurement of the designated page.

B. Availability SLA Activation

In order to activate the Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(B) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. SLA Escalation

In order to request a credit for a perceived service failure, Customer must, within five calendar days (120 hours) after the perceived failure, contact Akamai Customer Care in writing (which may be via email (ccare@akamai.com) or through the EdgeControl Management Center), specifying the time period in which the failure is believed to have occurred.

VI. Remedies

If the Service fails to meet the service levels in Section I or if an Outage, as defined in Section III, is identified, the Customer will receive (as its sole remedy) a credit equal to Customer's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

2.6 AKAMAI ION STANDARD NORTH AMERICA DOMESTIC SERVICE LEVEL AGREEMENT (MAY 2014)

I. Service Levels

Akamai agrees to provide a level of service for customers with origin infrastructure hosted in North America demonstrating:

(1) Improvement Over Origin: The daily average page delivery time using the Ion Standard service will be at least 30% faster than the daily average delivery time for the same page delivered from the customer's origin server (defined as the daily average page delivery time from origin being at least 1.3 times the daily average page delivery time using Akamai's Ion Standard service).

(2) 100% Availability: The Service will serve content 100% of the time.

II. SLA Monitoring Methodology

A. Performance SLA Monitoring Methodology

The following methodology will be employed to measure the performance improvement mprovided by the Service:

(1) A single static unauthenticated page (including HTML and associated embedded content), provided by Akamai, will be tested throughout the term of the Ion Standard Order Form for purposes of this SLA.

(2) Delivery times will be tested using Akamai's Site Analyzer testing services to measure the daily average page delivery time computed from the complete global set of available Site Analyzer measurement agents. The tests will be configured to take place once an hour from the measurement agents.

(3) This SLA assumes that there will be no material changes to the test content including, without limitation, metadata applied to the content, agent network used for testing, origin settings and origin infrastructure; any such changes will nullify a deficient test result.

(4) This SLA shall not apply if there is limited or no performance improvement due to causes originating from customer's infrastructure or a third party's infrastructure outside of Akamai's control, including the customer's DNS that provides the CNAME into the Akamai network.

Akamai reserves the right to remove DNS times from the speedup computation if deemed necessary.

B. Availability SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability:

Agents and Polling Frequency

(1) From at least six (6) geographically and network-diverse locations in major metropolitan areas,

(2) Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(3) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com). One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject). The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(4) The Akamaized test content must use a TTL of 2 hours or greater.

(5) The test content will be a file of approximately 10 KB in size.

(6) Polling will occur at approximately 6-minute intervals.

(7) Based on the http GET operations described in II-B above, the response times received from the two sources, (a) the Customer server (directly), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outage

An availability outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly).

IV. SLA Activation

A. Performance SLA Activation

To activate this SLA, Customer and Akamai will designate the applicable page as contemplated by II-A (1) above. Customer shall be responsible for configuring such page to enable testing by the measurement agents. Customer shall be responsible for configuring the Site Analyzer service for measurement of the designated page.

B. Availability SLA Activation

In order to activate the Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(B) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. SLA Escalation

In order to request a credit for a perceived service failure, Customer must, within five calendar days (120 hours) after the perceived failure, contact Akamai Customer Care in writing (which may be via email (ccare@akamai.com) or through the EdgeControl Management Center), specifying the time period in which the failure is believed to have occurred.

VI. Remedies

If the Service fails to meet the service levels in Section I or if an Outage, as defined in Section III, is identified, the Customer will receive (as its sole remedy) a credit equal to Customer's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

2.7 PROGRESSIVE MEDIA DOWNLOADS SERVICE LEVEL AGREEMENT (02/11/2014)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

(a) 100% Uptime: The Service will have no Outages (as defined herein below).

(b) Measurable Performance Enhancement: The Service will deliver content faster than the Customer's web site.

(c) Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject), or via an Akamai SiteShield region if the customer is using the SiteShield feature.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a SiteShield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a SiteShield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its

sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a SiteShield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate the Progressive Media Downloads Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and, for customers using SiteShield, Akamai Professional Services. Customers using Remote SiteShield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the SiteShield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

2.8 AKAMAI MEDIA CONTENT DELIVERY SERVICE LEVEL AGREEMENT (01/09/2014)

Applies to Adaptive Media Delivery, Download Delivery, Object Delivery, and Dynamic Site Delivery products.

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

(a) 100% Uptime: The Service will serve content 100% of the time.

(b) Measurable Performance Enhancement: The Service will deliver content measurably faster than the Customer's web site.

(c) Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Site Shield feature.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a Site Shield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate this Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and for customers using Akamai Professional Services. Customers using Remote Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

2.9 AKAMAI DYNAMIC SITE ACCELERATOR (DSA) SERVICE AND WEB APPLICATION ACCELERATOR (WAA) NORTH AMERICA DOMESTIC SERVICE LEVEL AGREEMENT (10/31/2007)

I. Service Levels

Akamai agrees to provide a level of service for customers with origin infrastructure hosted in North America demonstrating:

(1) Improvement Over Origin: The daily average page delivery time using the WAA or DSA service will be at least 20% faster than the daily average delivery time for the same page delivered from the customer's origin server (defined as the daily average page delivery time from origin being at least 1.2 times the daily average page delivery time using the WAA or DSA service).
(2) 100% Availability: The Service will serve content 100% of the time.

II. SLA Monitoring Methodology

A. Performance SLA Monitoring Methodology

The following methodology will be employed to measure the performance improvement provided by the Service:

(1) A single static unauthenticated page (including HTML and associated embedded content), selected and agreed upon by Akamai and Customer, will be tested throughout the term of the WAA or DSA Order Form for purposes of this SLA. Akamai will make a sample static page (including HTML and associated embedded content) available for use by customers.

(2) Delivery times will be tested using Akamai's Site Analyzer testing services to measure the daily average page delivery time computed from the complete North American set of available Site Analyzer measurement agents. The tests will be configured to take place once an hour from the measurement agents.

(3) This SLA assumes that there will be no material changes to the test content including, without limitation, metadata applied to the content, agent network used for testing, origin settings and origin infrastructure; any such changes will nullify a deficient test result.

(4) This SLA shall not apply if there is limited or no performance improvement due to causes originating from customer's infrastructure or a third party's infrastructure outside of Akamai's control, including the customer's DNS that provides the CNAME into the Akamai network. Akamai reserves the right to remove DNS times from the speedup computation if deemed necessary.

B. Availability SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability:

Agents and Polling Frequency

(1) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(2) The polling mechanism will perform two (2) simultaneous http GET operations:

(3) A test file will be placed on the customer's origin server (i.e., origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (i.e., http://origin.customer.com/testobject).

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (i.e., http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(4) The Akamaized test content must use a TTL of 2 hours or greater.

(5) The test content will be a file of approximately 10 KB in size.

(6) Polling will occur at approximately 6-minute intervals.

(7) Based on the http GET operations described in II-B above, the response times received from the two sources, (a) the Customer server (directly), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outage

An availability outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly).

IV. SLA Activation

A. Performance SLA Activation

To activate this SLA, Customer and Akamai will designate the applicable page as contemplated by II-A (1) above. Customer shall be responsible for configuring such page to enable testing by the measurement agents. Customer shall be responsible for configuring the Site Analyzer service for measurement of the designated page.

B. Availability SLA Activation

In order to activate the Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(B) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on

https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. SLA Escalation

In order to request a credit for a perceived service failure, Customer must, within five calendar days (120 hours) after the perceived failure, contact Akamai Customer Care in writing (which may be via email (ccare@akamai.com) or through the EdgeControl Management Center), specifying the time period in which the failure is believed to have occurred.

VI. Remedies

If the Service fails to meet the service levels in Section I or an Outage, as defined in Section III, is identified, the Customer will receive (as its sole remedy) a credit equal to Customer's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

2.10 KSD CLIENT REPUTATION OPTION SERVICE LEVEL AGREEMENT (02/08/2016)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

a. 100% Uptime: The Service will serve content 100% of the time.

b. Measurable Performance Enhancement: The Service will deliver content measurably faster than the Customer's web site.

c. Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

a. From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

b. The polling mechanism will perform two (2) simultaneous http GET operations:

c. A test file will be placed on the customer's origin server (i.e., origin.customer.com).

d. One GET operation will be performed to retrieve the file directly from the origin server (i.e., http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Site Shield feature.

   The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (i.e., http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

e. The Akamaized test content must use a TTL of 2 hours or greater.

f. The test content will be a file of approximately 10 KB in size.

g. Polling will occur at approximately 6-minute intervals.

h. Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a Site Shield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate the Kona Site Defender Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) using the Provision Test option found with the Configure menu under SLA Management within the LUNA portal. Detailed instructions regarding the setup of SLA Test Objects are found in the Support menu of LUNA; in addition, assistance is available from the Customer's Account Manager and, for customers using Site Shield, Akamai Professional Services. Customers using Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

VI. Governing SLA

This SLA shall apply solely to delivery configurations using the Kona Site Defender service (noted by "Site_Defender::Site_Defender" in the cpcode within the contract). For delivery configurations using other services, the SLA for those services shall apply and supersede this SLA.

The Web Application Firewall option does not have an SLA. Instead, the SLA for the base delivery solution to which WAF is added is the applicable SLA.

The Client Reputation option does not have an SLA. Instead, the SLA for the base delivery solution to which Client Reputation is added is the applicable SLA.

2.11 BOT MANAGER SERVICE LEVEL AGREEMENT (02/08/2016)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

a. Availability SLA: Akamai offers a service level ("Service Level") committing to 100% availability of the contracted security service.

b. Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

a. From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

b. The polling mechanism will perform two (2) simultaneous http GET operations:

c. A test file will be placed on the customer's origin server (i.e., origin.customer.com).

d. One GET operation will be performed to retrieve the file directly from the origin server (i.e., http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Site Shield feature.

   The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (i.e., http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

e. The Akamaized test content must use a TTL of 2 hours or greater.

f. The test content will be a file of approximately 10 KB in size.

g. Polling will occur at approximately 6-minute intervals.

h. Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

IV. SLA Activation

In order to activate the Kona Site Defender Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) using the Provision Test option found with the Configure menu under SLA Management within the LUNA portal. Detailed instructions regarding the setup of SLA Test Objects are found in the Support menu of LUNA; in addition, assistance is available from the Customer's Account Manager and, for customers using Site Shield, Akamai Professional Services. Customers using Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. Governing SLA

In order for Akamai to issue a credit in accordance with this SLA, Customer must have an account that is current with payments and in good standing with Akamai, and must be able to confirm that Customer has completed the Integration process for the applicable Service and, if applicable, all Competing Mitigation Hardware and Techniques have been disabled or removed during any mitigation services.

Credits shall only apply for Services provided pursuant to the Monthly Service Fee and/or Monthly Service Overage Fee, and will not apply to any other Service. The aggregate credits to be provided in any calendar month shall not exceed 25% of the Monthly Service Fee in respect of the affected Service(s).

This SLA shall apply solely to delivery configurations using the Kona Site Defender service (noted by "Bot Manager::Bot Manager" in the cpcode within the contract). For delivery configurations using other services, the SLA for those services shall apply and supersede this SLA.

3.0 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 Update 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).

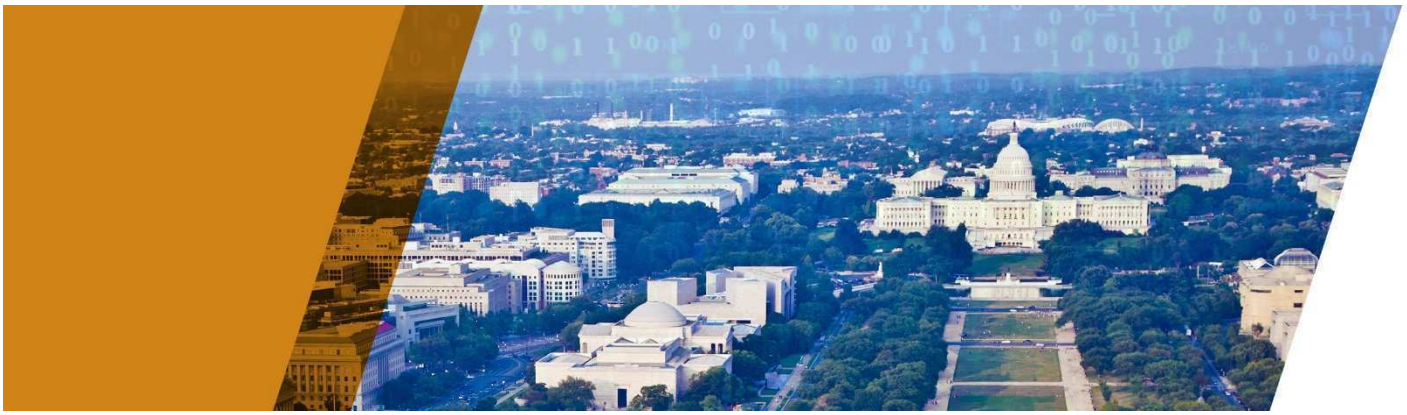

THE STATE OF UTAH • DIVISION OF PURCHASING

NASPO ValuePoint Master Agreement for Cloud Solutions

Utah Solicitation Number CH16012

Appendix 4
AWS Cloud Services Links to Publicly Available Information

PRESENTED BY

March 10, 2016

DLT Solutions / 2411 Dulles Corner Park, Suite 800 / Herndon VA 20171 / 703.709.7172 / www.dlt.com
TIN 54-1599882 / DUNS # 78-6468199 / CAGE Code 0S0H9 / GSA Schedule 70 GS-35F-4543G

AWS CLOUD SERVICES AND LINKS TO PUBLICLY AVAILABLE INFORMATION

Link to AWS Cloud Services Information	Service Description	Link to Documentation
Compute		
Amazon EC2	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.	http://aws.amazon.com/documentation/ec2/
Amazon EC2 Container Service	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.	http://aws.amazon.com/ecs/faqs/
AWS Lambda	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, in-app activity, website click, or output from a connected device.	http://aws.amazon.com/documentation/lambda/
Auto Scaling	Auto Scaling is a web service designed to launch or terminate Amazon EC2 instances automatically based on user-defined policies, schedules, and health checks.	http://aws.amazon.com/documentation/autoscaling/
Elastic Load Balancing	Elastic Load Balancing automatically distributes your incoming application traffic across multiple Amazon EC2 instances. It detects unhealthy instances and reroutes traffic to healthy instances until the unhealthy instances have been restored. Elastic Load Balancing	http://aws.amazon.com/documentation/elasticloadbalancing/

	automatically scales its request handling capacity in response to incoming traffic.	
Networking		
Amazon VPC	Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS.	http://aws.amazon.com/documentation/vpc/
Amazon Route 53	Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service.	http://aws.amazon.com/documentation/route53/
AWS Direct Connect	AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 GB or 10 GB Ethernet fiberoptic 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.	http://aws.amazon.com/documentation/directconnect/
Storage and Content Delivery		
Amazon S3	Amazon Simple Storage Service (Amazon S3) is storage for the Internet. You can use Amazon S3 to store and retrieve any amount of data, at any time, from anywhere on the web. You can accomplish these tasks using the simple and intuitive web interface of the AWS Management Console.	http://aws.amazon.com/documentation/s3/
Amazon Glacier	Amazon Glacier is a storage service optimized for infrequently used data, or "cold data." The service	http://aws.amazon.com/documentation/glacier/

	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 even 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 EBS	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 any 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.	http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html
Amazon CloudFront	Amazon CloudFront is a content delivery web service. It integrates with other AWS cloud services to give developers and businesses an easy way to distribute content to end users with low latency, high data transfer speeds, and no commitments.	http://aws.amazon.com/documentation/cloudfront/
AWS Import/Export	AWS Import/Export accelerates transferring large amounts of data between the cloud and portable storage devices that you mail to us. AWS transfers data directly onto and off of 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, we return your storage device. For	http://aws.amazon.com/documentation/importexport/

	large data sets, AWS Import/Export is significantly faster than Internet transfer and more cost effective than upgrading your connectivity.	
AWS Storage Gateway	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 AWS's storage infrastructure.	http://aws.amazon.com/documentation/storagegateway/
Databases		
Amazon RDS	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.	http://aws.amazon.com/documentation/rds/
Amazon DynamoDB	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 any amount of data, and serve any level 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.	http://aws.amazon.com/documentation/dynamodb/
Amazon Redshift	Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse solution that makes it simple and cost-effective to efficiently analyze all your data	http://aws.amazon.com/documentation/redshift/

	using your existing business intelligence tools. You can start small for just $0.25 per hour with no commitments or up-front costs and scale to a petabyte or more for $1,000 per terabyte per year, less than a tenth of most other data warehousing solutions.	
Amazon ElastiCache	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.	http://aws.amazon.com/documentation/elasticache/
Analytics		
Amazon EMR	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 such tasks as web indexing, data mining, log file analysis, machine learning, scientific simulation, and data warehousing.	http://aws.amazon.com/documentation/elasticmapreduce/
Amazon Kinesis	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.”	http://aws.amazon.com/documentation/kinesis/
AWS Data Pipeline	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	http://aws.amazon.com/documentation/data-pipeline/

	AWS Data Pipeline, you can 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.	
Amazon Mobile Analytics	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.	http://aws.amazon.com/documentation/mobileanalytics/
Administration & Security		
AWS Identity & Access Management	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 users can access.	http://aws.amazon.com/documentation/iam/
AWS Directory Service	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, and once this connection is established, all users can access AWS resources and applications with their existing corporate credentials.	http://aws.amazon.com/documentation/directory-service/

AWS Service Catalog	AWS Service Catalog is a service that allows administrators to create and manage approved catalogs of resources that end users can then access via 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.	http://aws.amazon.com/servicecatalog/
AWS Config	AWS Config is a fully managed service that provides you with an AWS 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 all configuration details, and determine how a resource was configured at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.	http://aws.amazon.com/documentation/config/
AWS CloudHSM	AWS CloudHSM provides secure cryptographic key storage to customers by making Hardware Security Modules (HSMs) available in the AWS cloud.	http://aws.amazon.com/documentation/cloudhsm/
AWS Key Management Service	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 including Amazon EBS, Amazon S3, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.	http://aws.amazon.com/documentation/kms/
AWS CloudTrail	With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made	http://aws.amazon.com/documentation/cloudtrail/

	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 CloudWatch	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 trends, and take automated action based on the state of your cloud environment.	http://aws.amazon.com/documentation/cloudwatch/
Deployment & Management		
AWS Management Console	Access and manage Amazon cloud services through a simple and intuitive web-based user interface. You can also use the AWS Console mobile app to quickly view resources on-the-go.	http://aws.amazon.com/console/faqs/
AWS Command Line Interface	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.	http://aws.amazon.com/documentation/cli/
APIs	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.	https://aws.amazon.com/tools/

AWS Elastic Beanstalk	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 health monitoring.	http://aws.amazon.com/documentation/elasticbeanstalk/
AWS CloudFormation	AWS CloudFormation gives 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 fashion. You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.	http://aws.amazon.com/documentation/cloudformation/
AWS CodeDeploy	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 to one Amazon EC2 instance or thousands.	http://aws.amazon.com/documentation/codedeploy/
AWS CodeCommit	AWS CodeCommit is a secure, highly scalable, managed source control 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	http://aws.amazon.com/codecommit/

	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.	
AWS CodePipeline	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 into staging, testing it, and releasing it to production. You can integrate third-party tools into any step of your release process or you can use AWS CodePipeline as an end-to-end solution.	http://aws.amazon.com/codepipeline/
AWS OpsWorks	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.	http://aws.amazon.com/documentation/opsworks/
Enterprise Applications		
Amazon WorkDocs	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 resort to emailing multiple versions of their files as attachments.	http://aws.amazon.com/documentation/workdocs/
Amazon Workspaces	Amazon WorkSpaces is a fully managed desktop computing service in the cloud. Amazon WorkSpaces allows customers to easily provision cloud-based desktops that allow end users to access the documents, applications, and resources they need with the device of their choice, including laptops, iPad, Kindle Fire, or Android tablets. With a few clicks in the	http://aws.amazon.com/documentation/workspaces/

	AWS Management Console, customers can provision a high-quality cloud desktop experience for any number of users at a cost that is highly competitive with traditional desktops and half the cost of most Virtual Desktop Infrastructure (VDI) solutions.	
Application Services		
Amazon AppStream	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.	http://aws.amazon.com/documentation/appstream/
Amazon CloudSearch	Amazon CloudSearch is a fully 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. 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.	http://aws.amazon.com/documentation/cloudsearch/
Amazon SWF	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.	http://aws.amazon.com/documentation/swf/

	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 and maintaining their state.	
Amazon SQS	Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles messages or workflows between other components in a system.	http://aws.amazon.com/documentation/sqs/
Amazon SES	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.	http://aws.amazon.com/documentation/ses/
Amazon SNS	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.	http://aws.amazon.com/documentation/sns/
Amazon Elastic Transcoder	Amazon Elastic Transcoder lets you convert media files that you have stored in Amazon S3 into media files in the formats required by consumer playback devices. For example, you can convert large, high-quality digital media files into formats that users can play back on mobile devices, tablets, web browsers, and connected televisions.	http://aws.amazon.com/documentation/elastictranscoder/
Amazon Cognito	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 also support unauthenticated guests.	http://aws.amazon.com/cognito/dev-resources/

Amazon FPS	Amazon Flexible Payments Service facilitates the digital transfer of money between any two entities, humans or computers.	https://payments.amazon.com/developer
AWS Support		
AWS Support	AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced and technical support engineers to help customers of all sizes and technical abilities successfully utilize the products and features provided by AWS.	https://aws.amazon.com/documentation/aws-support/
AWS Trusted Advisor	AWS Trusted Advisor acts like your customized cloud expert, and it helps you provision your resources by following best practices. AWS Trusted Advisor inspects your AWS environment and finds opportunities to save money, improve system performance and reliability, or help close security gaps. Since 2013, customers have viewed over 1.7 million best-practice recommendations and realized over $300 million in estimated cost reductions.	https://aws.amazon.com/premiumsupport/faqs/
AWS Marketplace		
AWS Marketplace	AWS Marketplace is an online store that helps customers find, buy, and immediately start using the software and services they need to build products and run their businesses.	https://aws.amazon.com/marketplace/help/

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

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.

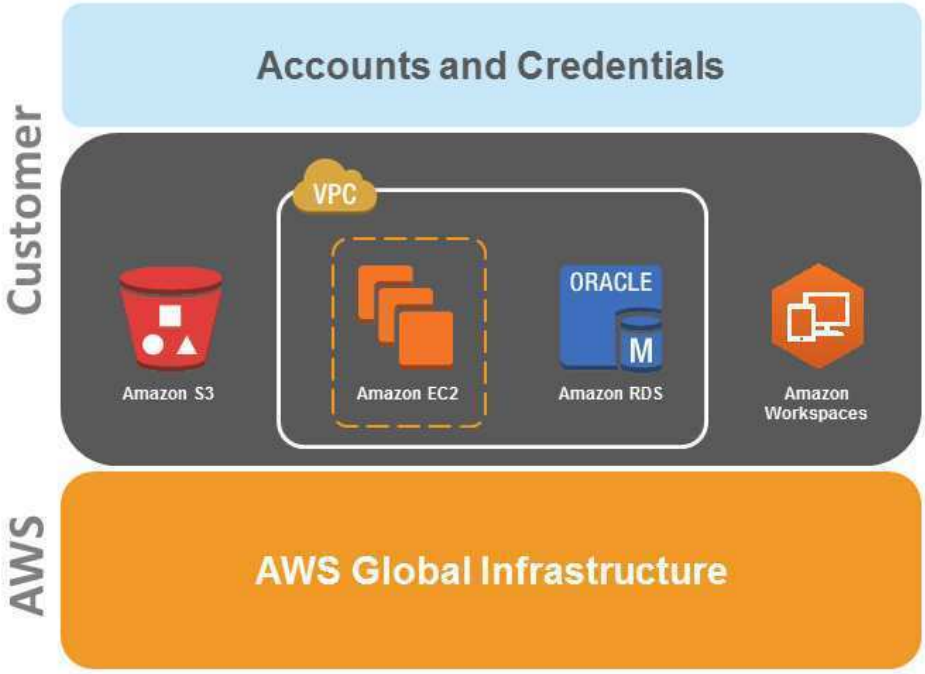

Figure 1: AWS Shared Security Responsibility Model

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/

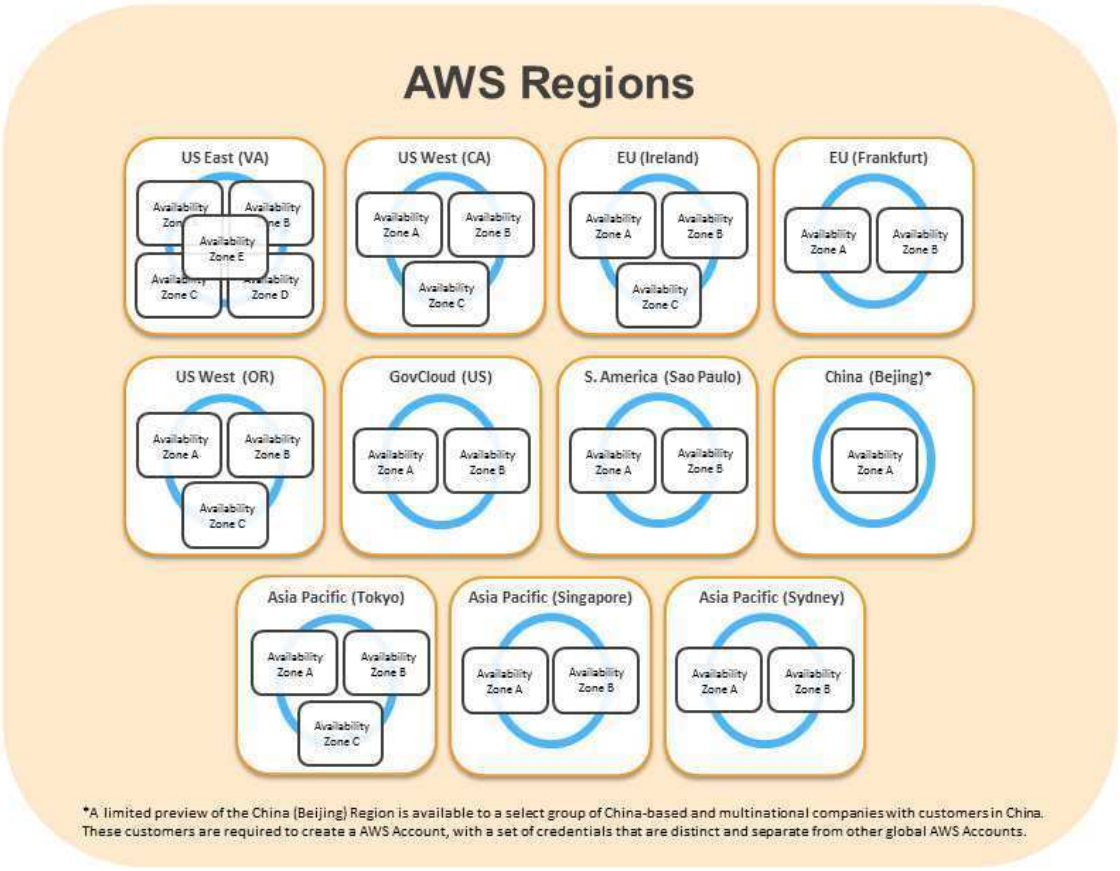

Figure 2: Regions and Availability Zones

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/AWSSecurityPenTestRequest
- **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.

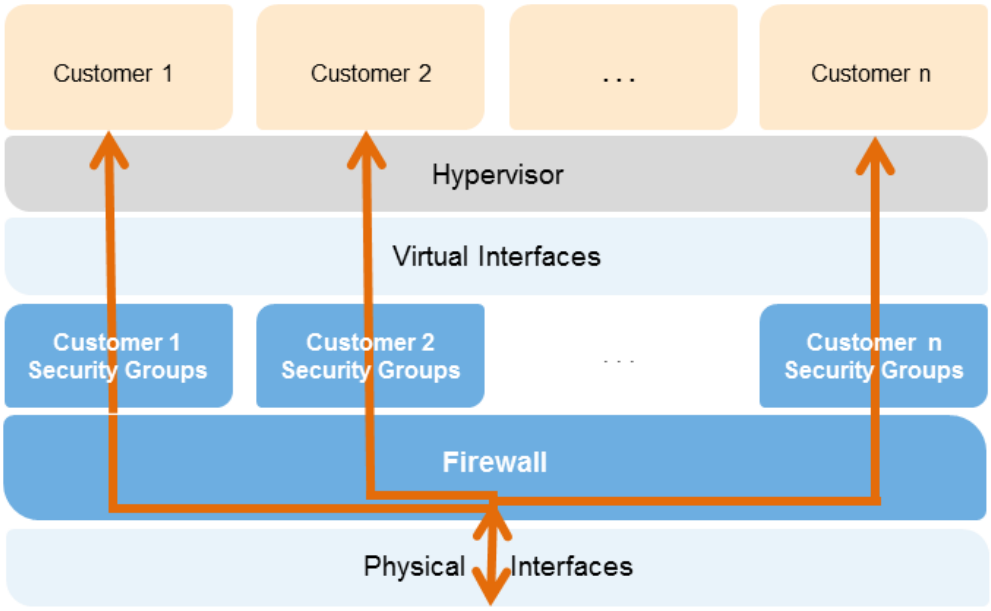

Figure 3: Amazon EC2 Multiple Layers of Security

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:

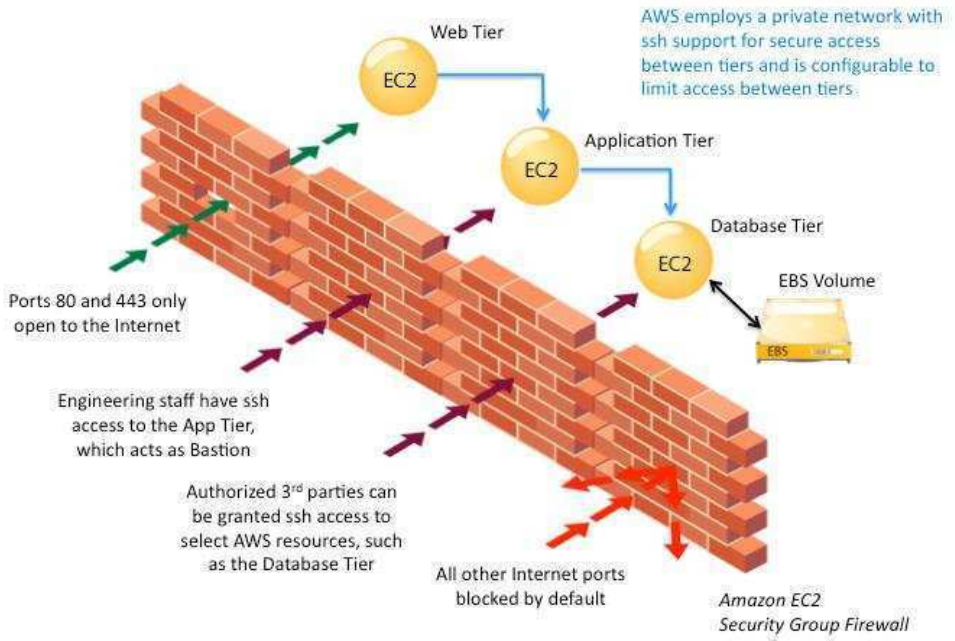

Figure 4: Amazon EC2 Security Group Firewall

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.

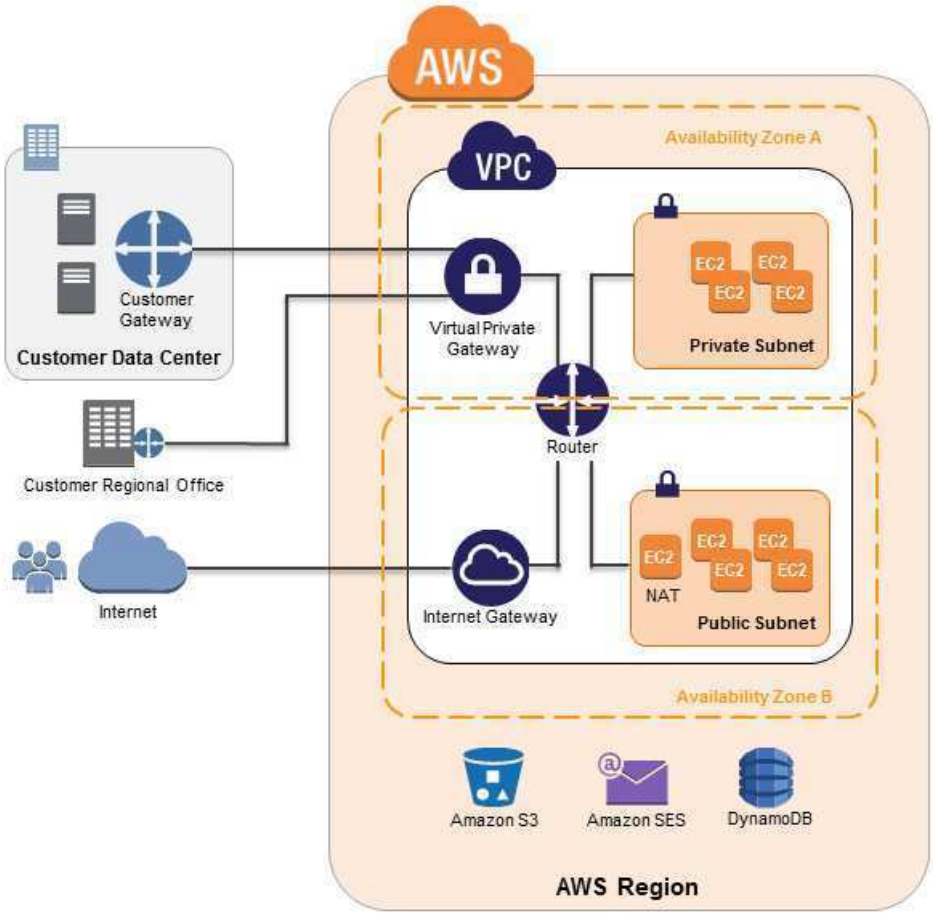

Figure 5: Amazon VPC Network Architecture

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.

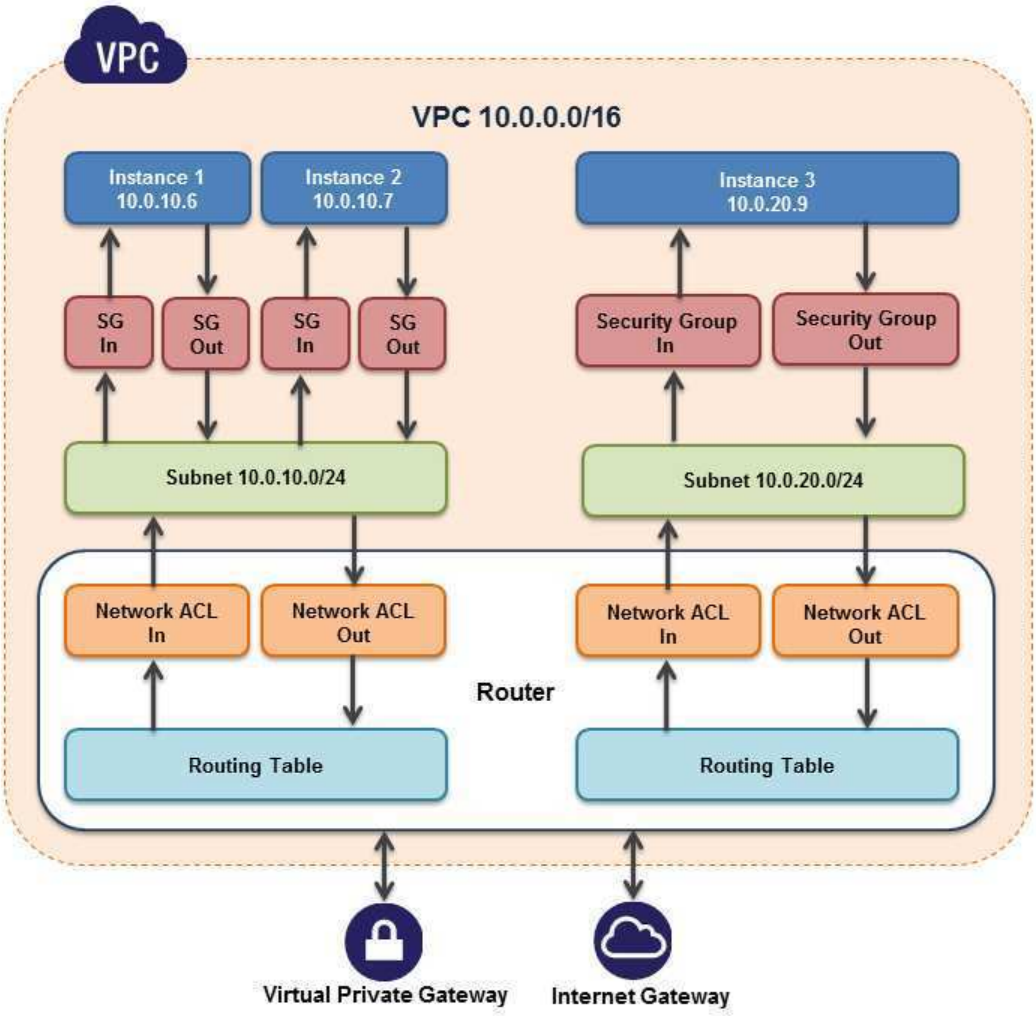

Figure 6: Flexible Network Topologies

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.

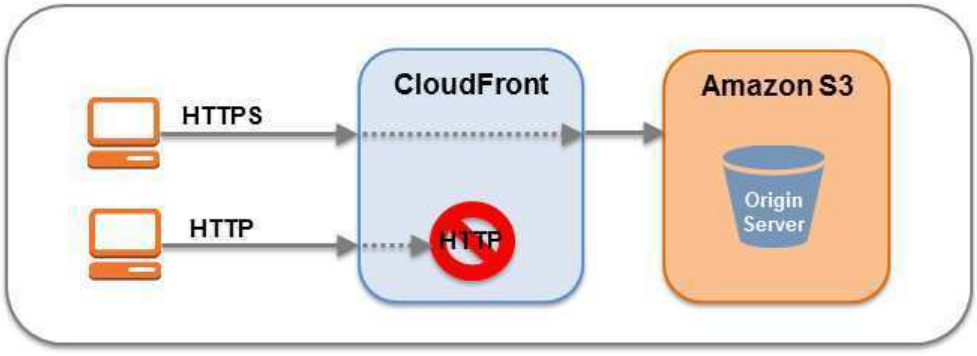

Figure 7: Amazon CloudFront Encrypted Transmission

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.

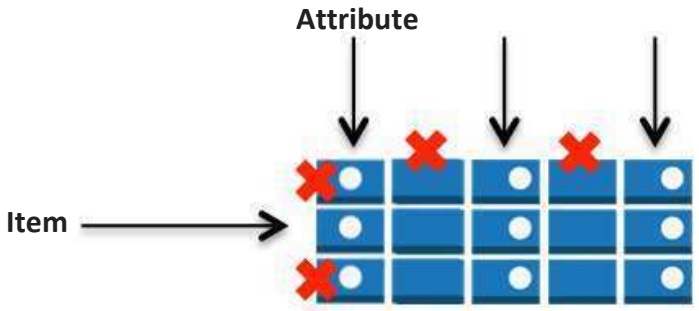

You can optionally use web identity federation to control access by application users who are authenticated by Login with Amazon, Facebook, or Google. Web identity federation removes the need for creating individual IAM users; instead, users can sign in to an identity provider and then obtain temporary 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 us 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 accidently 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 ensures 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 KHMAC), 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

Google Terms of Service

Last modified: April 14, 2014 (view archived versions)

Welcome to Google!

Thanks for using our products and services ("Services"). The Services are provided by Google Inc. ("Google"), located at 1600 Amphitheatre Parkway, Mountain View, CA 94043, United States.

By using our Services, you are agreeing to these terms. Please read them carefully.

Our Services are very diverse, so sometimes additional terms or product requirements (including age requirements) may apply. Additional terms will be available with the relevant Services, and those additional terms become part of your agreement with us if you use those Services.

Using our Services

You must follow any policies made available to you within the Services.

Don't misuse our Services. For example, don't interfere with our Services or try to access them using a method other than the interface and the instructions that we provide. You may use our Services only as permitted by law, including applicable export and re-export control laws and regulations. We may suspend or stop providing our Services to you if you do not comply with our terms or policies or if we are investigating suspected misconduct.

Using our Services does not give you ownership of any intellectual property rights in our Services or the content you access. You may not use content from our Services unless you obtain permission from its owner or are otherwise permitted by law. These terms do not grant you the right to use any branding or logos used in our Services. Don't remove, obscure, or alter any legal notices displayed in or along with our Services.

Our Services display some content that is not Google's. This content is the sole responsibility of the entity that makes it available. We may review content to determine whether it is illegal or violates our policies, and we may remove or refuse to display content that we reasonably believe violates our policies or the law. But that does not necessarily mean that we review content, so please don't assume that we do.

In connection with your use of the Services, we may send you service announcements, administrative messages, and other information. You may opt out of some of those communications.

Some of our Services are available on mobile devices. Do not use such Services in a way that distracts you and prevents you from obeying traffic or safety laws.

Your Google Account

You may need a Google Account in order to use some of our Services. You may create your own Google Account, or your Google Account may be assigned to you by an administrator, such as your employer or educational institution. If you are using a Google Account assigned to you by an administrator, different or additional terms may apply and your administrator may be able to access or disable your account.

To protect your Google Account, keep your password confidential. You are responsible for the activity that happens on or through your Google Account. Try not to reuse your Google Account password on third-party applications. If you learn of any unauthorized use of your password or Google Account, follow these instructions.

Privacy and Copyright Protection

Google's privacy policies explain how we treat your personal data and protect your privacy when you use our Services. By using our Services, you agree that Google can use such data in accordance with our privacy policies.

We respond to notices of alleged copyright infringement and terminate accounts of repeat infringers according to the process set out in the U.S. Digital Millennium Copyright Act.

We provide information to help copyright holders manage their intellectual property online. If you think somebody is violating your copyrights and want to notify us, you can find information about submitting notices and Google's policy about responding to notices in our Help Center.

Your Content in our Services

Some of our Services allow you to upload, submit, store, send or receive content. You retain ownership of any intellectual property rights that you hold in that content. In short, what belongs to you stays yours.

When you upload, submit, store, send or receive content to or through our Services, you give Google (and those we work with) a worldwide license to use, host, store, reproduce, modify, create derivative works (such as those resulting from translations, adaptations or other changes we make so that your content works better with our Services), communicate, publish, publicly perform, publicly display and distribute such content. The rights you grant in this license are for the limited purpose of operating, promoting, and improving our Services, and to develop new ones. This license continues even if you stop using our Services (for example, for a business listing you have added to Google Maps). Some Services may offer you ways to access and remove content that has been provided to that Service. Also, in some of our Services, there are terms or settings that narrow the scope of our use of the content submitted in those Services. Make sure you have the necessary rights to grant us this license for any content that you submit to our Services.

Our automated systems analyze your content (including emails) to provide you personally relevant product features, such as customized search results, tailored advertising, and spam and malware detection. This analysis occurs as the content is sent, received, and when it is stored.

If you have a Google Account, we may display your Profile name, Profile photo, and actions you take on Google or on third-party applications connected to your Google Account (such as +1's, reviews you write and comments you post) in our Services, including displaying in ads and other commercial contexts. We will respect the choices you make to limit sharing or visibility settings in your Google Account. For example, you can choose your settings so your name and photo do not appear in an ad.

You can find more information about how Google uses and stores content in the privacy policy or additional terms for particular Services. If you submit feedback or suggestions about our Services, we may use your feedback or suggestions without obligation to you.

About Software in our Services

When a Service requires or includes downloadable software, this software may update automatically on your device once a new version or feature is available. Some Services may let you adjust your automatic update settings.

Google gives you a personal, worldwide, royalty-free, non-assignable and non-exclusive license to use the software provided to you by Google as part of the Services. This license is for the sole purpose of enabling you to use and enjoy the benefit of the Services as provided by Google, in the manner permitted by these terms. You may not copy, modify, distribute, sell, or lease any part of our Services or included software, nor may you reverse engineer or attempt to extract the source code of that software, unless laws prohibit those restrictions or you have our written permission.

Open source software is important to us. Some software used in our Services may be offered under an open source license that we will make available to you. There may be provisions in the open source license that expressly override some of these terms.

Modifying and Terminating our Services

We are constantly changing and improving our Services. We may add or remove functionalities or features, and we may suspend or stop a Service altogether.

You can stop using our Services at any time, although we'll be sorry to see you go. Google may also stop providing Services to you, or add or create new limits to our Services at any time.

We believe that you own your data and preserving your access to such data is important. If we discontinue a Service, where reasonably possible, we will give you reasonable advance notice and a chance to get information out of that Service.

Our Warranties and Disclaimers

We provide our Services using a commercially reasonable level of skill and care and we hope that you will enjoy using them. But there are certain things that we don't promise about our Services.

OTHER THAN AS EXPRESSLY SET OUT IN THESE TERMS OR ADDITIONAL TERMS, NEITHER GOOGLE NOR ITS SUPPLIERS OR DISTRIBUTORS MAKE ANY SPECIFIC PROMISES ABOUT THE SERVICES. FOR EXAMPLE, WE DON'T MAKE ANY COMMITMENTS ABOUT THE CONTENT WITHIN THE SERVICES, THE SPECIFIC FUNCTIONS OF THE SERVICES, OR THEIR RELIABILITY, AVAILABILITY, OR ABILITY TO MEET YOUR NEEDS. WE PROVIDE THE SERVICES "AS IS".

SOME JURISDICTIONS PROVIDE FOR CERTAIN WARRANTIES, LIKE THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. TO THE EXTENT PERMITTED BY LAW, WE EXCLUDE ALL WARRANTIES.

Liability for our Services

WHEN PERMITTED BY LAW, GOOGLE, AND GOOGLE'S SUPPLIERS AND DISTRIBUTORS, WILL NOT BE RESPONSIBLE FOR LOST PROFITS, REVENUES, OR DATA, FINANCIAL LOSSES OR INDIRECT, SPECIAL, CONSEQUENTIAL, EXEMPLARY, OR PUNITIVE DAMAGES.

TO THE EXTENT PERMITTED BY LAW, THE TOTAL LIABILITY OF GOOGLE, AND ITS SUPPLIERS AND DISTRIBUTORS, FOR ANY CLAIMS UNDER THESE TERMS, INCLUDING FOR ANY IMPLIED WARRANTIES, IS LIMITED TO THE AMOUNT YOU PAID US TO USE THE SERVICES (OR, IF WE CHOOSE, TO SUPPLYING YOU THE SERVICES AGAIN).

IN ALL CASES, GOOGLE, AND ITS SUPPLIERS AND DISTRIBUTORS, WILL NOT BE LIABLE FOR ANY LOSS OR DAMAGE THAT IS NOT REASONABLY FORESEEABLE.

Business uses of our Services

If you are using our Services on behalf of a business, that business accepts these terms. It will hold harmless and indemnify Google and its affiliates, officers, agents, and employees from any claim, suit or action arising from or related to the use of the Services or violation of these terms, including any liability or expense arising from claims, losses, damages, suits, judgments, litigation costs and attorneys' fees.

About these Terms

We may modify these terms or any additional terms that apply to a Service to, for example, reflect changes to the law or changes to our Services. You should look at the terms regularly. We'll post notice of modifications to these terms on this page. We'll post notice of modified additional terms in the applicable Service. Changes will not apply retroactively and will become effective no sooner than fourteen days after they are posted. However, changes addressing new functions for a Service or changes made for legal reasons will be effective immediately. If you do not agree to the modified terms for a Service, you should discontinue your use of that Service.

If there is a conflict between these terms and the additional terms, the additional terms will control for that conflict.

These terms control the relationship between Google and you. They do not create any third party beneficiary rights.

If you do not comply with these terms, and we don't take action right away, this doesn't mean that we are giving up any rights that we may have (such as taking action in the future).

If it turns out that a particular term is not enforceable, this will not affect any other terms.

The laws of California, U.S.A., excluding California's conflict of laws rules, will apply to any disputes arising out of or relating to these terms or the Services. All claims arising out of or relating to these terms or the Services will be litigated exclusively in the federal or state courts of Santa Clara County, California, USA, and you and Google consent to personal jurisdiction in those courts.

For information about how to contact Google, please visit our contact page.

Data Processing Amendment

The Analytics Data Processing Amendment is meant for businesses established in the territory of a member state of the European Economic Area or Switzerland, or that, for other reasons, are subject to the territorial scope of the national implementations of Directive 95/46/EC.

Access to the Data Processing Amendment is available in the Account Settings in the Administration section. Users with edit permission for an Analytics account can review the Data Processing Amendment and accept it if they so choose. Once accepted, it is possible to review the contents of the accepted amendment.

Acceptance of the Data Processing Amendment is optional and does not affect the functionality of Analytics in any way.

- 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

Google Apps

Google Apps for Business (Online) Agreement

Go to the Additional Terms for services made available with the new accounts infrastructure

*The terms below are for monthly postpay billing. Please click here to see the terms for annual prepay customers.

This Google Apps for Business (Online) Agreement (the "Agreement") is entered into by and between Google Inc., a Delaware corporation, with offices at 1600 Amphitheatre Parkway, Mountain View, California 94043 ("Google") and the entity agreeing to these terms ("Customer"). This Agreement is effective as of the date you click the "I Accept" button below or, if applicable, the date the Agreement is countersigned (the "**Effective Date**"). If you are accepting on behalf of your employer or another entity, you represent and warrant that: (i) you have full legal authority to bind your employer, or the applicable entity, to these terms and conditions; (ii) you have read and understand this Agreement; and (iii) you agree, on behalf of the party that you represent, to this Agreement. If you don''t have the legal authority to bind your employer or the applicable entity, please do not click the "I Accept" button below (or, if applicable, do not sign this Agreement). This Agreement governs Customer''s access to and use of the Services.

1. **Services**

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 Google stores and processes its own information of a similar type. Google 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 Services Google may transfer store and process Customer Data in the United States or any other country in which Google or its agents maintain facilities. By using the Services Customer consents to this transfer, processing and storage of Customer Data.

1.2 **Modifications**

a. To the Services. Google may make commercially reasonable changes to the Services from time to time. If Google makes a material change to the Services Google will inform Customer, provided that Customer has subscribed with Google to be informed about such change.

b. **To URL Terms.** Google may make commercially reasonable changes to the URL Terms from time to time. If Google makes a material change to the URL Terms, Google will inform Customer by either sending an email to the Notification Email Address or alerting Customer 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 Google via the Help Center within thirty days after receiving notice of the change. If Customer notifies Google as required, then Customer will remain governed by the terms in effect immediately prior to the change until the end of the then-current Services Term for the affected Services. If the affected Services are renewed, they will be renewed under Google's then current URL Terms.

1.3 **Customer Domain Name Ownership.** Prior to providing the Services Google may verify that Customer owns or controls the Customer Domain Names. If Customer does not own or control the Customer Domain Names, then Google will have no obligation to provide Customer with the Services.

1.4 **Ads.** The default setting for the Services is one that does not allow Google to serve Ads. Customer may change this setting in the Admin Console which constitutes Customer's authorization for Google to serve Ads. If Customer enables the serving of Ads it may revert to the default setting at any time and Google will cease serving Ads.

1.5 **Google Apps Vault.** If Customer purchases Google Apps Vault, the following additional terms apply:

a. Retention. Google will have no obligation to retain any archived Customer Data beyond the retention period specified by Customer (other than for any legal holds). If Customer does not renew Google Apps Vault, Google will have no obligation to retain any archived Customer Data.

b. **Additional Purchases.** Unless Google allows otherwise, with each additional purchase of End User Accounts for the Services after Customer has purchased Google Apps Vault, Customer will receive access to, and will be invoiced for, Google Apps Vault for that same number of End User Accounts.

2. **Customer Obligations.**

2.1 **Compliance.** Customer will use the Services in accordance with the Acceptable Use Policy. Google may make new applications features or functionality for the Services available from time to time the use of which may be contingent upon Customer's agreement to additional terms. In addition, Google will make other Non-Google Apps Products (beyond the Services) available to Customer and its End Users in accordance with the Non-Google Apps Product Terms and the applicable product-specific Google 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 **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 Google may monitor emails sent to these aliases for Customer Domain Names to allow Google to identify Services abuse.

2.3 **Customer Administration of the Services.** 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 is 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 the Agreement. Customer agrees that Google's responsibilities do not extend to the internal management or administration of the Services for Customer and that Google is merely a data-processor.

2.4 **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 Google providing Customer with the ability to do so and (ii) Google to provide the Services.

2.5 **Unauthorized Use.** Customer will use commercially reasonable efforts to prevent unauthorized use of the Services' and to terminate any unauthorized use. Customer will promptly notify Google of any unauthorized use of, or access to, the Services of which it becomes aware.

2.6 **Restrictions on Use.** Unless Google 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 Services to a third party (unless expressly authorized in this Agreement); (b) attempt to reverse engineer the Services or any component; (c) attempt to create a substitute or similar service through use of, or access to, the Services; (d) use the Services for High Risk Activities; or (e) use the Services 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.

2.7 **Third Party Requests.** Customer is responsible for responding to Third Party Requests. Google 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 Google only if it cannot reasonably obtain such information.

3. **Billing and Payment.**

3.1 **Billing**Customer may elect one of the following billing options when placing its order for the Services.

a. Flexible Plan. If Customer selects this option, Customer will not be committed to purchase the Services for a pre-defined term, but will pay for the Services on a monthly basis. Google will bill Customer: (i) Fees based upon Customer's daily usage of the Services during the preceding month; and (ii) monthly in arrears for its use of the Services. Google will provide Customer with the monthly rate for the Services when Customer orders the Services, and will use this rate to calculate the Fees, on a prorated basis, for Customer's daily usage during that month. Any partial day of Services usages will be rounded up to a full day of Services usage for the purposes of calculating Fees. Customer may pay for the Services using the payment options listed below.

b. Annual Plan If Customer selects this option, Customer will be committed to purchasing the Services from Google for an annual term, and in exchange will receive a discount on the Services which will be reflected in Customer's monthly payment. Google will still bill Customer monthly in arrears for its use of the Services when Customer has an annual commitment for the Services with Google. Customer may pay for the Services using the payment options listed below.

3.2 **Payment.** All payments due are in U.S. dollars unless otherwise indicated on the Order Page or invoice.

a. Credit Card or Debit Card. Fees for orders where Customer is paying with a credit card, debit card or other non-invoice form of payment, are due at the end of the month during which Customer received the Services. For credit cards, or debit cards, as applicable: (i) Google will charge Customer for all applicable Fees when due and (ii) these Fees are considered delinquent thirty days after the end of the month during which Customer received the Services.

b. Invoices. Payments for invoices are due thirty days after the invoice date, unless otherwise specified on the Order Page, and are considered delinquent after such date.

c. Other Forms of Payment. Customer may change its payment method to those available within the Admin Console. Google may enable other forms of payment by making them available in the Admin Console. These other forms of payment may be subject to additional terms which Customer may have to accept prior using the additional forms of payment.

3.3 **Delinquent Payments** Delinquent payments may bear interest at the rate of one-and-one-half percent per month (or the highest rate permitted by law, if less) from the payment due date until paid in full. Customer will be responsible for all reasonable expenses (including attorneys' fees) incurred by Google in collecting such delinquent amounts, except where such delinquent amounts are due to Google's billing inaccuracies.

3.4 **Suspension for Non-Payment.**

a. Automatic Suspension. Customer will have thirty days to pay Google delinquent Fees. If Customer does not pay Google delinquent Fees within thirty days, Google will automatically suspend Customer's use of the Services. The duration of this suspension will be until Customer pays Google all outstanding Fees.

b. During Suspension. If Customer is on a monthly billing plan, and Customer is suspended for non-payment, Google will stop charging Customer monthly Fees during Customer's suspension for non-payment. If Customer has an annual commitment to Google for the Services, Google will continue to charge Customer monthly Fees during Customer's suspension for non-payment and Customer must pay all outstanding Fees in order to resume its use of the Services.

c. Termination After Suspension. If Customer remains suspended for non-payment for more than sixty days, Google may terminate Customer for breach pursuant to Section 11.

3.5 **Taxes.** Customer is responsible for any Taxes, and Customer will pay Google for the Services without any reduction for Taxes. If Google is obligated to collect or pay Taxes, the Taxes will be invoiced to Customer, unless Customer provides Google with a valid tax exemption certificate authorized by the appropriate taxing authority. If Customer is required by law to withhold any Taxes from its payments to Google, Customer must provide Google with an official tax receipt or other appropriate documentation to support such payments.

3.6 **Purchase Orders.** If Customer requires a purchase order number on its invoice, Customer will inform Google and Google will include such purchase order number on invoices following receipt. If Customer does not provide a purchase order number, Customer waives any purchase order requirement and (a) Google will invoice Customer without a purchase order number; and (b) Customer agrees to pay invoices without a purchase order number referenced. Any terms and conditions on a purchase order do not apply to this Agreement and are null and void.

4. **Technical Support Services.**

4.1 **By Customer.** Customer 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 Services. Customer will use commercially reasonable efforts to resolve support issues before escalating them to Google.

4.2 **By Google.** If Customer cannot resolve a support issue consistent with the above, then Customer may escalate the issue to Google in accordance with the TSS Guidelines. Google will provide TSS to Customer in accordance with the TSS Guidelines.

5. **Suspension**

5.1 **Of End User Accounts by Google.** If Google becomes aware of an End User's violation of the Agreement, then Google may specifically request that Customer Suspend the applicable End User Account. If Customer fails to comply with Google's request to Suspend an End User Account, then Google may do so. The duration of any Suspension by Google will be until the applicable End User has cured the breach which caused the Suspension.

5.2 **Emergency Security Issues.** Notwithstanding the foregoing, if there is an Emergency Security Issue, then Google 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 Google Suspends an End User Account for any reason without prior notice to Customer, at Customer's request, Google will provide Customer the reason for the Suspension as soon as is reasonably possible.

6. **Confidential Information.**

6.1 **Obligations.** Each party will: (a) protect the other party'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. Each party (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 this Agreement, while using reasonable care to protect it. Each party is responsible for any actions of its Affiliates' employees and agents in violation of this Section.

6.2 **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 another party.

6.3 **Required Disclosure.** Each party may disclose the other party's Confidential Information when required by law but only after it, if legally permissible: (a) uses commercially reasonable efforts to notify the other party; and (b) gives the other party the chance to challenge the disclosure.

7. **Intellectual Property Rights; Brand Features.**

7.1 **Intellectual Property Rights.** Except as expressly set forth herein, this Agreement does not grant either party any rights, implied or otherwise, to the other's content or any of the other's intellectual property. As between the parties, Customer owns all Intellectual Property Rights in Customer Data, and Google owns all Intellectual Property Rights in the Services.

7.2 **Display of Brand Features.** Google may display those Customer Brand Features authorized by Customer (such authorization is provided by Customer uploading its Brand Features into the Services) within designated areas of the Service Pages. Customer may specify the nature of this use using the Admin Console. Google may also display Google Brand Features on the Service Pages to indicate that the Services are provided by Google. Neither party may display or use the other party’s Brand Features beyond what is allowed in this Agreement without the other party’s prior written consent.

7.3 **Brand Features Limitation.** Any use of a party's Brand Features will inure to the benefit of the party holding Intellectual Property Rights in those Brand Features. A party may revoke the other party's right to use its Brand Features pursuant to this Agreement with written notice to the other and a reasonable period to stop the use.

8. **Publicity**.

Customer agrees that Google may include Customer's name or Brand Features in a list of Google customers, online or in promotional materials. Customer also agrees that Google may verbally reference Customer as a

customer of the Google products or services that are the subject of this Agreement. This section is subject to Section 7.3 (Brand Features Limitation).

9. **Representations, Warranties and Disclaimers.**

9.1 **Representations and Warranties.** Each party represents that it has full power and authority to enter into the Agreement. Each party warrants that it will comply with all laws and regulations applicable to its provision, or use, of the Services, as applicable (including applicable security breach notification law). Google warrants that it will provide the Services in accordance with the applicable SLA.

9.2 **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. GOOGLE MAKES NO REPRESENTATIONS ABOUT ANY CONTENT OR INFORMATION MADE ACCESSIBLE BY OR THROUGH THE SERVICES. CUSTOMER ACKNOWLEDGES THAT THE SERVICES ARE NOT A TELEPHONY SERVICE AND THAT THE SERVICES ARE NOT CAPABLE OF PLACING OR RECEIVING ANY CALLS, INCLUDING EMERGENCY SERVICES CALLS, OVER PUBLICLY SWITCHED TELEPHONE NETWORKS.

10. **Term**

10.1 **Agreement Term.** This Agreement will remain in effect for the Term.

10.2 **Services Term and Purchases During Services Term.** Google will provide the Services to Customer during the Services Term. Unless the parties agree otherwise in writing, End User Accounts purchased during any Services Term will have a prorated term ending on the last day of that Services Term.

10.3 **Renewal.**

a. With a Flexible Plan. With a flexible plan Customer is not committed to purchase the Services for a pre-defined term, but pays for the Services on a monthly basis. As a result, there is no renewal event for the flexible plan. Rather, Google will simply continuing billing Customer Fees based upon Customer's daily usage of the Services during the preceding month, and Customer can cancel their service at any time.

b. With an Annual Plan. At the end of each Services Term, the Services (and all End User Accounts previously purchased) will automatically renew for an additional monthly Services Term. In addition, after Customer's initial annual commitment has concluded, Customer's annual commitment will switch to the Flexible Plan. If Customer wants to renew the Annual Plan, then Customer must change the renewal settings in the Admin Console to reflect this change before their annual commitment has ended.

c. Generally. Customer may alter the number of End User Accounts to be renewed by communicating the appropriate number of accounts to be renewed to Google via the Admin Console. Customer will continue to pay Google the then-current Fees for each renewed End User Account unless Customer and Google mutually agree otherwise. If Google does not want the Services to renew, then it will provide Customer written notice to this effect at least fifteen days prior to the end of the then current Services Term. This notice of non renewal will be effective upon the conclusion of the then current Services Term.

10.4 **Requesting End User Accounts.** Customer may request End User Accounts by: (i) notifying its designated Google Account Manager; or (ii) ordering End User Accounts via the Admin Console.

10.5 **Revising Rates.** Google may revise its rates for the following Services Term by providing Customer written notice (which may be by email) at least thirty days prior to the start of the following Services Term.

11. **Termination.**

11.1 **Termination for Breach.** Either party may suspend performance or terminate this Agreement if: (i) the other party is in material breach of the Agreement and fails to cure that breach within thirty days after receipt of written notice; (ii) the other party ceases its business operations or becomes subject to insolvency proceedings and the proceedings are not dismissed within ninety days; or (iii) the other party is in material breach of this Agreement more than two times notwithstanding any cure of such breaches.

11.2 **Effects of Termination.** If this Agreement terminates, then: (i) the rights granted by one party to the other will cease immediately (except as set forth in this Section); (ii) Google will provide Customer access to, and the ability to export, the Customer Data for a commercially reasonable period of time at Google's then-current rates for the applicable Services; (iii) after a commercially reasonable period of time, Google will delete Customer Data by removing pointers to it on Google's active servers and overwriting it over time; and (iv) upon request each party will promptly use commercially reasonable efforts to return or destroy all other Confidential Information of the other party. If a Customer on an annual plan terminates the Agreement prior to the conclusion of its annual plan, Google will bill Customer, and Customer is responsible for paying Google, for the remaining unpaid amount of Customer's annual commitment.

12. **Indemnification.**

12.1 **By Customer.** Customer will indemnify, defend, and hold harmless Google from and against all liabilities, damages, and costs (including settlement costs and reasonable attorneys' fees) arising out of a third party claim: (i) regarding Customer Data or Customer Domain Names; (ii) that Customer Brand Features infringe or misappropriate any patent, copyright, trade secret or trademark of a third party; or (iii) regarding Customer's use of the Services in violation of the Acceptable Use Policy.

12.2 **By Google.** Google will indemnify, defend, and hold harmless Customer from and against all liabilities, damages, and costs (including settlement costs and reasonable attorneys' fees) arising out of a third party claim that Google's technology used to provide the Services or any Google Brand Feature infringe or misappropriate any patent, copyright, trade secret or trademark of such third party. Notwithstanding the foregoing, in no event shall Google have any obligations or liability under this Section arising from: (i) use of any Services or Google Brand Features in a modified form or in combination with materials not furnished by Google, and (ii) any content, information or data provided by Customer, End Users or other third parties.

12.3 **Possible Infringement.**

a. Repair, Replace, or Modify. If Google reasonably believes the Services infringe a third party's Intellectual Property Rights, then Google will: (a) obtain the right for Customer, at Google's expense, to continue using the Services; (b) provide a non-infringing functionally equivalent replacement; or (c) modify the Services so that they no longer infringe.

b. Suspension or Termination. If Google does not believe the foregoing options are commercially reasonable, then Google may suspend or terminate Customer's use of the impacted Services. If Google

terminates the impacted Services, then Google will provide a pro-rata refund of the unearned Fees actually paid by Customer applicable to the period following termination of such Services.

12.4 **General.** The party seeking indemnification will promptly notify the other party of the claim and cooperate with the other party in defending the claim. The indemnifying party has full control and authority over the defense, except that: (a) any settlement requiring the party seeking indemnification to admit liability or to pay any money will require that party's prior written consent, such consent not to be unreasonably withheld or delayed; and (b) the other party may join in the defense with its own counsel at its own expense. THE INDEMNITIES ABOVE ARE A PARTY'S ONLY REMEDY UNDER THIS AGREEMENT FOR VIOLATION BY THE OTHER PARTY OF A THIRD PARTY'S INTELLECTUAL PROPERTY RIGHTS.

13. **Limitation of Liability.**

13.1 **Limitation on Indirect Liability.** NEITHER PARTY WILL BE LIABLE UNDER THIS AGREEMENT FOR LOST REVENUES OR INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, EXEMPLARY, OR PUNITIVE DAMAGES, EVEN IF THE PARTY KNEW OR SHOULD HAVE KNOWN THAT SUCH DAMAGES WERE POSSIBLE AND EVEN IF DIRECT DAMAGES DO NOT SATISFY A REMEDY.

13.2 **Limitation on Amount of Liability.** NEITHER PARTY MAY BE HELD LIABLE UNDER THIS AGREEMENT FOR MORE THAN THE AMOUNT PAID BY CUSTOMER TO GOOGLE HEREUNDER DURING THE TWELVE MONTHS PRIOR TO THE EVENT GIVING RISE TO LIABILITY.

13.3 **Exceptions to Limitations.** These limitations of liability apply to the fullest extent permitted by applicable law but do not apply to breaches of confidentiality obligations, violations of a party's Intellectual Property Rights by the other party, or indemnification obligations.

14. **Miscellaneous.**

14.1 **Notices.** Unless specified otherwise herein, (a) all notices must be in writing and addressed to the attention of the other party's legal department and primary point of contact and (b) notice will be deemed given: (i) when verified by written receipt if sent by personal courier, overnight courier, or when received if sent by mail without verification of receipt; or (ii) when verified by automated receipt or electronic logs if sent by facsimile or email.

14.2 **Assignment.** Neither party may assign or transfer any part of this Agreement without the written consent of the other party, except to an Affiliate, but only if: (a) the assignee agrees in writing to be bound by the terms of this Agreement; and (b) the assigning party remains liable for obligations incurred under the Agreement prior to the assignment. Any other attempt to transfer or assign is void.

14.3 **Change of Control.** Upon a change of control (for example, through a stock purchase or sale, merger, or other form of corporate transaction): (a) the party experiencing the change of control will provide written notice to the other party within thirty days after the change of control; and (b) the other party may immediately terminate this Agreement any time between the change of control and thirty days after it receives the written notice in subsection (a).

14.4 **Force Majeure.** Neither party will be liable for inadequate performance to the extent caused by a condition (for example, natural disaster, act of war or terrorism, riot, labor condition, governmental action, and Internet disturbance) that was beyond the party's reasonable control.

14.5 **No Waiver.** Failure to enforce any provision of this Agreement will not constitute a waiver.

14.6 **Severability.** If any provision of this Agreement is found unenforceable, the balance of the Agreement will remain in full force and effect.

14.7 **No Agency.** The parties are independent contractors, and this Agreement does not create an agency, partnership or joint venture.

14.8 **No Third-Party Beneficiaries.** There are no third-party beneficiaries to this Agreement.

14.9 **Equitable Relief.** Nothing in this Agreement will limit either party's ability to seek equitable relief.

14.10 **Governing Law.** This Agreement is governed by California law, excluding that state's choice of law rules. FOR ANY DISPUTE ARISING OUT OF OR RELATING TO THIS AGREEMENT, THE PARTIES CONSENT TO PERSONAL JURISDICTION IN, AND THE EXCLUSIVE VENUE OF, THE COURTS IN SANTA CLARA COUNTY, CALIFORNIA.

14.11 **Amendments.** Any amendment must be in writing and expressly state that it is amending this Agreement.

14.12 **Survival.** The following sections will survive expiration or termination of this Agreement: Section 3, 6, 7.1, 11.2, 12, 13, 14, and 15.

14.13 **Entire Agreement.** This Agreement, and all documents referenced herein, is the parties' entire agreement relating to its subject and supersedes any prior or contemporaneous agreements on that subject. The terms located at a URL and referenced in this Agreement are hereby incorporated by this reference.

14.14 **Interpretation of Conflicting Terms.** If there is a conflict between the documents that make up this Agreement, the documents will control in the following order: the Order Page, the Agreement, and the terms located at any URL. If Customer signs a physical agreement with Google to receive the Services, the physical agreement will override this online Agreement.

14.15 **Counterparts.** The parties may enter into this Agreement in counterparts, including facsimile, PDF or other electronic copies, which taken together will constitute one instrument.

15. **Definitions.**

"Acceptable Use Policy" means the acceptable use policy for the Services available at https://www.google.com/a/help/intl/en/admins/use_policy.html or such other URL as Google may provide.

"Account Manager" means the Google business person working with Customer regarding Customer's purchase of the Services.

"Admin Account(s)" means the administrative account(s) provided to Customer by Google for the purpose of administering the Services. The use of the Admin Account(s) requires a password, which Google will provide to Customer.

"Admin Console" means the online tool provided by Google to Customer for use in reporting and certain other administration functions.

"Administrators" mean the Customer-designated technical personnel who administer the Services to End Users on Customer's behalf.

"Ads" means online advertisements displayed by Google to End Users.

"Affiliate" means any entity that directly or indirectly controls, is controlled by, or is under common control with a party.

"Brand Features" means the trade names, trademarks, service marks, logos, domain names, and other distinctive brand features of each party, respectively, as secured by such party from time to time.

"Confidential Information" means information disclosed by a party to the other party under this Agreement 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 Services by Customer or End Users.

"Customer Domain Names" mean the domain names owned or controlled by Customer, which will be used in connection with the Services and specified in the Order Page.

"Emergency Security Issue" means either: (a) Customer's use of the Services in violation of the Acceptable Use Policy, which could disrupt: (i) the Services; (ii) other customer's use of the Services; or (iii) the Google network or servers used to provide the Services; or (b) unauthorized third party access to the Services.

"End Users" means the individuals Customer permits to use the Services.

"End User Account" means a Google-hosted account established by Customer through the Services for an End User.

"Export Control Laws" means all applicable export and reexport control laws and regulations, including the Export Administration Regulations (“EAR”) maintained by the U.S. Department of Commerce, 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.

"Fees" means the amounts invoiced to Customer by Google for the Services as described in an Order Page.

"Help Center" means the Google help center accessible at https://www.google.com/support/, or other such URL as Google may provide.

"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 Services 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.

"Initial Services Term" means the term for the applicable Services beginning on the Service Commencement Date and continuing for the duration set forth on the Order Page.

"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 Google products which are not part of the Services, 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: https://www.google.com/support/a/bin/answer.py?hl=en&answer=181865, or such other URL as Google may provide.

"Non-Google Apps Product Terms" means the terms found at the following URL: https://www.google.com/apps/intl/en/terms/additional_services.html, or such other URL as Google may provide.

"Notification Email Address" means the email address designated by Customer to receive email notifications from Google. Customer may change this email address through the Admin Console.

"Order Page" means the online order page Customer completes in signing up for the Services or attached to this Agreement, and which contains: (i) the Services being ordered; (ii) Fees; (iii) number of, and Initial Services Term for, End User Accounts; (iv) the applicable form of payment; and (v) Customer Domain Names.

"Service Commencement Date" is the date upon which Google makes the Services available to Customer, and will be within one week of Google's receipt of the completed Order Page, unless otherwise agreed by the parties.

"Service Pages" mean the web pages displaying the Services to End Users.

"Services" means the applicable Google Apps Core Services (e.g. Google Apps Premier Edition or Google Apps for Business and Google Apps Vault) provided by Google and used by Customer under this Agreement. The Services are as described here: https://www.google.com/apps/intl/en/terms/user_features.html, or such other URL as Google may provide.

"Services Term" means the Initial Services Term and all renewal terms for the applicable Services.

"SLA" means the Service Level Agreement located here: https://www.google.com/a/help/intl/en/admins/sla.html, or such other URL as Google may provide.

"Suspend" means the immediate disabling of access to the Services, or components of the Services, as applicable, to prevent further use of the Services.

"Taxes" means any duties, customs fees, or taxes (other than Google's income tax) associated with the sale of the Services, including any related penalties or interest.

"Term" means the term of the Agreement, which will begin on the Effective Date and continue until the earlier of (i) the end of the last Services Term or (ii) the Agreement is terminated as set forth herein.

"Third Party Request" means a request from a third party for records relating to an End User's use of the Services. 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 services provided by Google to the Administrators during the Term pursuant to the TSS Guidelines.

"TSS Guidelines" means Google's technical support services guidelines then in effect for the Services. TSS Guidelines are at the following URL: https://www.google.com/a/help/intl/en/admins/tssg.html or such other URL as Google may provide.

"URL Terms" means the "Acceptable Use Policy," the "SLA," and the "TSS Guidelines."

Version: March 28, 2012

Log Delivery Service User Guide

Notice

As the global leader in Content Delivery Network (CDN) services, Akamai makes the Internet fast, reliable and secure. The company's advanced web performance, cloud security and media delivery solutions are revolutionizing how businesses optimize the consumer, enterprise or entertainment experience to any device, anywhere. To learn how Akamai solutions, and its team of Internet experts, are helping businesses move faster forward, please visit www.akamai.com and follow @Akamai on Twitter.

Table of Contents

Introduction

Akamai understands that you rely on us to deliver your content to give your end users a better overall online experience. We also understand that delivering your content means that we have the server logs that previously existed on your servers – and that this is often mission-critical information to your organization.

The Akamai Log Delivery Service provides you with the server logs from the various services that you are using, and is available for the following services:

- HTTP Content Delivery, HTTP Downloads, HTTP Downloads Client Side Delivery, Site Delivery, Web Application Accelerator, and Site Accelerator
- Streaming (QuickTime Streaming, Real Media Streaming, Windows Media Streaming, and Streaming for Flash)
- NetStorage (FTP)
- Fast File Upload module for Web Application Accelerator and Site Accelerator. Fast File Upload is for HTTP-based uploads using PUT/POST methods as opposed to FTP uploads.
- IPA/SXL, IP Application Accelerator and Session Accelerator
- Fast DNS
- Traffic Management
- AnswerX Managed

Intended Audience

This document is intended for those who receive logs from Akamai and are responsible for processing them.

How does Log Delivery Service work?

Akamai's infrastructure is constantly gathering the log entries from the thousands of edge servers around the world. Log Delivery Service creates a copy of these logs, separates your logs from other customer logs, and then delivers your logs based on a predetermined schedule. Most of the log files will be delivered within a 24-hour period. Due to the distributed nature of the network, some number of log lines can be delayed and be part of a later delivery.

You must be a customer of Log Delivery Service to configure the service and begin receiving logs from that point forward. Logs are not available retroactively, and logs for other customers who do not subscribe to Log Delivery Service are not retrievable.

Configuring Log Delivery Service

Log Delivery Service can be configured through Luna Control Center. Access to this configuration tool is available by default to users who have Technical or Admin role types and who are associated with contracts that include Log Delivery Service.

You can create one LDS configuration for each service associated with a CP code or DNS for Fast DNS or domain for Traffic Management. Services include HTTP Content Delivery, HTTP Downloads, HTTP Downloads for Client Side Delivery, Site Delivery, Streaming for Flash, QuickTime Streaming, Real Media Streaming, Windows Media Streaming, NetStorage (FTP) and the Fast File Upload module for Web Application Accelerator, Site Accelerator, Fast DNS and Traffic Management, and AnswerX Managed.

Once a service is initially configured for LDS, the tool shows whether the configuration is active or suspended, and allows you to view or edit it.

Access Log Delivery Service Configuration

To access Log Delivery Service configuration:

1. Log in to Luna Control Center.
2. Click **Configure** to expand the menu.
3. Click **Log Delivery**.

The Log Delivery Service page appears, showing the objects and service(s) associated with your login. This page indicates whether there is any existing log delivery configuration, or whether the current log delivery configurations are enabled or suspended. You can create new configurations and view or modify current configurations from this page, as well as view as summary of the configurations that are active. You can configure how many objects are included in the view at the bottom of the screen.

If you have configured Log Delivery Service for FTP delivery, you can test the FTP connection for a log delivery configuration under Action > Test FTP Connection.

Create a new Log Delivery Service configuration

When creating a new Log Delivery Service configuration, you can create a configuration by defining all options You can also copy an existing configuration by using the Copy button to configure one object, or the Bulk Copy button to configure multiple objects based on the one chosen configuration.

To create a new log delivery configuration for a service:

1. Find the object ID/service pair to be configured on the Log Delivery Service page.
2. Select the Begin Log Delivery action. Only services that lack a current log delivery configuration will have this action. The Begin Log Delivery page appears.
3. Define the options for your log delivery configuration, delivery method, and contact information. You will click Next three times before you have completed all options. See *Appendix B: Overview of Log Delivery Options*on page 10 for a complete discussion of available options.
4. Click Finish to initiation this Log Delivery Service configuration.

Because it is difficult to time the LDS provisioning process, a best practice is to set the start date to two days before you need the service enabled.

Editing an existing Log Delivery Service configuration

To edit an existing Log Delivery Service configuration:

1. Find the Object ID/service pair whose configuration you want to modify on the Log Delivery Service page.
2. Select the Change Settings action
3. Edit the configuration options as desired.
4. Click Finish.

Because it is difficult to time the LDS provisioning process, a best practice is to introduce changes two days before you need the service to be modified.

Enabling or suspending a Log Delivery Service configuration

To enable or suspend a Log Delivery Service configuration:

1. Find the Object ID/service pair whose configuration you want to suspend or resume on the Log Delivery Service page.
2. The service will have a Resume or a Suspend action, as appropriate, that allows you to perform that function on this configuration. Click the associated Resume or Suspend action.

Because it is difficult to time the LDS provisioning process, a best practice is to suspend or resume delivery two days before you need the service enabled.

Requesting redelivery of a log

The Luna Control Center enables you to request redelivery of logs. You can select the Object IDs for which you wish to redeliver the logs, and the time period of the desired data. Note that log redeliveries are sent at lower priority than normal log deliveries, and you may not request a redelivery until the original log has been completed. Logs are available for redelivery for the log retention period indicated next to each configuration.

Redelivery Intervals

The intervals at which redelivery occurs depends on the type of log aggregation associated with your object. The Types of Aggregation section discusses these options as part of configuring Log Delivery.

Aggregated by log arrival time: For objects configured for aggregation by log arrival time, you can select a date, start time, and end time for the redelivery. Log data will be redelivered for all time intervals in the selected time period, based on the delivery frequency that is configured for that CP code. For example, if you select a 24-hour time period for redelivery and your object is configured for a 24-hour delivery interval, you will receive one redelivery encompassing 24 hours of log data. If your object is configured for a one-hour delivery internal, you will receive 24 redeliveries, each encompassing one hour of log data.

Aggregated by calendar day: For objects configured for aggregation by calendar day (that is, hit time), you can select the date for the redelivery. Calendar day redeliveries can only be requested for entire days, so the start and end time fields are not displayed. All log data that has been collected for that date will be redelivered. Since some hits may have arrived after the original delivery, it is possible that the redelivery may include more log lines than the initial delivery.

Note that if you select multiple objects, the redeliveries for each object will occur according to that object's log configuration as described here.

To request a redelivery:

1. Select the **Redelivery** tab where you will see the list of already requested redeliveries. Select **New Redelivery**.
2. Select the product and configurations to be redelivered.
3. Identify the dates, start time, and end time in GMT. For calendar day deliveries, select the date for redelivery. Please note that all redeliveries will be based on the delivery frequency you're configured for, but will be optimized to include the interval between the selected start time and end time.
4. Click Redeliver.

If you are requesting a redelivery because the original delivery failed, first be sure that the cause of the original failure is corrected, and verify that the redelivery was successful once it is complete.

Overview of Log Delivery Options

This section describes the options you must define when configuring Log Delivery Service.

Log Delivery Options

Message size

The maximum message size applies to both email and FTP. For all encoding options other than MIME encoded email deliveries, Log Delivery Service attempts to send logs in parts no greater than the maximum size when compressed. However, logs are broken into parts before compression, so the estimation of message size is based on the uncompressed file size.

Estimation of compressed size currently assumes a 10-to-1 compression ratio. The same ratio is assumed for all customers and is based on historic log compressibility statistics. Because actual compressibility can vary from customer to customer, delivery parts can be larger or smaller than the configured size. For MIME encoded email deliveries, the logs will be split after compression.

Encoding

Log Delivery Service supports three mail formats: gzipped and uuencoded, MIME with base-64 encoding, and GPG Encrypted. Note that GPG Encrypted requires a key that you must upload. Otherwise, the main difference is support by mail clients. Most customers using email delivery are currently configured for gzipped and uuencoded logs.

There are four types of encoding, two of which are only available in email:

Type	Available via	Description
gz	FTP or email	GNU gzip
gpg	FTP or email	Gzipped and GPG encrypted
gzu	Email only	Gzipped and uuencoded
gzm	Email only	Mime (base64)

Status

Log delivery configurations, once created, will have a status of either "Suspended" or "Active" status. Services that have no log delivery configured will show a Begin Log Delivery action, through which you can set up a configuration. Log delivery configurations that are currently active provide you the option to suspend them.

Start date

The start date affects both the beginning of log collection and the delivery of logs, but is unrelated to the timestamps of the logs collected. Log collection will begin within no more than 6-8 hours of the start date, and the first delivery will be scheduled within no more than two days.

Because it is difficult to time the LDS provisioning process, a best practice is to set the start date to two days before you need the service enabled.

NOTE: LDS relies on data from the Domain Validation Tool. This works the same way in LDS as it does for Akamai's billing and reporting systems. Logs will only be collected if the origin server domain in the logs matches a domain on the valid domains list. The domain list is independent of any other LDS provisioning, and will go live on the LDS system automatically. It may take up to 3 hours for a domain to become live for Log Delivery Service configurations from the time it is entered in the Valid Domains area of Luna Control Center.

End date

The end date can be set to either an actual end date, or "indefinite" by leaving the field blank. If it is set to a specific end date, you might get a few empty deliveries before it ends. Like the start date, the end date is unrelated to the timestamps of the individual log lines, and refers only to log collection and delivery.

Administrative contact information

To complete the configuration, you must provide a name, phone number, and email address. This information will be used both for announcements about changes to the product as well as automated delivery failure notices.

Do not use the same email address to which emailed logs are sent. We strongly recommend that you also avoid using an address in the same domain or hosted on the same mail server as the delivery address; should that domain or server become unavailable, you will then still receive announcements and failure notices.

You can enter multiple email addresses, delimited by commas. Please make certain that all contacts that might need access to delivery failure notices are included in the email. You might find it simpler to give the address of a mailing list to which multiple users are subscribed, rather than updating the LDS configuration each time the email list needs to be changed.

It is essential that the admin email address is checked on a regular basis. Akamai only guarantees availability of logs for seven calendar days after they were originally delivered. If delivery failures are responded to promptly, five business days should be adequate to diagnose and correct the causes of delivery failures, as well as to contact Akamai to arrange a redelivery. For more on delivery failures and requesting redeliveries, see Troubleshooting.

Type of Delivery

Log Delivery Service offers delivery via email or FTP.

Email

If you are using email for your log delivery, be sure that your email server is able to handle the volume of mail that LDS will produce. We do not recommend using a standard corporate email server for handling logs from LDS.

Advantages of email:

- Widely available.
- Incorporates a robust delivery mechanism that queues messages and retries when servers become temporarily unavailable.
- Eliminates the need to configure an FTP server and manage usernames and passwords.

Disadvantages of email:

- Some email servers might not be able to handle high volumes.
- Cannot be used in conjunction with NetStorage.
- If log delivery emails are rejected by your email servers for some reason, Akamai has no way of knowing this or notifying you about the failed delivery.

FTP

Advantages of FTP:

- Can handle higher log volumes than the average mail server.

- Can be used in conjunction with NetStorage .

Disadvantages of FTP:

- Less widely available than email.
- More vulnerable to delivery failures, due to the lack of a built-in redelivery mechanism.
- Requires more initial configuration than email.

Type of Aggregation

Logs can be aggregated by arrival time or by calendar day.

NOTE:

1. All time periods referenced in the Aggregation discussion are GMT.

Aggregate by log arrival time

Aggregating logs by arrival time allows you to receive log data in the order in which it arrives from the edge servers. This is the fastest way to get log data, but it is more difficult to process. For example, a delivery that covers a 24-hour period will usually contain some data from the previous several days.

Logs can be delivered at intervals of 1, 2, 3, 4, 6, 8, 12 and 24 hours, with the default being every 24 hours. Less frequent deliveries are often more convenient to work with. More frequent deliveries will help ensure that you get your logs as soon as possible. Large logs, however, may require more frequent deliveries.

Aggregate by calendar day (hit time)

Aggregating logs by calendar day (hit time) allows you to receive most of the hits for each GMT calendar day (a 24-hour period) in a single bulk delivery. You may set the threshold, that is, the percentage of log completeness (described below), upon which you would like the log data sent. Please note that the higher you set the threshold, the longer you will have to wait for your log data. Residual data will contain any hits that were not processed in time to make it into the bulk delivery. This data is sent at regular intervals after each day for a period of four days if you check the "Deliver residual data" checkbox.

Completeness Threshold

Due to the distributed nature of the edge network, it is impossible to determine in advance that all hits that occurred on a particular day have been processed and are ready for delivery by LDS. For calendar day deliveries, you can set a completion threshold that is used by the LDS system to determine when the data is complete enough to send the bulk delivery. The completion threshold determines the percentage of expected logs that must be processed before delivery. Setting the threshold value to a high percentage will ensure that your delivery has data that is complete as possible, but may cause deliveries to be delayed to

meet the threshold. Setting the threshold value to a lower percentage will provide for more timely deliveries, but more of the data may be deferred to a leftover delivery.

Residual Delivery Frequency

You can also choose to receive leftover deliveries containing log lines that were not included in either the bulk delivery or a previous leftover delivery. Most of the log data will be included in the bulk delivery, so the leftover deliveries will often contain no log lines. However, the delivery will be sent, even if the file does not contain any log lines. These residual deliveries are sent once every 24 hours.

Log Identifier String

The log identifier string is a unique label that is added to the file names of log files. If you have multiple LDS configurations, the log file identifier string can make it much easier to distinguish them visually. Log identifier strings are limited to alphanumeric characters (numbers and letters) only, and may be no longer than 80 characters. They are case-sensitive, meaning that the case of names entered in the configuration tool will be preserved in log file names.

See *Appendix B: Filenames and Email Subject Lines* on page 25 for a complete description of log file names and email subject lines.

Email Address

You may only have logs delivered to a single email address. If you need logs to go to multiple email addresses, consider sending logs to a mailing list managed on your local mail server.

The email address is limited to 80 characters. You may see a phone number in previously configured email addresses; this information is pulled from the person's portal login information, if it's available.

FTP Login

You will need the FTP server name, the login, the password, and the directory path (if the account used for LDS does not automatically default to this directory).

Before accepting a new configuration, the portal will run a series of tests against your FTP server. These tests consist of trying to:

1. connect to the customer FTP server (uses active mode)
2. login with the userid and password given
3. cd to the log directory specified in the configuration
4. put a file on the server
5. delete a file (del)
6. do a directory listing on the log directory (dir)

7. overwrite a file (put the same file twice)
8. rename a file (mv)

Note that these tests will come from a machine associated with Luna Control Center, and not an actual LDS delivery machine.

Administrative contact information

To complete the configuration, you must provide a name, phone number, and email address. This information will be used both for announcements about changes to the product as well as automated delivery failure notices.

Do not use the same email address to which emailed logs are sent. We strongly recommend that you also avoid using an address in the same domain or hosted on the same mail server as the delivery address; should that domain or server become unavailable, you will then still receive announcements and failure notices.

You can enter multiple email addresses, delimited by commas. Please make certain that all contacts that might need access to delivery failure notices are included in the email. You might find it simpler to give the address of a mailing list to which multiple users are subscribed, rather than updating the LDS configuration each time the email list needs to be changed.

It is essential that the admin email address is checked on a regular basis. Akamai only guarantees availability of logs for seven calendar days after they were originally delivered. If delivery failures are responded to promptly, five business days should be adequate to diagnose and correct the causes of delivery failures, as well as to contact Akamai to arrange a redelivery. For more on delivery failures and requesting redeliveries, see Troubleshooting.

Formats and Example Log Lines

This chapter describes the available log formats and gives sample log lines for you to review. The "cs" and "sc" in log field names indicate the direction of communication, that is, client-to-server and server-to-client, respectively.

Luna Format Name	Description
dns	Fields: <ID> - <timestamp> <readable timestamp>,<requestor ip>,<requestor port>,<request name>,IN,<request type>,<"E" if EDNS0 request or blank>,<edns0 size if EDNS0 or blank>,<"D" if DNSSEC request or blank>,<"T" if TCP request or blank>,<answers>
Client Side Downloads	Application Version Information: Timestamp GUID ClientIP CPCODE AppVersion AppInfo StreamingFlags Certified Download: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize StreamingFlags StreamEvent StreamId HTTPSourceNum PeerSourceNum PeerOverheadBytes Download Status Update: Timestamp GUID ClientIP CPCODE TransId RL HTTPSourceDeltaBytes PeerDeltaBytes ValidDeltaBytes FailedDeltaBytes ObjectStatus ContentType StreamingFlags InstallCPCODE InstallTimestamp StreamEvent StreamId HTTPSourceNum PeerSourceNum PeerOverheadDeltaBytes EULA Acceptance Status: Timestamp GUID ClientIP CPCODE EulaTimestamp EulaVersion EulaAuthenticator

Luna Format Name	Description
	Failure - Client-Detected: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize Reason ExtraInfo StreamingFlags PeerOverheadBytes Initiated Download: Timestamp GUID ClientIP CPCODE TransId RL APIVersion StreamingFlags Installed a NetSession Interface Client: Timestamp GUID ClientIP CPCODE InstallTimestamp Paused Download: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize PauseReason StreamingFlags PeerOverheadBytes Resumed Download: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize PauseReason StreamingFlags PeerOverheadBytes
Combined **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent"
Combined + Edge tokenization **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "x-edge-tokenization"

Luna Format Name	Description
Combined + HDD Delivery **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" guid bitrate playout-time encrypted-bytes stream-format
Combined + Web App FireWall Log record format:	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" "waf-info"
Common **Log record format:**	client_ip - - [date] "http_method arl HTTP/1.1" status_code total_bytes
W3C + Edge Tokenization **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" "x-edge-tokenization"
W3C + HDD **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \tguid \t bitrate \tplayout-time \tencrypted-bytes \tstream-format
W3C + WebApp Firewall Log record format:	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"x-wafinfo
W3C **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl \t status_code \t total_bytes \t transfer_time
WMS9	Fields: c-ip date time c-dns cs-url c-starttime x-duration c-rate c-status c-playerid c-playerversion c-playerlanguage cs(User-Agent) cs(Referer) c-hostexe c-hostexever c-os c-osversion c-cpu filelength filesize avgbandwidth protocol transport audi-ocodec videocodec channelURL sc-bytes c-bytes s-pkts-sent c-pkts-received c-pkts-lost-client c-pkts-lost-net c-pkts-lost-cont-net c-resendreqs c-pkts-recovered-ECC c-pkts-recovered-resent c-buffercount c-totalbuffertime c-quality s-ip s-dns s-totalclients s-cpu-util cs-user-name s-session-id s-con-tent-path cs-media-name c-max-bandwidth cs-media-role s-proxied
NetStorage	Fields: log_ts region leader_ip (fields[4-12] : Varying values in these fields) server_ip protocol client_ip user_id log_id flags network_id fwd_host session_id start_ts end_ts transfer_ts u_ flags request_id proto_ver host_header error overhead_bytes bytes_read bytes_send
Combined + ARL without domain	client_ip domain - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"

Luna Format Name	Description
Log record format:	
Combined + Cookie **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"
Combined + Cookie + Accept Language **Log record format:**	client_ip - - [date] "http_method arl-stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "accept_language"
Combined + Cookie + ARL without domain **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"
Combined+Cookie+ARL without domain+Host Header+Transfer Time+Custom Field **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" host_header transfer_time "custom"
Combined + Cookie + Custom Field **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "custom"
Combined + Cookie + Host Header **Log record format:**	client_ip - - [date] host_header "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"
Combined + Cookie + Total Size + Byte Range Log record format:	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" total_object_size byte_range
Combined + Cookie + Total Size + Byte Range + Completion Status **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" total_object_size byte_range last_byte_served_flag_1_or_0
Combined + Referrer + User Agent **Log record format:**	client_ip - - [date] "http_method arl HTTP/1.1" status_code total_bytes "referrer" "user_agent"

Luna Format Name	Description
Combined + Time(ms) + Cookie + Custom Field **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "custom"
Combined + Time(ms) + Cookie + Host Header + Custom Field **Log record format:**	client_ip - - [date] host_header "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "custom"
Combined + UBB **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "UBB_GUID"
Combined + URI-stem Log record format:	Fields: date time c-ip cs-username cs-method cs-uri-stem cs-uri-query sc-status sc-bytes cs(User-Agent) cs(Cookie) cs (Referer)
Combined + Web App FireWall + Custom Field **Log record format**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "waf-info" "customfield"
Combined + Web App FireWall + Host Header **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "host header"
Extended **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie"
Extended + Accept Language **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t "accept_language"
Extended + Astraeus	Fields: date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken astraeus-used reliability-mode
Extended + Byte Range **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer" \t "user_agent" \t "cookie" \t total_object_size \t byte_range
Extended + Completion Flag	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time

Luna Format Name	Description
Log record format:	\t "referrer" \t "user_agent" \t "cookie" \t total_object_size \t byte_range \t last_byte_served_flag
Extended + Content-Type + Cache-related	Fields: start-time-unixtimestamp-dot-ms cs-ip cs-method ssl-version cs-uri sc-status sc-content-bytes sc-total-bytes s-object-size s-uncompressed-size http-overhead-bytes cs (Referer) cs(User-Agent) cs(Cookie) cs(Host) sc(Content-Type) x-time-ssloverhead-ms x-time-turnaround-ms x-time-transfer-ms x-request-id x-max-age x-cache-status x-cache-refresh-source x-last-byte-served-flag x-no-store-flag
Extended + CPU Stats	Fields: date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) normalized-cpu
Extended + Custom Field **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie"\t "custom"
Extended + Custom Field + Server Content Bytes **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"-customfield"\t"content-bytes"
Extended + Custom Field + Without Client IP **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t --- \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"customfield"
Extended + EdgeJava	Fields: date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) dca-app-str
Extended + Entity Size **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" "entity-size"
Extended + Hashed Client IP **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \client ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"-customfield"\t"content-bytes"
Extended + Host Header **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" "host_header"
Extended + Host Header + ARL without domain **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie"

Luna Format Name	Description
Extended + Origin Retrieval Status	Fields: date time cs-ip cs-method cs-uri-origin cs-uri-stem cs-uri-query sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) cs(Host-Header) cs-request-id cached
Extended + Request End Time	Fields: start-time-unixtimestamp-dot-ms end-time-ms cs-ip cs-method ssl-version cs-uri sc-status sc-content-bytes sc-total-bytes s-object-size s-uncompressed-size http-overhead-bytes cs(Referer) cs(User-Agent) cs(Cookie) cs(Host) sc(Content-Type) x-time-ssloverhead-ms x-time-turnaround-ms x-time-transfer-ms x-request-id x-max-age x-cache-status x-cache-refresh-source x-last-byte-served-flag x-no-store-flag Note: end-time-ms is https://-docs.akamai.com/esp/user/edgesuite/log-format.xml#r4
Extended + UBB **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t x-ubbguid
Extended + URI-Query Log record format	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t cs_stem \t cs_query \t status_code \t "user_agent"
Extended + Without Client IP **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t --- \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"-customfield"\t"content-bytes"
IIS 6.0 **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" "referer", status, -, S_ TOTAL_BYTES_SERVED, S_CLIENT_REQ_HDR_SIZE, S_ TRANSFER_TIME
Speedera **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)
Speedera + ARL without domain **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)
Speedera + ARL without domain + Last Byte Status **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)
Speedera + Last Byte Status **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)

Luna Format Name	Description
W3C + Web App Firewall + Host Header **Log record format`**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-wafinfo cs-host
Extended + Custom Field + Country code For Ghost IP + Sub Customer ID **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom country_code sub_customer_id
Extended + Custom Field + Country code For Ghost IP + Sub Customer ID + Host Header **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom country_code sub_customer_id host_header
Combined + Web App Firewall + Custom Field Without Query string **Log record format**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" "waf-info" "customfield"
Extended + Custom field + GhostIP + Cache Status **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom ghostip cache_status
Extended + Error Message **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) err-msg
W3C + Custom Field + GeoCode + Bot **Log record format**	starttime(msec) endtime(msec) cs-method cs-host cs-path cs-status cs-query cs(User-Agent) client-ip cs-ref x-custom
Extended + Custom field + GhostIP + CacheStatus + WAF **Log Record Format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom ghostip cache_status x-wafinfo
W3C + Custom Field + GeoCode + Bot + Full Filename in cs-path	starttime(msec) endtime(msec) cs-method cs-host cs-path cs-status cs-query cs(User-Agent) client-ip cs-ref x-custom

Luna Format Name	Description
Log Record Format	
Extended + Custom field + GhostIP + Cache Status + Host Header **Log Record Format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom ghostip cache_status cs(host-header)
BIND	<date> <time> "client" <requester IP>#<requester port> (<question name>): <query or response> <question name> <class> <question type> <flags> Flags include: • +/- (whether recursion desired or not) • "T" if TCP • "C" if CD set

Filenames and Email Subject Lines

Overview

Log filenames in both email and FTP are the same. Email subject lines are the same as filenames.

The format is:

[Identifier String]_[CP Code].[Format]_[Sorting Status].[Date and Time]-[Part Number].[Encoding]

For example:

customer_1234.esclf_S.200401250000-2400-0.gz

Log Identifier String

The log identifier string is a unique label that is added to the file names of log files. If you have multiple LDS configurations, the log file identifier string can make it much easier to distinguish them visually. Log identifier strings are limited to alphanumeric characters (numbers and letters) only, and may be no longer than 80 characters. They are case-sensitive, meaning that the case of names entered in the configuration tool will be preserved in log file names.

See *Appendix B: Filenames and Email Subject Lines*on page 25 for a complete description of log file names and email subject lines.

CP code

LDS can be per CP code. If you have multiple services associated with a single CP code (for example, multiple streaming formats), you might have several Log Delivery Service configurations using the same CP code (but only one per service per CP code).

Sorting status

Log files are labeled either "S" for sorted or "U" for unsorted. See *Appendix B: Sorting of Logs*on page 28 for more details on sorting.

Date and time

NOTE:

1. Times referenced in Log Delivery file names use GMT.

The date and time field is in the format YYYYMMDDHHMM-HHMM.

For example:

customer_1234.esclf_S.200401241200-2400-0.gz

will be followed by

customer_1234.esclf_S.200401250000-1200-0.gz

Part number

Log files may be split into parts to keep the individual parts below the configured maximum file size. Part numbering begins at 0.

Encoding

Log Delivery Service supports three mail formats: gzipped and uuencoded, MIME with base-64 encoding, and GPG Encrypted. Note that GPG Encrypted requires a key that you must upload. Otherwise, the main difference is support by mail clients. Most customers using email delivery are currently configured for gzipped and uuencoded logs.

There are four types of encoding, two of which are only available in email:

Type	Available via	Description
gz	FTP or email	GNU gzip
gpg	FTP or email	Gzipped and GPG encrypted
gzu	Email only	Gzipped and uuencoded
gzm	Email only	Mime (base64)

File names for deliveries aggregated by hit time

NOTE:

1. Times referenced in Log Delivery file names use GMT.

File names for deliveries using aggregation by hit time also include information about the time range that was examined for hits that occurred on the specified day and the completion percentage that was met when the delivery was scheduled.

The format is:

[Identifier String]_[CP Code].[Format]_[Sorting Status].[Time Range Start]-[Time Range End]-[Date]-[CompletionPercentage]-[PartNumber].[Encoding]

For example:

customer_1234.esclf_S.200401250000-200401260500-20040125-99.5c-0.gz

This file name indicates the bulk delivery of the data for 2004.01.25, using a completion threshold of 99.5 percent, and checking logs that were processed between 2004.01.25 00:00 GMT and 2004.01.26 06:00 GMT for hits that occurred on 2004.01.25.

Sorting of Logs

Within a single delivery, logs are sorted one of two ways:

- The logs are sorted before being broken into parts, which is indicated by an "S" in the filename. Reassembling the parts produces a complete, sorted log.
- The logs are broken into parts, and then each part is sorted separately, which is indicated by a "U" in the filename. Reassembling the parts does not produce a complete, sorted log.

Only deliveries where the number of log lines is above a certain threshold will be delivered with parts sorted separately; all others will be completely sorted. The current threshold is based on uncompressed log file size, and usually approximates 20 million log lines per delivery. This might change in the future as log volumes increase or more resources become available for sorting.

Note that despite sorting within deliveries, logs are processed and delivered in the order in which they are received. Nearly all logs arrive in the LDS system in a timely fashion and are processed very quickly. However, it is normal to see considerable overlap of a few hours between deliveries accompanied by a much smaller trickle of logs for up to a few days. Small variations in log arrival time guarantee that there will always be some overlap between the time periods covered by each delivery, and arrival of some logs to the LDS system can be delayed for longer due to network-related factors outside of Akamai's control. The exact pattern will vary depending on the geographic distribution of your end users.

Any automated processes or log-analysis software you are using will need to be configured to take this potential time overlap into account. See *Appendix B: Using 3rd-Party Log Analysis Tools* on page 29 for some further recommendations on working with LDS and log-analysis software.

Using 3rd-Party Log Analysis Tools

The following are general suggestions or workarounds to ensure the best possible compatibility between Akamai logs and 3rd-party log analysis tools.

Format Recognition

Akamai log formats generally follow widely accepted standards. Occasionally, log-analysis software has trouble with W3C or Combined Log Format. Switching to another format often solves the problem.

Chronological Ordering of Logs

Most log analysis software is designed to have all logs processed in chronological order, and will throw out log lines that are not in order. However, LDS is designed in such a way that there will always be overlap between deliveries (Akamai logs are sorted within deliveries – see Sorting of Logs for details).

Your software configuration will involve a trade-off between completeness of data and availability of reports. The longer you wait for all remaining logs to arrive so that you can merge and sort them with previous deliveries, the longer you have to wait for your reports. Because time-range overlap will vary from customer to customer, we recommend some initial examination of your logs from LDS before choosing a software configuration scheme.

WebTrends

Akamai has a close relationship with NetIQ, makers of WebTrends Reporting Center and WebTrends Analysis Suite. We have done some testing to make sure that our logs are as compatible with their software as possible.

To feed deliveries sorted in parts into WebTrends software, customers can use the following workaround. Define a WRC "profile" that considers the parts of an LDS delivery as coming from (in NetIQ terminology) a "server cluster", a web site hosted on multiple machines. Each part can be considered as the data from one server in the cluster. WRC has support to merge these files and to produce aggregate statistics on the entire collection of log files.

We have tested this workaround at Akamai and have verified with NetIQ that this "cluster support" is available in WebTrends Analysis Suite (advanced edition) and in WebTrends Reporting Center v2.0+. WebTrends Analysis Suite (standard) does not have this support.

Log Delivery Times

Log delivery times can vary depending on the load placed on the LDS infrastructure. The times of daily deliveries are especially vulnerable to load-induced variation due to the size and number of logs being delivered. It is usually possible to increase the predictability of delivery times by reducing the frequency of delivery. In any case, automated processes set up to manage incoming logs should allow for varying delivery times.

Troubleshooting

The following section provides recommendations and instructions for avoiding or fixing problems with Log Delivery Service.

Failure Notification Emails

If Akamai is unable to deliver your logs, an email will be sent to the admin contact you have specified in the configuration tool on Luna Control Center. It will give you the object (CP code or domain), date, and time range, some explanatory text about the error, and a transcript of the delivery session (either FTP or SMTP). It will be your responsibility to determine the cause of the failure and correct it if necessary before requesting a redelivery.

The help messages in the emails are based primarily upon the response codes returned, and are fairly general. Because of the variety of FTP and mail server configurations available, it is not always possible to tell what has gone wrong from the help message alone. In most cases it will be necessary to read the transcript.

To help you with this, we have provided some troubleshooting examples from delivery transcripts along with likely interpretations. All are taken from real delivery failures that we have seen, modified to make them anonymous. Remember that not all email and FTP servers produce the same error messages, and that response codes don't always correspond logically to the events that trigger them.

After reading the transcript, you might find it helpful to replicate the behavior of the LDS system. Try sending an email with a similar subject line and size, or delivering a file to the same FTP account. Note that many common problems affecting LDS are transient or intermittent, and might not be detectable during troubleshooting.

Example email errors

In many cases the symptoms of an email problem are several failed attempts ("bounces") concluding in failure after several days. LDS will send a failure notice on the initial bounce, even though the delivery might later succeed.

The mail server is rejecting the deliveries because there is no user with that address.

h80.67.64.10 does not like recipient.
Remote host said: 550 <user@lds-customer.com>: User unknown
DGiving up on 80.67.64.10

There is no known user with that address.

Either create the user, or change the LDS configuration.

The mail server is unreachable.

ZSorry, I wasn't able to establish an SMTP connection. (#4.4.1)

ZConnected to 80.67.64.10 but connection died. (#4.4.2)

This could have a variety of potential causes, including excessive packet loss, routing problems, or the email server being down.

Try sending mail to the address and checking connectivity to and from the mail server.

There are too many active connections to the mail server and it is rejecting new ones.

ZConnected to 80.67.64.10 but greeting failed.
Remote host said: 421 lds-customer.com connection limit reached
DConnected to 80.67.64.10 but sender was rejected.
Remote host said: 505 Authentication required

This is unusual, but can happen when multiple large deliveries are still in progress and overloading the email server. Anything that causes deliveries to take a long time could be the cause.

Note that in this case several previous deliveries might be effectively "hung", and might need to be terminated and redelivered.

LDS can't determine the IP address of the mail server due to DNS problems.

DSorry, I couldn't find any host named lds-customer.com. (#5.1.2)

This could be due to a single incident affecting your DNS server, intermittent DNS performance problems, connectivity problems to the DNS server, or a configuration problem with DNS records for the domain.

Try querying the DNS directly (possibly from a location outside of your corporate network), or checking connectivity to and from it.

Example FTP errors

DNS times out when trying to resolve the FTP server.

open: host name resolve timeout

This could be due to connectivity problems between LDS and the DNS server, or simply a DNS server that fails to respond.

Try querying the DNS directly (possibly from a location outside of your corporate network), or running a traceroute to or from it.

The FTP server can't upload a file of the same name.

Error 553 in FTP put: ftp_test: Permission denied on server. (Overwrite)

Please make sure the user has permission to upload a file into the destination

directory.

The FTP server might be unable to upload a file if the filename already exists. Rename the file and try to upload the file again.

The FTP server is full, or a quota is exceeded.

[8526] FTP.lds-customer.com <--- 452 Error writing file: No space left on device.

[22595] FTP.lds-customer.com <--- 425 Data connection error: No space left on device

[25814] FTP.lds-customer.com <--- 452 Transfer aborted. No space left on device

[26026] FTP.lds-customer.com <--- 452-Maximum quota exceeded. Transfer aborted.

The FTP server is full, or a quota is exceeded.

It is important that you have some method for ensuring that there will always be space for log deliveries. For many customers, that takes the form of automated removal of logs to another server. If this is being done on Akamai NetStorage, you can set up age-based deletion in Luna Control Center.

The FTP server is refusing connections.

[12620] FTP.lds-customer.com **** Socket error (Connection refused) – reconnecting

This error can occur because of a misconfigured access control list.

Check the FTP server. Try connecting to it yourself. Look for a misconfigured access control list.

The FTP server is unexpectedly closing the connection.

[9036] FTP1.lds-customer.com **** Peer closed connection

[22397] FTP2.lds-customer.com <--- 421
[22397] FTP2.lds-customer.com **** remote end closes connection
[22397] FTP2.lds-customer.com ---- Closing control socket

[32620] FTP3.lds-customer.com <--- 426 Connection closed; transfer aborted.

LDS can establish a connection, but it consistently times out before logs can be delivered.

[28859] FTP.lds-customer.com <--- 150 Opening BINARY mode data connection for customer_1234.esclf.200401250000-2400-0.gz.
[28859] FTP.lds-customer.com **** Timeout – reconnecting
[28859] FTP.lds-customer.com ---> ABOR
[28859] FTP.lds-customer.com ---- Closing control socket

[5019] 80.67.64.10 <--- 150 Opening BINARY mode data connection for customer_1234.esclf.200401250000-2400-0.gz.
[5019] 80.67.64.10 **** Broken pipe
[5019] 80.67.64.10 ---> ABOR
[5019] 80.67.64.10 ---- Closing control socket

The connection to the FTP server consistently times out before logs can be delivered. This error tends to be intermittent.

Check connectivity to your FTP server over time.

LDS Public GPG Key

The following is the public GPG key for Log Delivery Service:

-----BEGIN PGP PUBLIC KEY BLOCK-----

Version: GnuPG v1.4.11 (GNU/Linux)

mQGiBFTJAgsRBAD4RqtHK6AJjzenMk6JwaA6a7f09ZKObQ6mc61LblFml/BzqeIA

CZTvSGv31dPnEddyLnSFlAoVJCWIWl6DJ3rkanVqmX/kxCioGT3l9rmJWcEigHtj

GjGKW0YNWeWhz0WlzjuYP0yvGrLSxn6g8PRiBa6VT04T5MBgf4rivxSnhwCgxfgn

5lhVmHP/XNA+bWGnMCLkYYUD/29FIJEkw1IuNyBmf3RWaJzwsxHrz504G0phPqf/

D0k2RSDvIKp6rDrYDHHqRmRl7dHNOAz5qcGjQjEXP4SOFFQ1wxsVV+XJtMs27XWm

2E0Ge3AiC8ctQNhPca8uyyPOHbw4NrYoJgRGib1ZdiXdq9VB5PTxZsaaCniQ/tRm

16tgA/9If7k6ifJSTCL1X06MOONcuCnlwMEIHbyIpY6uZUSeJfMUuJp1MFBzoEDQ

fd1ph00e61IsosPOUqgfu7wuIIkpfvEKzlX3DLjHmgY/YBDDgqySAngtNy9mkS9I

MBs7zEn1q8JbIQrAAKvhNNZKGO/zoYs5hMui5tLoQ/JNR1bG2rRcQWthbWFpIExv

ZyBEZWxpdmVyeSBTZXJ2aWNlIChTZXJ2aWNlIGFkZHJlc3Mgb25seS4gIERvIG5v

dCByZXBseS4pIDxsZHNAbGRzLWxvZ3MuYWthbWFpLmNvbT6IYAQTEQIAIAUCVMkC

CwIbAwYLCQgHAwIEFQIIAwQWAgMBAh4BAheAAAoJEJ2+a/Qkpjjomt IAn132fwvu

kgBp6aupXSTNt52xQ095AJ93BAdwoW6Bq5Cm5QrbNUg954ouLrkCDQRUyQIMEAgA

pa7RIAjxtsCV/82Byi1Gv90jdIDEpD22XYo3kZ5TyPQqgG4rnx7yD6NrLeUPtHsY

XJhalqCNEPqn6+0u8sh4gwh9HghrooweY3Ly8CV/peW0DazvsHhYoYQIM4IxbUmY

zBCIJS5Pf7rHV9dFf+foAx0G80ZWyXxga/bBvTnZ1NWFJ56dkQBtTORng061mlUY

dbmyjqHWiQ+w0YDvC/+6pQn2n3DJCQ7kiazWpCQcab3I/4pv4t7bt7JmBhbPSDP5

M6piC8RBdDdWut1Kf8eboEwXrAxw0EaSal7vF981KjWGuo70i9EZWav++JIWl5uM

so2nIav1f3FBtYcGOUyoswADBQgAj3Githc66p8/hm1gGkYhyaZ9AHgqoHlWHDyu

WqTaNAwOyXqfg3fKCaD49kACFc/FM440NS24rJf11Js++UtBznzbr+1sxuRAHpT3

/+ZrdGvpS8ZwvTUpPdootqtqvC0simXh6Q0456HMnuOfE077BHXUe8OdY2tkAjSP

rZ7taZzNNCN5zf2+JzWMn7Obk6slKI/bV3eQpKZgqIbRC14FkuzdlTlfOW/WXcPL

/ekIXQ3T16LXjut1+SRUz+b/Zy3AS+TJ/ldCCSNqLnG2LtXtd52VmfXv2keoR7PT

5lcWFwMVxaoSh0XyJXNcIVoILN54D+auJpfLKk9COo3z/xCgvohJBBgRAgAJBQJU

yQIMAhsMAAoJEJ2+a/QkpjjohqUAnikbkUTbwkqp0ttFyRO+kHS4MyeSAJ9LoEKT

BWwW1DvqzY3Hka7RJkb5kA==

=WY8v

-----END PGP PUBLIC KEY BLOCK-----

Procmail Configuration Examples

Procmail can be used to automatically put incoming logs from email in a certain directory. Postfix uses procmail and sendmail as its default mail delivery agents. Therefore, you can use either Postfix or Sendmail and the following instructions will work.

(http://www.redhat.com/support/docs/faqs/RH-postfix-FAQ/x69.html)

Suppose that mail fromLog Delivery Service had a header that looks like this:

From: lds@<*>.akamai.com

Subject: Customer_100.wm_S.200010080400-0100-0.log.gz

Where the subject contains a filename, which is unique for the date. Suppose that you have a "virtual" user set up to receive the email called

incoming-logs@yoursite

Suppose you want the incoming log to be put into

/home/ftp/available/Customer_100.wm_S.200010080400-0100-0.log.gz

so that your software can grab it off of your in-house FTP server customers can automatically put incoming logs from email to a certain directory.

Then you would set up a procmailrc file in ~incoming-logs/.procmailrc as follows:[1]

---------- cut here for .procmailrc ----------

LOGFILE=procmail.log

SUBJECT=`formail -xSubject:`

[1]Disclaimer: Akamai provides these examples to help you set up your own systems to receive and handle the Log Delivery Service email. The above code is not guaranteed to function and should only serve as a guideline. See man procmailex for additional useful examples.

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

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.

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.

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.

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 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.

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.

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.

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.

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

Appendix 9

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.

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.

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.

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.

```
:0:

* ^From:.* lds@<*>.akamai.com

| cat > /home/ftp/available/${SUBJECT}

----------- cut here ------------
```

amazon web services | Partner Network
ADVANCED CONSULTING PARTNER
AUTHORIZED GOVERNMENT PARTNER
CHANNEL RESELLER PARTNER

DLT

Standard Support Package Overview for Amazon Web Services

DLT Solutions offers a comprehensive Standard Support Package for Amazon Web Services (AWS), utilizing our subject matter expertise and in-depth experience with IT infrastructure, cloud strategy, architecture development, and program management support. Our Standard Support Package includes access to:

- *The DLT Cloud Advisory Group* – a specialty knowledge-based practice designed to help our customers understand available cloud technologies and how best to leverage them. The DLT Cloud Advisory Group provides strategy, architectural and design support relating to all cloud service and cloud delivery models. Members of the Cloud Advisory Group average over 20 years of experience in the IT and services industries.
- *The DLT Service Center* – a 24x7x365 US-Citizen, US-Soil-based technical call center providing certified engineering resources in support of our customers' technical issues. The Service Center specializes in government support, serving the entire public sector, including Department of Defense and Intelligence communities.
- *The DLT Program Management Office (PMO)* – a team of experienced professionals who assist our government customers by providing a programmatic overview of service utilization and related financial spend.
- DLT's Standard Support Package also provides access to AWS' business level support resources, including backline" engineering, when needed.

All orders placed with DLT Solutions for AWS are required to include a DLT Standard Support Package.

Once a customer order is accepted, DLT Solutions creates a unique customer account for the Ordering Activity. The Ordering Activity has the ability to add, delete, enable, and/or disable services as desired. The Ordering Activity will be billed for actual usage, based on current unit prices, for the services activated or used during each billing period.

Pricing for the Standard Support Package is applied to the pre-discounted order total.

Product #	Description	Price
DLT-AWS-STD-SUPPT	DLT Standard Support Package	10% of total AWS spend

The Standard Support Package Overview

Telephone Support

The DLT Service Center (DSC) provides telephone support as a key component of DLT's Standard Support Package for AWS. Telephone support is available 24x7x365 with all calls answered by a live operator within 10 seconds.

Online Support

The DSC also supports email-based inquiries for opening new tickets or requests for updates on existing tickets.

Customer Training

The Standard Support Package is specifically designed to support request relating to:

- "How to" questions about AWS services and features
- Best practices to help integrate, deploy, and manage AWS services
- Troubleshooting operational or systemic problems with AWS resources
- Issues with the AWS Management Console or other AWS tools

The DLT Cloud Advisory Group and certified Architects are available to provide pre-sales consultative information relating to the AWS platform as well as other cloud services, their models, their usage, and best practices. If desired, DLT Solutions can develop custom training courses specifically tailored to the AWS resources being purchased. AWS also offers a wide collection of documents, articles, forums, and webcasts that provide everything from high level overviews of AWS services to technical deep dives at http://aws.amazon.com

Availability

The DSC is available for technical support 24x7x365.

Escalation, Responses, and Timefames

The DSC has defined the following incident severity levels and response times:

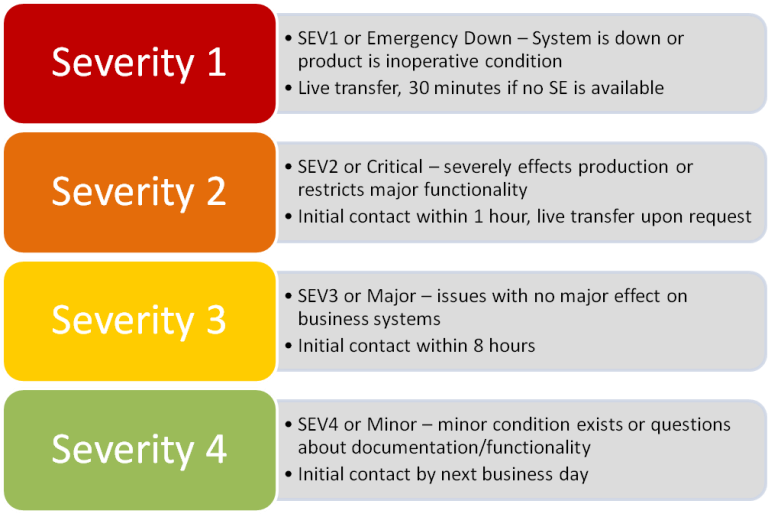

Account Creation Requests & Timelines

DLT Solutions handles all of the activities necessary to setup a new AWS account for our customers. Each customer account created is linked to a DLT AWS account for consolidated billing purposes.

It is important to note that DLT transfers account credentials to our customers/the Ordering Activity and retains *no access to our customers' data or the AWS services being consumed* within that account. The linkage is for consolidated billing purposes only and is used to provide DLT Solutions with utilization metrics associated with AWS service consumption for proper invoicing.

The high level process that DLT follows to create new customer/Ordering Activity accounts is as follows:

1. Gather pertinent information about the customer and project, including:
    a. Customer/Dept/Agency Name
    b. Customer Project Name
    c. Contract Officer Name, Email, and Phone
    d. Contract Officer Technical Resource (COTR) Name, Email, Phone (if applicable)
    e. Lead Technical Contact Name, Email, Phone
2. Establish an email distribution list in the form of aws-<CUSTOMER>-<PROJECT>@dlt.com, where <CUSTOMER> is the agency or organization name and <PROJECT> is the unique project identifier within that agency or organization.
3. Create a new customer account. The login name will be the same as the email forwarder name (aws-<CUSTOMER>-<PROJECT>@dlt.com).
4. Initiate invoicing and consolidated billing linkages. This step is a manual step that is performed by both AWS and DLT Solutions staff and can take up to 48 hours upon receipt of all required information described in Step 1. Please note that this is a manual process due to the fact that an AWS API for account creation is not yet available.
5. Send a Welcome Package to the customer/Ordering Activity. Once the account is created, DLT Solutions will send a Welcome Package that includes all pertinent information needed to begin using the account. This package will be accompanied by a Support Package detailing the information our customers need to initiate support requests through the DLT Service Center for AWS related questions and incidents.

Provisioning Requests

Once a customer's account has been created, they can access the AWS Management Console or AWS platform via API to provision and manage AWS services themselves. Most provisioning request are completed within minutes, but times may vary depending on the service or size of instance (EC2, RDS, SimpleDB, DynamoDB, etc.).

Real-time Reporting

AWS publishes its most up-to-the-minute information on service availability at http://status.aws.amazon.com. Customers may also subscribe to RSS feeds from this page to integrate service outage information with other monitoring tools.

Contact Information

For additional questions about the Standard Support Package for AWS, please contact your sales representative directly or contact The DLT Cloud Advisory Group:

cloud@dlt.com
1-855-CLOUD01
www.dlt.com/aws

2411 Dulles Corner Park
Suite 800
Herndon, VA 20171

User Traffic Report

Unique Visitors Report

Unique Visitors Report

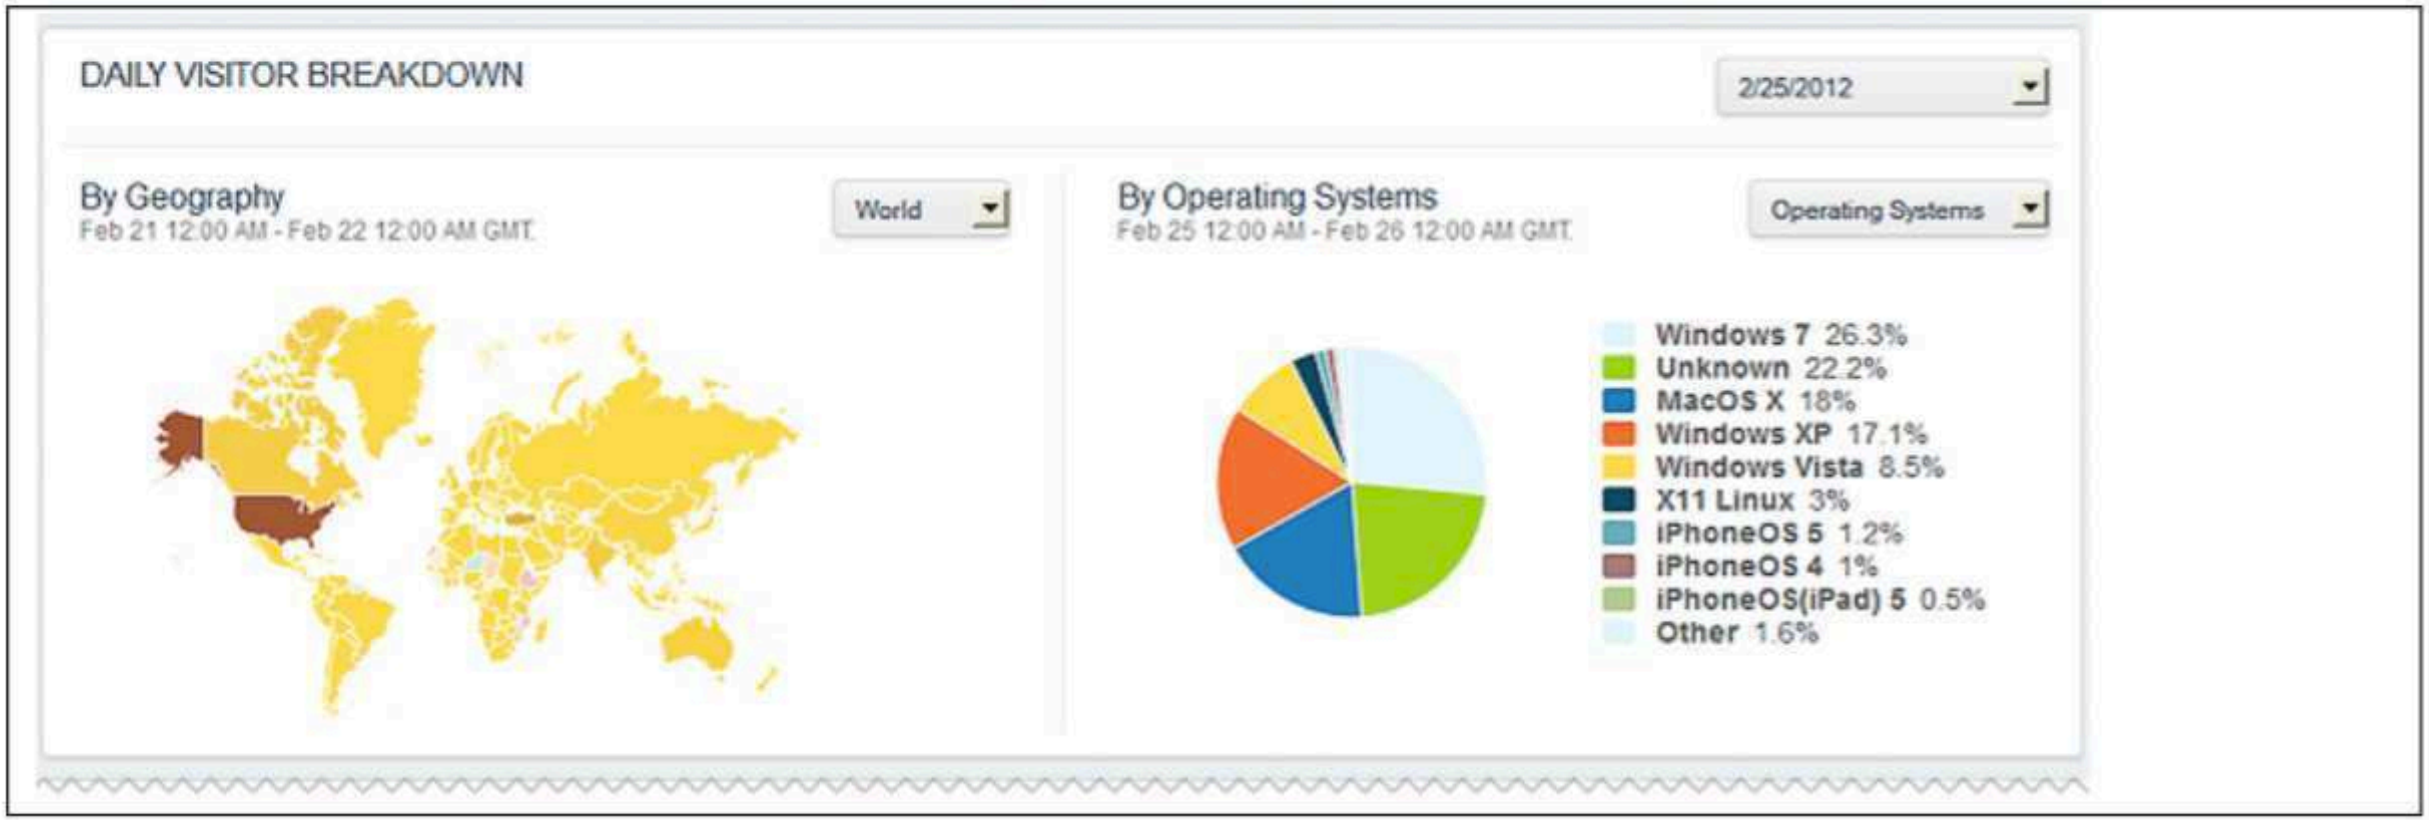

Offload Report

Responses Report

Attachment E

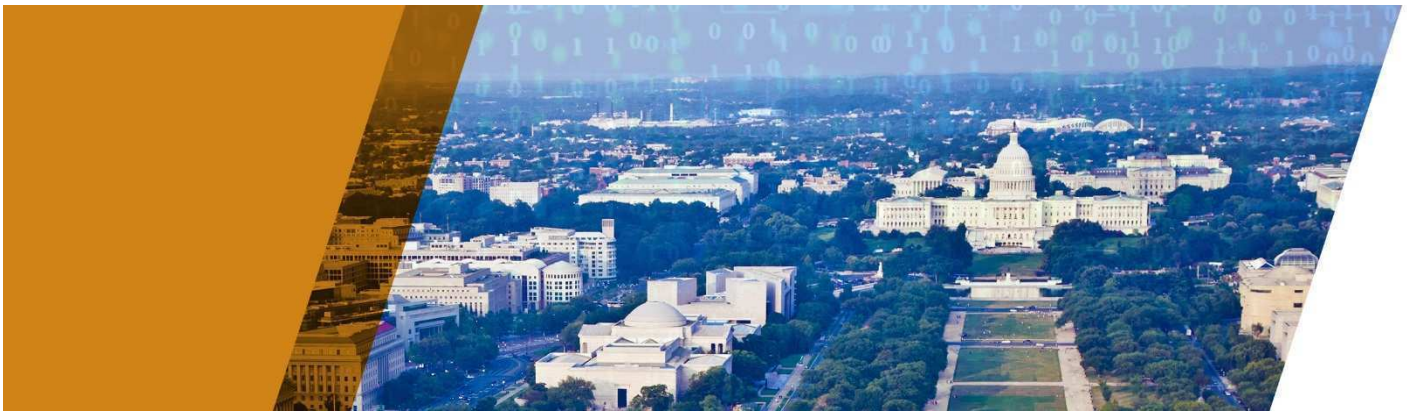

THE STATE OF UTAH
DIVISION OF PURCHASING

NASPO ValuePoint Master Agreement for Cloud Solutions

Utah Solicitation Number CH16012

Technical Proposal – Appendix 3

PRESENTED BY

March 10, 2016

DLT Solutions / 2411 Dulles Corner Park, Suite 800 / Herndon VA 20171 / 703.709.7172 / www.dlt.com
TIN 54-1599882 / DUNS # 78-6468199 / CAGE Code 0S0H9 / GSA Schedule 70 GS-35F-4543G

CONTENTS

1.0 AMAZON WEB SERVICES (AWS) SERVICE LEVEL AGREEMENTS

NOTE: Due to the rapidly evolving nature of AWS's product offerings, the most current SLAs are best reviewed directly on the AWS website via the links below:

- Amazon EC2 SLA: http://aws.amazon.com/ec2-sla/
- Amazon S3 SLA: http://aws.amazon.com/s3-sla
- Amazon CloudFront SLA: http://aws.amazon.com/cloudfront/sla
- Amazon Route 53 SLA: http://aws.amazon.com/route53/sla
- Amazon RDS SLA: http://aws.amazon.com/rds-sla

1.1 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:
    - For Amazon EC2, when all of your running instances have no external connectivity.
    - 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 by the end of the second billing cycle after which the incident occurred and must include:

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.

1.2 AMAZON S3 SERVICE LEVEL AGREEMENT

Last Updated September 16, 2015

This Amazon S3 Service Level Agreement ("SLA") is a policy governing the use of Amazon Simple Storage Service ("Amazon S3") 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 S3. 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 S3 available with the applicable Monthly Uptime Percentage (as defined below) during any monthly billing cycle (the "Service Commitment"). In the event Amazon S3 does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Error Rate" means: (i) the total number of internal server errors returned by Amazon S3 as error status "InternalError" or "ServiceUnavailable" divided by (ii) the total number of requests for the applicable request type during that five minute period. We will calculate the Error Rate for each Amazon S3 account as a percentage for each five minute period in the monthly billing cycle. The calculation of the number of internal server errors will not include errors that arise directly or indirectly as a result of any of the Amazon S3 SLA Exclusions (as defined below).

- "Monthly Uptime Percentage" is calculated by subtracting from 100% the average of the Error Rates from each five minute period in the monthly billing cycle.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon S3 account.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for Amazon S3 for the billing cycle in which the error occurred in accordance with the schedule below.

For all requests not otherwise specified below:

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99.0% but less than 99.9%	10%
Less than 99.0%	25%

For requests to Amazon S3 Standard – Infrequent Access (Standard-IA):

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 98.0% but less than 99.0%	10%
Less than 98.0%	25%

We will apply any Service Credits only against future Amazon S3 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 error 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 S3 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 by the end of the second billing cycle after which the incident occurred and must include:

1. the words “SLA Credit Request” in the subject line;

2. the dates and times of each incident of non-zero Error Rates that you are claiming; and

3. 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 applicable to the month of such request is confirmed by us and is less than the applicable 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 S3 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon S3, or any other Amazon S3 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 S3; (iii) that result from any actions or inactions of you or any third party; (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); or (v) arising from our suspension and termination of your right to use Amazon S3 in accordance with the AWS Agreement (collectively, the “Amazon S3 SLA Exclusions”). If availability is impacted by factors other than those used in our calculation of the Error Rate, then we may issue a Service Credit considering such factors at our discretion.

1.3 AMAZON CLOUDFRONT SERVICE LEVEL AGREEMENT

Last Updated June 1, 2013

This Amazon CloudFront Service Level Agreement ("SLA") is a policy governing the use of Amazon CloudFront 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 CloudFront. 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 CloudFront available with a Monthly Uptime Percentage (defined below) of at least 99.9% during any monthly billing cycle (the "Service Commitment"). In the event Amazon CloudFront does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Error Rate" means: (i) the total number of internal server errors returned by Amazon CloudFront divided by (ii) the total number of requests during that five minute period. We will calculate the Error Rate for each Amazon CloudFront account as a percentage for each five minute period in the monthly billing cycle. The calculation of the number of internal server errors will not include errors that arise directly or indirectly as a result of any of the Amazon CloudFront SLA Exclusions (as defined below).

- "Monthly Uptime Percentage" is calculated by subtracting from 100% the average of the Error Rates from each five minute period in the monthly billing cycle.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for Amazon CloudFront for the billing cycle in which the error occurred in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99% but less than 99.9%	10%
less than 99%	25%

We will apply any Service Credits only against future Amazon CloudFront 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 error 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 CloudFront 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 by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each incident of non-zero Error Rates that you are claiming; and

iii. 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 applicable to the month of such request is confirmed by us and is less than 99.9%, 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 CloudFront SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon CloudFront, or any other Amazon CloudFront 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 CloudFront; (iii) that result from any actions or inactions of you or any third party; (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) arising from our suspension and termination of your right to use Amazon CloudFront in accordance with the AWS Agreement; (vi) that result from exceeding usage limits stated in the Amazon CloudFront documentation; or (vii) that result from use of an origin server other than Amazon S3 (collectively, the **“Amazon CloudFront SLA Exclusions”**). If availability is impacted by factors other than those used in our calculation of the Error Rate, then we may issue a Service Credit considering such factors at our discretion.

1.4 AMAZON ROUTE 53 SERVICE LEVEL AGREEMENT

Last Updated May 15, 2015

This Amazon Route 53 Service Level Agreement ("SLA") is a policy governing the use of Amazon Route 53 (including Private DNS) 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 Route 53. 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 Route 53 100% Available (defined below). In the event Amazon Route 53 does not meet the foregoing commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "100% Available" means that Amazon Route 53 did not fail to respond to your DNS queries during a monthly billing cycle.
- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon Route 53 account.

Service Credits

Service Credits are calculated based on 1 day of Service Credit, which is equal to your average daily Amazon Route 53 query charges for the monthly billing cycle preceding the monthly billing cycle in which the period that Amazon Route 53 was not 100% Available occurred, and are available as follows:

Duration Amazon Route 53 was not 100% Available	Service Credit
5 - 30 minutes	1 day Service Credit
31 minutes - 4 hours	7 days Service Credit
More than 4 hours	30 days Service Credit

We will apply any Service Credits only against future Amazon Route 53 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 error 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 Route 53 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 by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each period that Amazon Route 53 was not 100% Available that you are claiming; and

iii. 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 period that Amazon Route 53 was not 100% Available is confirmed by us, 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 Route 53 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon Route 53, or any other Amazon Route 53 performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of Amazon Route 53’s reasonable control, including any force majeure event, denial-of-service attack or Internet access or related problems beyond the demarcation point of Amazon Route 53; (iii) that result from any actions or inactions of you or any third party; (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) arising from our suspension and termination of your right to use Amazon Route 53 in accordance with the AWS Agreement; (vi) that result from you exceeding usage limits stated in the Amazon Route 53 documentation (collectively, the “Amazon Route 53 SLA Exclusions”); or (vii) that, with respect to public DNS only, result during a period that you were not using all four virtual name servers (for example, ns123.awsdns.com, ns123.awsdns.net, ns123.awsdns.co.uk and ns123.awsdns.org) assigned to your “hosted zone”. If availability is impacted by factors other than those used in our calculation of 100% Available, then we may issue a Service Credit considering such factors at our discretion.

1.5 AMAZON RDS SERVICE LEVEL AGREEMENT

Last Updated June 1, 2013

This Amazon RDS Service Level Agreement ("SLA") is a policy governing the use of the Amazon Relational Database Service ("Amazon RDS") under the terms of the AWS 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 RDS. 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 Multi-AZ instances available with a Monthly Uptime Percentage (defined below) of at least 99.95% during any monthly billing cycle (the "Service Commitment"). In the event Amazon RDS does not meet the Monthly Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Monthly Uptime Percentage" for a given Multi-AZ instance is calculated by subtracting from 100% the percentage of 1 minute periods during the monthly billing cycle in which the Multi-AZ instance was "Unavailable". If you have been running that Multi-AZ instance for only part of the month, your Multi-AZ instance is assumed to be 100% available for the portion of the month that it was not running. Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon RDS SLA Exclusion (defined below).

- "Multi-AZ instance" means an Amazon RDS for MySQL,Oracle or PostgreSQL database instance with the Multi-AZ parameter set to true.

- "Unavailable" means that all connection requests to the running Multi-AZ instance fail during a 1 minute period.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Credits

Service Credits are calculated as a percentage of the charges paid by you for the Multi-AZ instances that did not meet the Monthly Uptime Percentage commitment in a billing cycle 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%	25%

We will apply any Service Credits only against future Amazon RDS 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 or non-performance or other failure by us to provide Amazon RDS 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 will need to submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each Unavailability incident you are claiming;

iii. the DB Instance IDs and the AWS Regions of the affected Multi-AZ instances; and

iv. 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 the request occurred. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon RDS SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon RDS, or any other Amazon RDS 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 RDS;

(iii) that result from any voluntary actions or inactions from you or any third party (e.g., rebooting a database instance, scaling compute capacity, not scaling storage when the storage is full, misconfiguring security groups, VPC configurations or credential settings, disabling encryption keys or making the encryption keys inaccessible, etc.);

(iv) that result from instances belonging to the Micro DB instance class or other instance classes which have similar CPU and memory resource limitations;

(v) that result from you not following the basic operational guidelines described in the Amazon RDS User Guide (e.g., overloading a database instance to the point it is inoperable, creating excessively large number of tables that significantly increase the recovery time etc.);

(vi) caused by underlying database engine software that lead to repeated database crashes or an inoperable database instance;

(vii) that result in long recovery time due to insufficient IO capacity for your database workload;

(viii) 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); or

(ix) that result from any maintenance as provided for pursuant to the AWS Agreement; or

(x) arising from our suspension and termination of your right to use Amazon RDS in accordance with the AWS Agreement (collectively, the “Amazon RDS SLA Exclusions”).

If availability is impacted by factors other than those explicitly used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

2.0 AKAMAI SERVICE LEVEL AGREEMENTS

2.1 AKAMAI CLOUD SECURITY SOLUTIONS SERVICE LEVEL AGREEMENTS (FEBRUARY 2016)

Definitions

"Attack Monitoring Services - Failure to Notify Event" is an event in which Akamai fails to take the defined steps to notify Customer within a period of 15 minutes from the time that Akamai's Security Operations Center (SOC) receives a Critical alert (applicable only to Prolexic Application-Based Monitoring and Prolexic Flow-Based Monitoring Services deployed at the Customer site).

"**Akamai Network**" means the distributed network owned and operated by Akamai.

"**Akamai Prolexic Network**" means the distributed network of specialized network of scrubbing centers owned and operated by Akamai.

"Availability Outage" (applicable only to Kona Site Defender and Kona DDoS Defender) is defined as a period of at least two consecutive failed attempts six (6) minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the contracted security service for the day in which the failure occurred, not to exceed 30 days of fees.

"Emergency Maintenance" means any activity that Akamai, in its sole discretion, deems necessary to correct an immediate threat to the ongoing availability and quality of Akamai's Service offerings

"Managed Kona Site Defender Response Service SLA Violation" – the inability of the Akamai support team to meet the Response time or Live Support Availability as defined.

"**Managed Kona Site Defender Service Initial Response Time"** (applicable only to Support Requests filed against the Kona Site Defender product under Managed Kona Site Defender Service) is the time it takes Customer to get a response on the reported issue from an Akamai technical support representative.

The measurement of the initial response time is the elapsed time from the start of the Security Incident Management process, to the response to Customer by an appropriate service resource to acknowledge the request, respond with a service request number and begin working the issue. This includes time until a response is received in the form of a call back or e-mail or any other customer facing communication. The degree of urgency can vary based upon the issue's priority level.

For Security Events identified through Akamai Managed Kona Site Defender Service Security Event monitoring, Security Incident Management begins once a Security Event has been observed and that event cannot be classified as a false positive and the issue is escalated to the customer.

For Security Events identified by the customer, the Security Incident Management process begins from the time the event is reported by the customer to Akamai Customer Care.

All Support Requests reported via e-mail will be considered as Severity 3.

"Security Severity Levels"

- **"Severity 1"** -- Critical Impact: This class exhibits: a) loss or outage on any portion of a protected property, b) data breach (exfiltration or infiltration) confirmed in progress, or c) defacement of a protected property.

- **"Severity 2"** – Major Impact: This class exhibits: a) degradation in performance on any portion of a protected property, b) suspected data breach, or c) excessive bot activity that may lead to intellectual property compromise.

- **"Severity 3"** – Low Impact: This class exhibits: a) signs of a potential small-scale security incident (log event evidence of malicious traffic that does not impact the origin and may be false positive, b) is a proactive action; "heightened attention" in response to a public threat, for instance, c) includes a possible fraud investigation without immediate evidence of data breach, or d) low-level site scraping activity.

"**Service Outage"** (applicable only to Prolexic Routed, Prolexic Connect, and Prolexic Proxy) means that Akamai's Prolexic Network did not respond to DNS or HTTP queries or the forwarding of IP traffic for more than sixty (60) consecutive seconds.

Akamai's Time To Mitigate and Consistency of Mitigation Service Levels (applicable only to Prolexic Routed, Prolexic Connect, and Prolexic Proxy)

With respect to Customers subscribing to Prolexic Routed, Prolexic Connect, and Prolexic Proxy, Akamai offers a service level ("Service Level") committing to the length of time that it will take Akamai to effectively deploy mitigation.

The Service Level begins at the time that a critical alert is generated by Akamai for Customers subscribed to the standard Always-On mitigation service or for Customers who are otherwise permitted to be running traffic through Akamai's Prolexic Network when a DDoS attack is identified. The time of the critical alert will be determined by relevant ticket correspondence and/or critical alerts in the Akamai portal.

The Service Level for Customers subscribed to an On-Demand mitigation service, if not currently routed through the Akamai Prolexic Network, begins after a Customer notifies Akamai and properly routes traffic through Akamai's Prolexic Network during a DDoS attack. The Time to Mitigate ("TTM") value for these On-Demand Customers depends upon the length of time for the Customer to properly route traffic through Akamai's Prolexic Network, and the length of time it takes for routes to propagate to the Internet at large.

Akamai's Service Level for te following attack types is available exclusively to Prolexic Routed, Prolexic Connect and Prolexic Proxy Services Customer. At a minimum, Service Validation is required once annually. Akamai commits to the following TTM, for each DDoS attack type, as categorized per following:

Attack Type	TTM – Time to Mitigate (typical)	TTM – Time to Mitigate – Guaranteed (Service Level)
* UDP/ICMP Floods	1 minute or less	5 minutes
* SYN Floods	1 minute or less	5 minutes
* TCP Flag Abuses	1 minute or less	5 minutes
* GET/POST Floods	10 minutes or less*	20 minutes
* DNS Reflection	5 minutes or less**	10 minutes
* DNS Attack	5 minutes or less**	10 minutes

** Mitigation requiring traffic analysis and custom signature deployment*

Akamai's Time To Mitigate and Consistency of Mitigation Service Levels (applicable only to Kona DDoS Defender and Managed Kona Site Defender Service)

With respect to Customers subscribing to Kona DDoS Defender and Managed Kona Site Defender Service, Akamai offers a service level ("Service Level") committing to the length of time that it will take Akamai to effectively deploy mitigation, meaning:

- Initial mitigations have been deployed
- They have been effective at mitigating the impact of the immediate attack.
- The benefits of the mitigation were evident within the time window of the SLA.

The Service Level begins at the time that a critical alert is generated by Akamai for Customers subscribed and integrated to the standard Always-On mitigation service when a DDoS attack is identified. The time of the critical alert will be determined by relevant ticket correspondence and/or critical alerts in the Akamai portal.

Akamai's Service Level only for the following attack types is available exclusively to Kona DDoS Defender and

Managed Kona Site Defender Service Customers. At a minimum, a Table Top Drill for Kona DDoS Defender, or a Threat Update Review and a Table Top Drill for Managed Kona Site Defender Service is required once annually and Akamai's Security Specialist recommendations must have been applied to the configuration. Akamai commits to the following TTM, for each DDoS attack type, as categorized per following:

Attack Type	TTM – Time to Mitigate (typical)	TTM – Time to Mitigate – Guaranteed (Service Level)
* UDP/ICMP Floods	1 minute or less	5 minutes
* SYN Floods	1 minute or less	5 minutes
* TCP Flag Abuses	1 minute or less	5 minutes
* GET/POST Floods	10 minutes or less*	20 minutes

Attack Type	TTM – Time to Mitigate (typical)	TTM – Time to Mitigate – Guaranteed (Service Level)
* DNS Reflection	5 minutes or less**	10 minutes
* DNS Attack	5 minutes or less**	10 minutes

** Mitigation requiring traffic analysis and custom signature deployment*
*** Applies to DNS attacks targeting Akamai IP addresses*

Akamai's Consistency of Mitigation Service Level (applicable only to Prolexic Routed, Prolexic Connect, Prolexic Proxy, Kona DDoS Defender and Managed Kona Site Defender Service)

Akamai offers a 95% Consistency of Mitigation Service Level. Consistency of Mitigation is measured by analyzing the ratio of clean traffic to attack traffic that is forwarded to the Customer. Measurement of the Consistency of Mitigation parameter begins after the committed TTM has elapsed. Claims against the Consistency of Mitigation Service Level must be submitted with a packet capture of at least one hour in duration, identifying the total amount of attack traffic forwarded during the event envelope. The event envelope is defined as all or part of the period between the TTM Service Level period and the end of the attack. Evidence of forwarding of attack traffic in excess of 5% of the total traffic volume qualifies for a credit under this Service Level clause.

Remedy for Time to Mitigate and Consistency of Mitigation Service Levels

The TTM is based from the time that traffic is properly routed through Akamai's Prolexic Network or Akamai Network for On-Demand Customers or from the time that a critical alert is generated for services that are AlwaysOn or already routed through Akamai's Prolexic Network or Akamai Network. The TTM is measured based upon the Consistency of Mitigation Service Level terms. During any given calendar month, if Akamai fails to meet the TTM Service Level as measured by the Consistency of Mitigation parameters set forth above, the following credits will be issued:

- Single event – in the event that the TTM Service Level is exceeded – with mitigation not meeting the Consistency of Mitigation Service Level, Akamai will credit Customer's account for such month for the pro-rated charges as follows:
    - Less than one hour: for (1) day of Monthly Service Fees due in respect of the affected Network Protection Services
    - For one hour or more, and less than 6 hours: two (2) days of Monthly Service Fees payable in respect of the affected DDoS Mitigation Services; and
- Multiple Events or Single Event lasting more than 6 hours – in the event that the Time to Mitigate Service Level is exceeded – with mitigation not meeting the Consistency of Mitigation Service Level for a period of six (6) hours or more, or during four (4) or more events within a calendar month, Customer will be credited with seven (7) days of Monthly Service Fees payable in respect of the affected DDoS Mitigation, or Managed Kona Site Defender Services.

All Customers must have successfully completed a Table Top Drill, with any prefix(es) affected, within the previous twelve months in order to qualify for remedy credit under the Time to Mitigate and Consistency of Mitigation Service Levels.

Akamai's Service Availability Service Level (applicable to Prolexic Connect, Prolexic Routed, and Prolexic Proxy Service Outage)

Akamai will provide any credits to Customer (or to Reseller for transfer to Customer, if Reseller is the contracting entity) per the following: Provided Customer reports a Service Outage to Akamai promptly following the occurrence of an event of interruption in Service that Customer believes is a Service Outage, but in any event no later than five (5) days after the event took place, Customer shall be entitled to receive a service credit for Customer's benefit in accordance with the schedule below. Whether an interruption in Services constitutes a Service Outage shall be determined solely by Akamai in its sole good faith discretion supported by records, data and other evidence. If a Service Outage has taken place and Customer notifies Akamai as provided in this Section, Akamai shall provide a credit to Customer as follows:

(i) If a particular Service Outage reported by Customer lasted for more than one minute but less than four (4) consecutive hours during a calendar month, Akamai will credit Customer for such month, the pro-rated charges for one (1) day of Monthly Service Fees of the amount of revenue Akamai receives from Customer with respect to the affected DDoS Mitigation Service(s); or

(ii) If a particular Service Outage reported by Customer lasted for four (4) or more consecutive hours during a calendar month, a credit equal to two (2) days of the Monthly Service Fees payable of the amount of revenue Akamai receives from Customer with respect to the affected DDoS Mitigation Service(s).

The above provision sets forth Customer's sole and exclusive remedy for Service Outages and any other interruptions or failures of Akamai's Managed DDoS Mitigation Service.

Remedy – Akamai Attack Monitoring Services (applicable only to Application-Based and Flow-Based Monitoring Services)

A Customer subscribing to the Akamai Application-Based Monitoring or Flow-Based Monitoring Service is entitled to remedy credit in accordance with this subsection should an Attack Monitoring Services - Failure to Notify Event occur, provided Customer reports the incident to Akamai promptly following the occurrence of an event that Customer believes is an Attack Monitoring Services - Failure to Notify Event, but in any event no later than five (5) calendar days after the event. Whether an incident constitutes an Attack Monitoring Services - Failure to Notify Event shall be determined by Akamai in its sole good faith discretion supported by records, data and other evidence.

(i) If an Attack Monitoring Services - Failure to Notify Event occurs once or more times during a calendar month, Akamai will credit Customer's account for the pro-rated charges for one (1) day's Monthly Service Fees due for each incident, in respect of the affected site(s)' Services; and

(ii) In addition to Customer being entitled to the above credits, in the event that three or more Attack Monitoring Services - Failure to Notify Events occur during a calendar month, Customer shall

have the right, for 30 days following the start of such incident, to terminate the affected Service, without liability.

Akamai's Availability and Performance Service Level (applicable only to Kona Site Defender, Kona DDoS Defender, Web Application Firewall, Bot Manager and Site Shield)

- **Availability SLA:** Akamai offers a service level ("Service Level") committing to 100% availability of the contracted security service.

  The Service Level begins at the time the customer has successfully completed integration and Service Validation for the protected properties, or the contract Billing Effective Date, whichever is later.

- **Performance SLA:** Akamai offers a service level ("Service Level") committing that the security service will not impede origin performance in any period that the protected digital property is not under attack.

  The Service Level begins at the time the customer has successfully completed integration and Service Validation for the protected properties, or the contract Billing Effective Date, whichever is later.

Activation of the Availability and Performance Service Level Agreements occurs once the Customer has successfully completed the following: Customer must enter and indicate the location of two valid test files for the same object (as described in II(c and (d above into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal . Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and, for customers using the optional Site Shield solution, Akamai Professional Services. Customers using Remote Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect five (5 business days after the Customer enters valid test files into the SLA Activation Tool.

Remedy for Availability and Performance Service Levels (applicable only to Kona Site Defender, Kona DDoS Defender, Web Application Firewall, Bot Manager and Site Shield)

If the Service fails to meet the defined service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly security service fee for the day for the protected origin(s) in which the failure occurs, not to exceed 30 days of fees.

The following methodology will be employed to measure the availability and performance of the security service:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's protected origin servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations using a test file on the customer's protected origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the protected origin server (ie, http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Kona Site Defender or the optional Site Shield solution.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the protected origin server (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the protected Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's protected production origin (measured directly, or via a Site Shield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the protected property, for that day in which the failure occurs, not to exceed 30 days of fees.

Akamai's Managed Kona Site Defender Service Response Service Level Agreement (applicable only to Managed Kona Site Defender Service)

Akamai agrees to provide a level of service to Customer's purchasing Managed Kona Site Defender Service as follows:

- Response Time
    - Severity 1 ≤ 30 minutes (must be opened via phone)
    - Severity 2 ≤ 1 hour
    - Severity 3 ≤ 1 business day
- Live Support Availability: An Akamai representative will be available live on the phone to respond to Severity 1 (Critical Impact) and Severity 2 (Major Impact) Service issues 24 hours a day, 7 days a week and 365 days a year. Live Support Availability for severity 3 (Low Impact) cases will be available during normal business hours, Monday through Friday, excluding local holidays, in the following geographies as follows:
    - North America (GMT – 05:00): 9:00 am to 9:00 pm ET
    - Europe (GMT): 08:00 am to 5:00 pm
    - Asia-India (GMT + 05:30): 9:00 am to 6:00 pm
    - Asia-Japan/Singapore (GMT + 08:30): 9:00 am to 6:00 pm

Remedy for Managed Kona Site Defender Service Response SLA Violation

In the event of a Managed Kona Site Defender Service Response SLA Violation, Customer must submit a written request for a credit (email request acceptable) to Customer's applicable Akamai relationship manager within seven days of the alleged SLA Violation. For acknowledged SLA Violations, Customer will receive (as its sole remedy) a credit equal to Customer's monthly Managed Kona Site Defender Service fee for the day in which the failure occurs, not to exceed 30 days of fees per month.

Remedy Terms – General

In order for Akamai to issue a credit in accordance with this SLA, Customer must have an account that is current with payments and in good standing with Akamai, and must be able to confirm that Customer has completed the Integration process for the applicable Service and, if applicable, all Competing Mitigation Hardware and Techniques have been disabled or removed during any mitigation services.

Credits shall only apply for Services provided pursuant to the Monthly Service Fee and/or Monthly Service Overage Fee, and will not apply to any other Service. Customers with subscriptions for more than one DDoS Mitigation Service will only receive credits for affected portion of DDoS Mitigation Service(s). The aggregate credits to be provided in any calendar month shall not exceed 25% of the Monthly Service Fee in respect of the affected Service(s).

2.2 DYNAMIC SITE DELIVERY SERVICE LEVEL AGREEMENT (07/18/2007)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

(a) 100% Uptime: The Service will serve content 100% of the time.

(b) Measurable Performance Enhancement: The Service will deliver content measurably faster than the Customer's web site.

(c) Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject), or via an Akamai SiteShield region if the customer is using the SiteShield feature.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a SiteShield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a SiteShield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a SiteShield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate the Dynamic Site Delivery Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and, for customers using SiteShield, Akamai Professional Services. Customers using Remote SiteShield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the SiteShield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

2.3 FASTDNS SERVICE SERVICE LEVEL AGREEMENT

I. Service Level and Credits

Akamai agrees to provide a level of service to the FastDNS customer providing:

(a) 100 % Uptime: For customer owned DNS zones configured to use the FastDNS service, Akamai will serve DNS resolutions 100% of the time.

(b) Service Credits: If the FastDNS service experiences an outage as defined below as a result of a failure of any component of the FastDNS service operated, managed or maintained by Akamai, the FastDNS customer will receive a Service Credit. A Service Credit is defined as equal to the pro rata FastDNS service fees for the continuous duration of the outage.

In the event that the FastDNS service SLA has not been met, the FastDNS customer shall request a credit via email or in writing to the customer's Account Manager at Akamai within 7 days of the outage.

II. FastDNS Outage Definition

(a) An outage is defined as a 12-minute period of consecutive failed attempts from multiple networks to resolve a DNS record configured for DNS resolution by the FastDNS service.

(b) Testing methodology for the FastDNS service must meet the following requirements:

- A minimum of five (5) testing agents in diverse geographic and network topological locations.
- A sample of at least two (2) data points per testing agent (i.e. a minimum polling interval less than 12 minutes on each testing agent).
- Test results must demonstrate that each testing agent was successful performing a DNS query for the DNS zone to the root DNS servers (.com, .net, .gov, etc.), but unsuccessful performing a DNS query for the DNS zone to the FastDNS service.

2.4 GLOBAL TRAFFIC MANAGEMENT SERVICES SERVICE LEVEL AGREEMENT (04/12/2006)

I. Service Levels and Credits

Akamai agrees to provide a level of service demonstrating:

a. 100% Uptime: Every client request through the Akamai Global Traffic Management service will be directed to a live Customer server 100% of the time.

b. Credits: If the Akamai Global Traffic Management service fails to meet the above service level, the Customer will receive (as its sole remedy) a credit equal to fees for the day in which the failure occurs.

II. SLA Monitoring Methodology

The following methodology will be employed to measure Global Traffic Management service availability:

Agents and Polling Frequency

a. Akamai will perform a customer-specified liveness test against each of the Customer's production servers from at least five (5) geographically and network-diverse locations in major metropolitan areas. Liveness testing occurs at the customer-specified frequency.

b. Akamai will also perform one DNS resolution via the Akamai Global Traffic Management service. Resolutions will take place at least once every 15 minutes.

c. There will be a separate test for each property (sub-domain) registered with Global Traffic Management.

III. Outages

An "Outage" is defined as a failed attempt by the Akamai Global Traffic Management service to resolve the name of a registered property to a live Customer server when there is at least one (1) live Customer web server. If an Outage is identified by this method, the customer will receive (as its sole remedy) a credit equivalent to the fees for the day in which the outage occurred.

2.5 AKAMAI ION PREMIER NORTH AMERICA DOMESTIC SERVICE LEVEL AGREEMENT (MAY 2014)

I. Service Levels

Akamai agrees to provide a level of service for customers with origin infrastructure hosted in North America demonstrating:

(a) Improvement Over Origin: The daily average page delivery time using the Ion Premier service will be at least 30% faster than the daily average delivery time for the same page delivered from the customer's origin server (defined as the daily average page delivery time from origin being at least 1.3 times the daily average page delivery time using Akamai's Ion Premier service).

(b) 100% Availability: The Service will serve content 100% of the time.

II. SLA Monitoring Methodology

A. Performance SLA Monitoring Methodology

The following methodology will be employed to measure the performance improvement provided by the Service:

(1) A single static unauthenticated page (including HTML and associated embedded content), provided by Akamai, will be tested throughout the term of the Ion Premier Order Form for purposes of this SLA.

(2) Delivery times will be tested using Akamai's Site Analyzer testing services to measure the daily average page delivery time computed from the complete global set of available Site Analyzer measurement agents. The tests will be configured to take place once an hour from the measurement agents.

(3) This SLA assumes that there will be no material changes to the test content including, without limitation, metadata applied to the content, agent network used for testing, origin settings and origin infrastructure; any such changes will nullify a deficient test result.

(4) This SLA shall not apply if there is limited or no performance improvement due to causes originating from customer's infrastructure or a third party's infrastructure outside of Akamai's control, including the customer's DNS that provides the CNAME into the Akamai network.

Akamai reserves the right to remove DNS times from the speedup computation if deemed necessary.

B. Availability SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability:

Agents and Polling Frequency

(1) From at least six (6) geographically and network-diverse locations in major metropolitan areas,

(2) Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(3) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com). One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject). The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(4) The Akamaized test content must use a TTL of 2 hours or greater.

(5) The test content will be a file of approximately 10 KB in size.

(6) Polling will occur at approximately 6-minute intervals.

(7) Based on the http GET operations described in II-B above, the response times received from the two sources, (a) the Customer server (directly), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outage

An availability outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly).

IV. SLA Activation

A. Performance SLA Activation

To activate this SLA, Customer and Akamai will designate the applicable page as contemplated by II-A (1) above. Customer shall be responsible for configuring such page to enable testing by the measurement agents. Customer shall be responsible for configuring the Site Analyzer service for measurement of the designated page.

B. Availability SLA Activation

In order to activate the Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(B) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. SLA Escalation

In order to request a credit for a perceived service failure, Customer must, within five calendar days (120 hours) after the perceived failure, contact Akamai Customer Care in writing (which may be via email (ccare@akamai.com) or through the EdgeControl Management Center), specifying the time period in which the failure is believed to have occurred.

VI. Remedies

If the Service fails to meet the service levels in Section I or if an Outage, as defined in Section III, is identified, the Customer will receive (as its sole remedy) a credit equal to Customer's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

2.6 AKAMAI ION STANDARD NORTH AMERICA DOMESTIC SERVICE LEVEL AGREEMENT (MAY 2014)

I. Service Levels

Akamai agrees to provide a level of service for customers with origin infrastructure hosted in North America demonstrating:

(1) Improvement Over Origin: The daily average page delivery time using the Ion Standard service will be at least 30% faster than the daily average delivery time for the same page delivered from the customer's origin server (defined as the daily average page delivery time from origin being at least 1.3 times the daily average page delivery time using Akamai's Ion Standard service).

(2) 100% Availability: The Service will serve content 100% of the time.

II. SLA Monitoring Methodology

A. Performance SLA Monitoring Methodology

The following methodology will be employed to measure the performance improvement mprovided by the Service:

(1) A single static unauthenticated page (including HTML and associated embedded content), provided by Akamai, will be tested throughout the term of the Ion Standard Order Form for purposes of this SLA.

(2) Delivery times will be tested using Akamai's Site Analyzer testing services to measure the daily average page delivery time computed from the complete global set of available Site Analyzer measurement agents. The tests will be configured to take place once an hour from the measurement agents.

(3) This SLA assumes that there will be no material changes to the test content including, without limitation, metadata applied to the content, agent network used for testing, origin settings and origin infrastructure; any such changes will nullify a deficient test result.

(4) This SLA shall not apply if there is limited or no performance improvement due to causes originating from customer's infrastructure or a third party's infrastructure outside of Akamai's control, including the customer's DNS that provides the CNAME into the Akamai network.

Akamai reserves the right to remove DNS times from the speedup computation if deemed necessary.

B. Availability SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability:

Agents and Polling Frequency

(1) From at least six (6) geographically and network-diverse locations in major metropolitan areas,

(2) Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(3) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com). One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject). The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(4) The Akamaized test content must use a TTL of 2 hours or greater.

(5) The test content will be a file of approximately 10 KB in size.

(6) Polling will occur at approximately 6-minute intervals.

(7) Based on the http GET operations described in II-B above, the response times received from the two sources, (a) the Customer server (directly), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outage

An availability outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly).

IV. SLA Activation

A. Performance SLA Activation

To activate this SLA, Customer and Akamai will designate the applicable page as contemplated by II-A (1) above. Customer shall be responsible for configuring such page to enable testing by the measurement agents. Customer shall be responsible for configuring the Site Analyzer service for measurement of the designated page.

B. Availability SLA Activation

In order to activate the Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(B) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. SLA Escalation

In order to request a credit for a perceived service failure, Customer must, within five calendar days (120 hours) after the perceived failure, contact Akamai Customer Care in writing (which may be via email (ccare@akamai.com) or through the EdgeControl Management Center), specifying the time period in which the failure is believed to have occurred.

VI. Remedies

If the Service fails to meet the service levels in Section I or if an Outage, as defined in Section III, is identified, the Customer will receive (as its sole remedy) a credit equal to Customer's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

2.7 PROGRESSIVE MEDIA DOWNLOADS SERVICE LEVEL AGREEMENT (02/11/2014)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

(a) 100% Uptime: The Service will have no Outages (as defined herein below).

(b) Measurable Performance Enhancement: The Service will deliver content faster than the Customer's web site.

(c) Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject), or via an Akamai SiteShield region if the customer is using the SiteShield feature.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a SiteShield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a SiteShield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its

sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a SiteShield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate the Progressive Media Downloads Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and, for customers using SiteShield, Akamai Professional Services. Customers using Remote SiteShield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the SiteShield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

2.8 AKAMAI MEDIA CONTENT DELIVERY SERVICE LEVEL AGREEMENT (01/09/2014)

Applies to Adaptive Media Delivery, Download Delivery, Object Delivery, and Dynamic Site Delivery products.

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

(a) 100% Uptime: The Service will serve content 100% of the time.

(b) Measurable Performance Enhancement: The Service will deliver content measurably faster than the Customer's web site.

(c) Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

(a) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(b) The polling mechanism will perform two (2) simultaneous http GET operations:

A test file will be placed on the customer's origin server (ie, origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (ie, http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Site Shield feature.

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (ie, http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(c) The Akamaized test content must use a TTL of 2 hours or greater.

(d) The test content will be a file of approximately 10 KB in size.

(e) Polling will occur at approximately 6-minute intervals.

(f) Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a Site Shield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate this Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager and for customers using Akamai Professional Services. Customers using Remote Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

2.9 AKAMAI DYNAMIC SITE ACCELERATOR (DSA) SERVICE AND WEB APPLICATION ACCELERATOR (WAA) NORTH AMERICA DOMESTIC SERVICE LEVEL AGREEMENT (10/31/2007)

I. Service Levels

Akamai agrees to provide a level of service for customers with origin infrastructure hosted in North America demonstrating:

(1) Improvement Over Origin: The daily average page delivery time using the WAA or DSA service will be at least 20% faster than the daily average delivery time for the same page delivered from the customer's origin server (defined as the daily average page delivery time from origin being at least 1.2 times the daily average page delivery time using the WAA or DSA service).
(2) 100% Availability: The Service will serve content 100% of the time.

II. SLA Monitoring Methodology

A. Performance SLA Monitoring Methodology

The following methodology will be employed to measure the performance improvement provided by the Service:

(1) A single static unauthenticated page (including HTML and associated embedded content), selected and agreed upon by Akamai and Customer, will be tested throughout the term of the WAA or DSA Order Form for purposes of this SLA. Akamai will make a sample static page (including HTML and associated embedded content) available for use by customers.

(2) Delivery times will be tested using Akamai's Site Analyzer testing services to measure the daily average page delivery time computed from the complete North American set of available Site Analyzer measurement agents. The tests will be configured to take place once an hour from the measurement agents.

(3) This SLA assumes that there will be no material changes to the test content including, without limitation, metadata applied to the content, agent network used for testing, origin settings and origin infrastructure; any such changes will nullify a deficient test result.

(4) This SLA shall not apply if there is limited or no performance improvement due to causes originating from customer's infrastructure or a third party's infrastructure outside of Akamai's control, including the customer's DNS that provides the CNAME into the Akamai network. Akamai reserves the right to remove DNS times from the speedup computation if deemed necessary.

B. Availability SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability:

Agents and Polling Frequency

(1) From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

(2) The polling mechanism will perform two (2) simultaneous http GET operations:

(3) A test file will be placed on the customer's origin server (i.e., origin.customer.com).

One GET operation will be performed to retrieve the file directly from the origin server (i.e., http://origin.customer.com/testobject).

The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (i.e., http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

(4) The Akamaized test content must use a TTL of 2 hours or greater.

(5) The test content will be a file of approximately 10 KB in size.

(6) Polling will occur at approximately 6-minute intervals.

(7) Based on the http GET operations described in II-B above, the response times received from the two sources, (a) the Customer server (directly), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outage

An availability outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly).

IV. SLA Activation

A. Performance SLA Activation

To activate this SLA, Customer and Akamai will designate the applicable page as contemplated by II-A (1) above. Customer shall be responsible for configuring such page to enable testing by the measurement agents. Customer shall be responsible for configuring the Site Analyzer service for measurement of the designated page.

B. Availability SLA Activation

In order to activate the Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(B) above) into the SLA Activation Tool located in the Provisioning Center on https://control.akamai.com (Akamai's Customer Portal). Detailed instructions are provided with the SLA Activation Tool on

https://control.akamai.com; in addition, assistance is available from the Customer's Account Manager. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. SLA Escalation

In order to request a credit for a perceived service failure, Customer must, within five calendar days (120 hours) after the perceived failure, contact Akamai Customer Care in writing (which may be via email (ccare@akamai.com) or through the EdgeControl Management Center), specifying the time period in which the failure is believed to have occurred.

VI. Remedies

If the Service fails to meet the service levels in Section I or an Outage, as defined in Section III, is identified, the Customer will receive (as its sole remedy) a credit equal to Customer's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

2.10 KSD CLIENT REPUTATION OPTION SERVICE LEVEL AGREEMENT (02/08/2016)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

a. 100% Uptime: The Service will serve content 100% of the time.

b. Measurable Performance Enhancement: The Service will deliver content measurably faster than the Customer's web site.

c. Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

a. From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

b. The polling mechanism will perform two (2) simultaneous http GET operations:

c. A test file will be placed on the customer's origin server (i.e., origin.customer.com).

d. One GET operation will be performed to retrieve the file directly from the origin server (i.e., http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Site Shield feature.

   The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (i.e., http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

e. The Akamaized test content must use a TTL of 2 hours or greater.

f. The test content will be a file of approximately 10 KB in size.

g. Polling will occur at approximately 6-minute intervals.

h. Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Performance Metric

The performance metric will be based on a daily average of performance for the Service and the Customer's production web server (measured directly, or via a Site Shield region if applicable), computed from data captured across all regions and hits. If on a given day the Akamai daily average time exceeds the Customer's daily average time, then the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for that day in which the failure occurs, not to exceed 30 days of fees.

IV. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

V. SLA Activation

In order to activate the Kona Site Defender Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) using the Provision Test option found with the Configure menu under SLA Management within the LUNA portal. Detailed instructions regarding the setup of SLA Test Objects are found in the Support menu of LUNA; in addition, assistance is available from the Customer's Account Manager and, for customers using Site Shield, Akamai Professional Services. Customers using Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

VI. Governing SLA

This SLA shall apply solely to delivery configurations using the Kona Site Defender service (noted by "Site_Defender::Site_Defender" in the cpcode within the contract). For delivery configurations using other services, the SLA for those services shall apply and supersede this SLA.

The Web Application Firewall option does not have an SLA. Instead, the SLA for the base delivery solution to which WAF is added is the applicable SLA.

The Client Reputation option does not have an SLA. Instead, the SLA for the base delivery solution to which Client Reputation is added is the applicable SLA.

2.11 BOT MANAGER SERVICE LEVEL AGREEMENT (02/08/2016)

I. Service Levels and Penalties

Akamai agrees to provide a level of service demonstrating:

a. Availability SLA: Akamai offers a service level ("Service Level") committing to 100% availability of the contracted security service.

b. Credits: If the Service fails to meet the above service levels, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurs, not to exceed 30 days of fees.

II. SLA Monitoring Methodology

The following methodology will be employed to measure the Service availability and performance enhancement:

Agents and Polling Frequency

a. From at least six (6) geographically and network-diverse locations in major metropolitan areas, Akamai will simultaneously poll a test file residing on the Customer's production servers and on Akamai's network

b. The polling mechanism will perform two (2) simultaneous http GET operations:

c. A test file will be placed on the customer's origin server (i.e., origin.customer.com).

d. One GET operation will be performed to retrieve the file directly from the origin server (i.e., http://origin.customer.com/testobject), or via an Akamai Site Shield region if the customer is using the Site Shield feature.

   The other GET operation will be performed to retrieve the file through the Service, by requesting the object from the appropriate customer hostname CNAMEd to Akamai (i.e., http://www.customer.com/testobject, where www.customer.com is CNAMEd to Akamai and configured to pull content from origin.customer.com)

e. The Akamaized test content must use a TTL of 2 hours or greater.

f. The test content will be a file of approximately 10 KB in size.

g. Polling will occur at approximately 6-minute intervals.

h. Based on the http GET operations described in II(b) above, the response times received from the two sources, (a) the Customer server (directly, or via a Site Shield region if applicable), and (b) the Akamai network, will be compared for the purpose of measuring performance metrics and outages.

III. Outages

An outage is defined as a period of at least two consecutive failed attempts six minutes apart by a single agent to GET the Customer test file from the Service while succeeding to GET the test file from the Customer Origin Server (directly, or via a Site Shield region if applicable). If an outage is identified by this method, the Customer will receive (as its sole remedy) a credit equal to Customer's or such domain's committed monthly service fee for the day in which the failure occurred, not to exceed 30 days of fees.

IV. SLA Activation

In order to activate the Kona Site Defender Service Level Agreement, the Customer must enter and indicate the location of two valid test files for the same object (as described in II(c) and (d) above) using the Provision Test option found with the Configure menu under SLA Management within the LUNA portal. Detailed instructions regarding the setup of SLA Test Objects are found in the Support menu of LUNA; in addition, assistance is available from the Customer's Account Manager and, for customers using Site Shield, Akamai Professional Services. Customers using Site Shield must ensure that their firewall configurations are updated to reflect changes made by Akamai to the Site Shield access control list no later than 60 days following notification by Akamai, via email or the https://control.akamai.com customer portal, of such a change. Failure to timely update such firewall configurations will invalidate the performance portion of this SLA. The SLA will go into effect within five business days after the Customer enters valid test files into the SLA Activation Tool.

V. Governing SLA

In order for Akamai to issue a credit in accordance with this SLA, Customer must have an account that is current with payments and in good standing with Akamai, and must be able to confirm that Customer has completed the Integration process for the applicable Service and, if applicable, all Competing Mitigation Hardware and Techniques have been disabled or removed during any mitigation services.

Credits shall only apply for Services provided pursuant to the Monthly Service Fee and/or Monthly Service Overage Fee, and will not apply to any other Service. The aggregate credits to be provided in any calendar month shall not exceed 25% of the Monthly Service Fee in respect of the affected Service(s).

This SLA shall apply solely to delivery configurations using the Kona Site Defender service (noted by "Bot Manager::Bot Manager" in the cpcode within the contract). For delivery configurations using other services, the SLA for those services shall apply and supersede this SLA.

3.0 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 Update 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).

Google Terms of Service

Last modified: April 14, 2014 (view archived versions)

Welcome to Google!

Thanks for using our products and services ("Services"). The Services are provided by Google Inc. ("Google"), located at 1600 Amphitheatre Parkway, Mountain View, CA 94043, United States.

By using our Services, you are agreeing to these terms. Please read them carefully.

Our Services are very diverse, so sometimes additional terms or product requirements (including age requirements) may apply. Additional terms will be available with the relevant Services, and those additional terms become part of your agreement with us if you use those Services.

Using our Services

You must follow any policies made available to you within the Services.

Don't misuse our Services. For example, don't interfere with our Services or try to access them using a method other than the interface and the instructions that we provide. You may use our Services only as permitted by law, including applicable export and re-export control laws and regulations. We may suspend or stop providing our Services to you if you do not comply with our terms or policies or if we are investigating suspected misconduct.

Using our Services does not give you ownership of any intellectual property rights in our Services or the content you access. You may not use content from our Services unless you obtain permission from its owner or are otherwise permitted by law. These terms do not grant you the right to use any branding or logos used in our Services. Don't remove, obscure, or alter any legal notices displayed in or along with our Services.

Our Services display some content that is not Google's. This content is the sole responsibility of the entity that makes it available. We may review content to determine whether it is illegal or violates our policies, and we may remove or refuse to display content that we reasonably believe violates our policies or the law. But that does not necessarily mean that we review content, so please don't assume that we do.

In connection with your use of the Services, we may send you service announcements, administrative messages, and other information. You may opt out of some of those communications.

Some of our Services are available on mobile devices. Do not use such Services in a way that distracts you and prevents you from obeying traffic or safety laws.

Your Google Account

You may need a Google Account in order to use some of our Services. You may create your own Google Account, or your Google Account may be assigned to you by an administrator, such as your employer or educational institution. If you are using a Google Account assigned to you by an administrator, different or additional terms may apply and your administrator may be able to access or disable your account.

To protect your Google Account, keep your password confidential. You are responsible for the activity that happens on or through your Google Account. Try not to reuse your Google Account password on third-party applications. If you learn of any unauthorized use of your password or Google Account, follow these instructions.

Privacy and Copyright Protection

Google's privacy policies explain how we treat your personal data and protect your privacy when you use our Services. By using our Services, you agree that Google can use such data in accordance with our privacy policies.

We respond to notices of alleged copyright infringement and terminate accounts of repeat infringers according to the process set out in the U.S. Digital Millennium Copyright Act.

We provide information to help copyright holders manage their intellectual property online. If you think somebody is violating your copyrights and want to notify us, you can find information about submitting notices and Google's policy about responding to notices in our Help Center.

Your Content in our Services

Some of our Services allow you to upload, submit, store, send or receive content. You retain ownership of any intellectual property rights that you hold in that content. In short, what belongs to you stays yours.

When you upload, submit, store, send or receive content to or through our Services, you give Google (and those we work with) a worldwide license to use, host, store, reproduce, modify, create derivative works (such as those resulting from translations, adaptations or other changes we make so that your content works better with our Services), communicate, publish, publicly perform, publicly display and distribute such content. The rights you grant in this license are for the limited purpose of operating, promoting, and improving our Services, and to develop new ones. This license continues even if you stop using our Services (for example, for a business listing you have added to Google Maps). Some Services may offer you ways to access and remove content that has been provided to that Service. Also, in some of our Services, there are terms or settings that narrow the scope of our use of the content submitted in those Services. Make sure you have the necessary rights to grant us this license for any content that you submit to our Services.

Our automated systems analyze your content (including emails) to provide you personally relevant product features, such as customized search results, tailored advertising, and spam and malware detection. This analysis occurs as the content is sent, received, and when it is stored.

If you have a Google Account, we may display your Profile name, Profile photo, and actions you take on Google or on third-party applications connected to your Google Account (such as +1's, reviews you write and comments you post) in our Services, including displaying in ads and other commercial contexts. We will respect the choices you make to limit sharing or visibility settings in your Google Account. For example, you can choose your settings so your name and photo do not appear in an ad.

You can find more information about how Google uses and stores content in the privacy policy or additional terms for particular Services. If you submit feedback or suggestions about our Services, we may use your feedback or suggestions without obligation to you.

About Software in our Services

When a Service requires or includes downloadable software, this software may update automatically on your device once a new version or feature is available. Some Services may let you adjust your automatic update settings.

Google gives you a personal, worldwide, royalty-free, non-assignable and non-exclusive license to use the software provided to you by Google as part of the Services. This license is for the sole purpose of enabling you to use and enjoy the benefit of the Services as provided by Google, in the manner permitted by these terms. You may not copy, modify, distribute, sell, or lease any part of our Services or included software, nor may you reverse engineer or attempt to extract the source code of that software, unless laws prohibit those restrictions or you have our written permission.

Open source software is important to us. Some software used in our Services may be offered under an open source license that we will make available to you. There may be provisions in the open source license that expressly override some of these terms.

Modifying and Terminating our Services

We are constantly changing and improving our Services. We may add or remove functionalities or features, and we may suspend or stop a Service altogether.

You can stop using our Services at any time, although we'll be sorry to see you go. Google may also stop providing Services to you, or add or create new limits to our Services at any time.

We believe that you own your data and preserving your access to such data is important. If we discontinue a Service, where reasonably possible, we will give you reasonable advance notice and a chance to get information out of that Service.

Our Warranties and Disclaimers

We provide our Services using a commercially reasonable level of skill and care and we hope that you will enjoy using them. But there are certain things that we don't promise about our Services.

OTHER THAN AS EXPRESSLY SET OUT IN THESE TERMS OR ADDITIONAL TERMS, NEITHER GOOGLE NOR ITS SUPPLIERS OR DISTRIBUTORS MAKE ANY SPECIFIC PROMISES ABOUT THE SERVICES. FOR EXAMPLE, WE DON'T MAKE ANY COMMITMENTS ABOUT THE CONTENT WITHIN THE SERVICES, THE SPECIFIC FUNCTIONS OF THE SERVICES, OR THEIR RELIABILITY, AVAILABILITY, OR ABILITY TO MEET YOUR NEEDS. WE PROVIDE THE SERVICES "AS IS".

SOME JURISDICTIONS PROVIDE FOR CERTAIN WARRANTIES, LIKE THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. TO THE EXTENT PERMITTED BY LAW, WE EXCLUDE ALL WARRANTIES.

Liability for our Services

WHEN PERMITTED BY LAW, GOOGLE, AND GOOGLE'S SUPPLIERS AND DISTRIBUTORS, WILL NOT BE RESPONSIBLE FOR LOST PROFITS, REVENUES, OR DATA, FINANCIAL LOSSES OR INDIRECT, SPECIAL, CONSEQUENTIAL, EXEMPLARY, OR PUNITIVE DAMAGES.

TO THE EXTENT PERMITTED BY LAW, THE TOTAL LIABILITY OF GOOGLE, AND ITS SUPPLIERS AND DISTRIBUTORS, FOR ANY CLAIMS UNDER THESE TERMS, INCLUDING FOR ANY IMPLIED WARRANTIES, IS LIMITED TO THE AMOUNT YOU PAID US TO USE THE SERVICES (OR, IF WE CHOOSE, TO SUPPLYING YOU THE SERVICES AGAIN).

IN ALL CASES, GOOGLE, AND ITS SUPPLIERS AND DISTRIBUTORS, WILL NOT BE LIABLE FOR ANY LOSS OR DAMAGE THAT IS NOT REASONABLY FORESEEABLE.

Business uses of our Services

If you are using our Services on behalf of a business, that business accepts these terms. It will hold harmless and indemnify Google and its affiliates, officers, agents, and employees from any claim, suit or action arising from or related to the use of the Services or violation of these terms, including any liability or expense arising from claims, losses, damages, suits, judgments, litigation costs and attorneys' fees.

About these Terms

We may modify these terms or any additional terms that apply to a Service to, for example, reflect changes to the law or changes to our Services. You should look at the terms regularly. We'll post notice of modifications to these terms on this page. We'll post notice of modified additional terms in the applicable Service. Changes will not apply retroactively and will become effective no sooner than fourteen days after they are posted. However, changes addressing new functions for a Service or changes made for legal reasons will be effective immediately. If you do not agree to the modified terms for a Service, you should discontinue your use of that Service.

If there is a conflict between these terms and the additional terms, the additional terms will control for that conflict.

These terms control the relationship between Google and you. They do not create any third party beneficiary rights.

If you do not comply with these terms, and we don't take action right away, this doesn't mean that we are giving up any rights that we may have (such as taking action in the future).

If it turns out that a particular term is not enforceable, this will not affect any other terms.

The laws of California, U.S.A., excluding California's conflict of laws rules, will apply to any disputes arising out of or relating to these terms or the Services. All claims arising out of or relating to these terms or the Services will be litigated exclusively in the federal or state courts of Santa Clara County, California, USA, and you and Google consent to personal jurisdiction in those courts.

For information about how to contact Google, please visit our contact page.

Data Processing Amendment

The Analytics Data Processing Amendment is meant for businesses established in the territory of a member state of the European Economic Area or Switzerland, or that, for other reasons, are subject to the territorial scope of the national implementations of Directive 95/46/EC.

Access to the Data Processing Amendment is available in the Account Settings in the Administration section. Users with edit permission for an Analytics account can review the Data Processing Amendment and accept it if they so choose. Once accepted, it is possible to review the contents of the accepted amendment.

Acceptance of the Data Processing Amendment is optional and does not affect the functionality of Analytics in any way.

Google Apps

Google Apps for Business (Online) Agreement

Go to the Additional Terms for services made available with the new accounts infrastructure

*The terms below are for monthly postpay billing. Please click here to see the terms for annual prepay customers.

This Google Apps for Business (Online) Agreement (the "Agreement") is entered into by and between Google Inc., a Delaware corporation, with offices at 1600 Amphitheatre Parkway, Mountain View, California 94043 ("Google") and the entity agreeing to these terms ("Customer"). This Agreement is effective as of the date you click the "I Accept" button below or, if applicable, the date the Agreement is countersigned (the "**Effective Date**"). If you are accepting on behalf of your employer or another entity, you represent and warrant that: (i) you have full legal authority to bind your employer, or the applicable entity, to these terms and conditions; (ii) you have read and understand this Agreement; and (iii) you agree, on behalf of the party that you represent, to this Agreement. If you don''t have the legal authority to bind your employer or the applicable entity, please do not click the "I Accept" button below (or, if applicable, do not sign this Agreement). This Agreement governs Customer''s access to and use of the Services.

1. **Services**

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 Google stores and processes its own information of a similar type. Google 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 Services Google may transfer store and process Customer Data in the United States or any other country in which Google or its agents maintain facilities. By using the Services Customer consents to this transfer, processing and storage of Customer Data.

1.2 **Modifications**

a. To the Services. Google may make commercially reasonable changes to the Services from time to time. If Google makes a material change to the Services Google will inform Customer, provided that Customer has subscribed with Google to be informed about such change.

b. **To URL Terms.** Google may make commercially reasonable changes to the URL Terms from time to time. If Google makes a material change to the URL Terms, Google will inform Customer by either sending an email to the Notification Email Address or alerting Customer 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 Google via the Help Center within thirty days after receiving notice of the change. If Customer notifies Google as required, then Customer will remain governed by the terms in effect immediately prior to the change until the end of the then-current Services Term for the affected Services. If the affected Services are renewed, they will be renewed under Google's then current URL Terms.

1.3 **Customer Domain Name Ownership.** Prior to providing the Services Google may verify that Customer owns or controls the Customer Domain Names. If Customer does not own or control the Customer Domain Names, then Google will have no obligation to provide Customer with the Services.

1.4 **Ads.** The default setting for the Services is one that does not allow Google to serve Ads. Customer may change this setting in the Admin Console which constitutes Customer's authorization for Google to serve Ads. If Customer enables the serving of Ads it may revert to the default setting at any time and Google will cease serving Ads.

1.5 **Google Apps Vault.** If Customer purchases Google Apps Vault, the following additional terms apply:

a. Retention. Google will have no obligation to retain any archived Customer Data beyond the retention period specified by Customer (other than for any legal holds). If Customer does not renew Google Apps Vault, Google will have no obligation to retain any archived Customer Data.

b. **Additional Purchases.** Unless Google allows otherwise, with each additional purchase of End User Accounts for the Services after Customer has purchased Google Apps Vault, Customer will receive access to, and will be invoiced for, Google Apps Vault for that same number of End User Accounts.

2. **Customer Obligations.**

2.1 **Compliance.** Customer will use the Services in accordance with the Acceptable Use Policy. Google may make new applications features or functionality for the Services available from time to time the use of which may be contingent upon Customer's agreement to additional terms. In addition, Google will make other Non-Google Apps Products (beyond the Services) available to Customer and its End Users in accordance with the Non-Google Apps Product Terms and the applicable product-specific Google 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 **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 Google may monitor emails sent to these aliases for Customer Domain Names to allow Google to identify Services abuse.

2.3 **Customer Administration of the Services.** 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 is 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 the Agreement. Customer agrees that Google's responsibilities do not extend to the internal management or administration of the Services for Customer and that Google is merely a data-processor.

2.4 **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 Google providing Customer with the ability to do so and (ii) Google to provide the Services.

2.5 **Unauthorized Use.** Customer will use commercially reasonable efforts to prevent unauthorized use of the Services' and to terminate any unauthorized use. Customer will promptly notify Google of any unauthorized use of, or access to, the Services of which it becomes aware.

2.6 **Restrictions on Use.** Unless Google 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 Services to a third party (unless expressly authorized in this Agreement); (b) attempt to reverse engineer the Services or any component; (c) attempt to create a substitute or similar service through use of, or access to, the Services; (d) use the Services for High Risk Activities; or (e) use the Services 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.

2.7 **Third Party Requests.** Customer is responsible for responding to Third Party Requests. Google 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 Google only if it cannot reasonably obtain such information.

3. **Billing and Payment.**

3.1 **Billing**Customer may elect one of the following billing options when placing its order for the Services.

a. Flexible Plan. If Customer selects this option, Customer will not be committed to purchase the Services for a pre-defined term, but will pay for the Services on a monthly basis. Google will bill Customer: (i) Fees based upon Customer's daily usage of the Services during the preceding month; and (ii) monthly in arrears for its use of the Services. Google will provide Customer with the monthly rate for the Services when Customer orders the Services, and will use this rate to calculate the Fees, on a prorated basis, for Customer's daily usage during that month. Any partial day of Services usages will be rounded up to a full day of Services usage for the purposes of calculating Fees. Customer may pay for the Services using the payment options listed below.

b. Annual Plan If Customer selects this option, Customer will be committed to purchasing the Services from Google for an annual term, and in exchange will receive a discount on the Services which will be reflected in Customer's monthly payment. Google will still bill Customer monthly in arrears for its use of the Services when Customer has an annual commitment for the Services with Google. Customer may pay for the Services using the payment options listed below.

3.2 **Payment.** All payments due are in U.S. dollars unless otherwise indicated on the Order Page or invoice.

a. Credit Card or Debit Card. Fees for orders where Customer is paying with a credit card, debit card or other non-invoice form of payment, are due at the end of the month during which Customer received the Services. For credit cards, or debit cards, as applicable: (i) Google will charge Customer for all applicable Fees when due and (ii) these Fees are considered delinquent thirty days after the end of the month during which Customer received the Services.

b. Invoices. Payments for invoices are due thirty days after the invoice date, unless otherwise specified on the Order Page, and are considered delinquent after such date.

c. Other Forms of Payment. Customer may change its payment method to those available within the Admin Console. Google may enable other forms of payment by making them available in the Admin Console. These other forms of payment may be subject to additional terms which Customer may have to accept prior using the additional forms of payment.

3.3 **Delinquent Payments** Delinquent payments may bear interest at the rate of one-and-one-half percent per month (or the highest rate permitted by law, if less) from the payment due date until paid in full. Customer will be responsible for all reasonable expenses (including attorneys' fees) incurred by Google in collecting such delinquent amounts, except where such delinquent amounts are due to Google's billing inaccuracies.

3.4 **Suspension for Non-Payment.**

a. Automatic Suspension. Customer will have thirty days to pay Google delinquent Fees. If Customer does not pay Google delinquent Fees within thirty days, Google will automatically suspend Customer's use of the Services. The duration of this suspension will be until Customer pays Google all outstanding Fees.

b. During Suspension. If Customer is on a monthly billing plan, and Customer is suspended for non-payment, Google will stop charging Customer monthly Fees during Customer's suspension for non-payment. If Customer has an annual commitment to Google for the Services, Google will continue to charge Customer monthly Fees during Customer's suspension for non-payment and Customer must pay all outstanding Fees in order to resume its use of the Services.

c. Termination After Suspension. If Customer remains suspended for non-payment for more than sixty days, Google may terminate Customer for breach pursuant to Section 11.

3.5 **Taxes.** Customer is responsible for any Taxes, and Customer will pay Google for the Services without any reduction for Taxes. If Google is obligated to collect or pay Taxes, the Taxes will be invoiced to Customer, unless Customer provides Google with a valid tax exemption certificate authorized by the appropriate taxing authority. If Customer is required by law to withhold any Taxes from its payments to Google, Customer must provide Google with an official tax receipt or other appropriate documentation to support such payments.

3.6 **Purchase Orders.** If Customer requires a purchase order number on its invoice, Customer will inform Google and Google will include such purchase order number on invoices following receipt. If Customer does not provide a purchase order number, Customer waives any purchase order requirement and (a) Google will invoice Customer without a purchase order number; and (b) Customer agrees to pay invoices without a purchase order number referenced. Any terms and conditions on a purchase order do not apply to this Agreement and are null and void.

4. **Technical Support Services.**

4.1 **By Customer.** Customer 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 Services. Customer will use commercially reasonable efforts to resolve support issues before escalating them to Google.

4.2 **By Google.** If Customer cannot resolve a support issue consistent with the above, then Customer may escalate the issue to Google in accordance with the TSS Guidelines. Google will provide TSS to Customer in accordance with the TSS Guidelines.

5. **Suspension**

5.1 **Of End User Accounts by Google.** If Google becomes aware of an End User's violation of the Agreement, then Google may specifically request that Customer Suspend the applicable End User Account. If Customer fails to comply with Google's request to Suspend an End User Account, then Google may do so. The duration of any Suspension by Google will be until the applicable End User has cured the breach which caused the Suspension.

5.2 **Emergency Security Issues.** Notwithstanding the foregoing, if there is an Emergency Security Issue, then Google 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 Google Suspends an End User Account for any reason without prior notice to Customer, at Customer's request, Google will provide Customer the reason for the Suspension as soon as is reasonably possible.

6. **Confidential Information.**

6.1 **Obligations.** Each party will: (a) protect the other party'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. Each party (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 this Agreement, while using reasonable care to protect it. Each party is responsible for any actions of its Affiliates' employees and agents in violation of this Section.

6.2 **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 another party.

6.3 **Required Disclosure.** Each party may disclose the other party's Confidential Information when required by law but only after it, if legally permissible: (a) uses commercially reasonable efforts to notify the other party; and (b) gives the other party the chance to challenge the disclosure.

7. **Intellectual Property Rights; Brand Features.**

7.1 **Intellectual Property Rights.** Except as expressly set forth herein, this Agreement does not grant either party any rights, implied or otherwise, to the other's content or any of the other's intellectual property. As between the parties, Customer owns all Intellectual Property Rights in Customer Data, and Google owns all Intellectual Property Rights in the Services.

7.2 **Display of Brand Features.** Google may display those Customer Brand Features authorized by Customer (such authorization is provided by Customer uploading its Brand Features into the Services) within designated areas of the Service Pages. Customer may specify the nature of this use using the Admin Console. Google may also display Google Brand Features on the Service Pages to indicate that the Services are provided by Google. Neither party may display or use the other party’s Brand Features beyond what is allowed in this Agreement without the other party’s prior written consent.

7.3 **Brand Features Limitation.** Any use of a party's Brand Features will inure to the benefit of the party holding Intellectual Property Rights in those Brand Features. A party may revoke the other party's right to use its Brand Features pursuant to this Agreement with written notice to the other and a reasonable period to stop the use.

8. **Publicity**.

Customer agrees that Google may include Customer's name or Brand Features in a list of Google customers, online or in promotional materials. Customer also agrees that Google may verbally reference Customer as a

customer of the Google products or services that are the subject of this Agreement. This section is subject to Section 7.3 (Brand Features Limitation).

9. **Representations, Warranties and Disclaimers.**

9.1 **Representations and Warranties.** Each party represents that it has full power and authority to enter into the Agreement. Each party warrants that it will comply with all laws and regulations applicable to its provision, or use, of the Services, as applicable (including applicable security breach notification law). Google warrants that it will provide the Services in accordance with the applicable SLA.

9.2 **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. GOOGLE MAKES NO REPRESENTATIONS ABOUT ANY CONTENT OR INFORMATION MADE ACCESSIBLE BY OR THROUGH THE SERVICES. CUSTOMER ACKNOWLEDGES THAT THE SERVICES ARE NOT A TELEPHONY SERVICE AND THAT THE SERVICES ARE NOT CAPABLE OF PLACING OR RECEIVING ANY CALLS, INCLUDING EMERGENCY SERVICES CALLS, OVER PUBLICLY SWITCHED TELEPHONE NETWORKS.

10. **Term**

10.1 **Agreement Term.** This Agreement will remain in effect for the Term.

10.2 **Services Term and Purchases During Services Term.** Google will provide the Services to Customer during the Services Term. Unless the parties agree otherwise in writing, End User Accounts purchased during any Services Term will have a prorated term ending on the last day of that Services Term.

10.3 **Renewal.**

a. With a Flexible Plan. With a flexible plan Customer is not committed to purchase the Services for a pre-defined term, but pays for the Services on a monthly basis. As a result, there is no renewal event for the flexible plan. Rather, Google will simply continuing billing Customer Fees based upon Customer's daily usage of the Services during the preceding month, and Customer can cancel their service at any time.

b. With an Annual Plan. At the end of each Services Term, the Services (and all End User Accounts previously purchased) will automatically renew for an additional monthly Services Term. In addition, after Customer's initial annual commitment has concluded, Customer's annual commitment will switch to the Flexible Plan. If Customer wants to renew the Annual Plan, then Customer must change the renewal settings in the Admin Console to reflect this change before their annual commitment has ended.

c. Generally. Customer may alter the number of End User Accounts to be renewed by communicating the appropriate number of accounts to be renewed to Google via the Admin Console. Customer will continue to pay Google the then-current Fees for each renewed End User Account unless Customer and Google mutually agree otherwise. If Google does not want the Services to renew, then it will provide Customer written notice to this effect at least fifteen days prior to the end of the then current Services Term. This notice of non renewal will be effective upon the conclusion of the then current Services Term.

10.4 **Requesting End User Accounts.** Customer may request End User Accounts by: (i) notifying its designated Google Account Manager; or (ii) ordering End User Accounts via the Admin Console.

10.5 **Revising Rates.** Google may revise its rates for the following Services Term by providing Customer written notice (which may be by email) at least thirty days prior to the start of the following Services Term.

11. **Termination.**

11.1 **Termination for Breach.** Either party may suspend performance or terminate this Agreement if: (i) the other party is in material breach of the Agreement and fails to cure that breach within thirty days after receipt of written notice; (ii) the other party ceases its business operations or becomes subject to insolvency proceedings and the proceedings are not dismissed within ninety days; or (iii) the other party is in material breach of this Agreement more than two times notwithstanding any cure of such breaches.

11.2 **Effects of Termination.** If this Agreement terminates, then: (i) the rights granted by one party to the other will cease immediately (except as set forth in this Section); (ii) Google will provide Customer access to, and the ability to export, the Customer Data for a commercially reasonable period of time at Google's then-current rates for the applicable Services; (iii) after a commercially reasonable period of time, Google will delete Customer Data by removing pointers to it on Google's active servers and overwriting it over time; and (iv) upon request each party will promptly use commercially reasonable efforts to return or destroy all other Confidential Information of the other party. If a Customer on an annual plan terminates the Agreement prior to the conclusion of its annual plan, Google will bill Customer, and Customer is responsible for paying Google, for the remaining unpaid amount of Customer's annual commitment.

12. **Indemnification.**

12.1 **By Customer.** Customer will indemnify, defend, and hold harmless Google from and against all liabilities, damages, and costs (including settlement costs and reasonable attorneys' fees) arising out of a third party claim: (i) regarding Customer Data or Customer Domain Names; (ii) that Customer Brand Features infringe or misappropriate any patent, copyright, trade secret or trademark of a third party; or (iii) regarding Customer's use of the Services in violation of the Acceptable Use Policy.

12.2 **By Google.** Google will indemnify, defend, and hold harmless Customer from and against all liabilities, damages, and costs (including settlement costs and reasonable attorneys' fees) arising out of a third party claim that Google's technology used to provide the Services or any Google Brand Feature infringe or misappropriate any patent, copyright, trade secret or trademark of such third party. Notwithstanding the foregoing, in no event shall Google have any obligations or liability under this Section arising from: (i) use of any Services or Google Brand Features in a modified form or in combination with materials not furnished by Google, and (ii) any content, information or data provided by Customer, End Users or other third parties.

12.3 **Possible Infringement.**

a. Repair, Replace, or Modify. If Google reasonably believes the Services infringe a third party's Intellectual Property Rights, then Google will: (a) obtain the right for Customer, at Google's expense, to continue using the Services; (b) provide a non-infringing functionally equivalent replacement; or (c) modify the Services so that they no longer infringe.

b. Suspension or Termination. If Google does not believe the foregoing options are commercially reasonable, then Google may suspend or terminate Customer's use of the impacted Services. If Google

terminates the impacted Services, then Google will provide a pro-rata refund of the unearned Fees actually paid by Customer applicable to the period following termination of such Services.

12.4 **General.** The party seeking indemnification will promptly notify the other party of the claim and cooperate with the other party in defending the claim. The indemnifying party has full control and authority over the defense, except that: (a) any settlement requiring the party seeking indemnification to admit liability or to pay any money will require that party's prior written consent, such consent not to be unreasonably withheld or delayed; and (b) the other party may join in the defense with its own counsel at its own expense. THE INDEMNITIES ABOVE ARE A PARTY'S ONLY REMEDY UNDER THIS AGREEMENT FOR VIOLATION BY THE OTHER PARTY OF A THIRD PARTY'S INTELLECTUAL PROPERTY RIGHTS.

13. **Limitation of Liability.**

13.1 **Limitation on Indirect Liability.** NEITHER PARTY WILL BE LIABLE UNDER THIS AGREEMENT FOR LOST REVENUES OR INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL, EXEMPLARY, OR PUNITIVE DAMAGES, EVEN IF THE PARTY KNEW OR SHOULD HAVE KNOWN THAT SUCH DAMAGES WERE POSSIBLE AND EVEN IF DIRECT DAMAGES DO NOT SATISFY A REMEDY.

13.2 **Limitation on Amount of Liability.** NEITHER PARTY MAY BE HELD LIABLE UNDER THIS AGREEMENT FOR MORE THAN THE AMOUNT PAID BY CUSTOMER TO GOOGLE HEREUNDER DURING THE TWELVE MONTHS PRIOR TO THE EVENT GIVING RISE TO LIABILITY.

13.3 **Exceptions to Limitations.** These limitations of liability apply to the fullest extent permitted by applicable law but do not apply to breaches of confidentiality obligations, violations of a party's Intellectual Property Rights by the other party, or indemnification obligations.

14. **Miscellaneous.**

14.1 **Notices.** Unless specified otherwise herein, (a) all notices must be in writing and addressed to the attention of the other party's legal department and primary point of contact and (b) notice will be deemed given: (i) when verified by written receipt if sent by personal courier, overnight courier, or when received if sent by mail without verification of receipt; or (ii) when verified by automated receipt or electronic logs if sent by facsimile or email.

14.2 **Assignment.** Neither party may assign or transfer any part of this Agreement without the written consent of the other party, except to an Affiliate, but only if: (a) the assignee agrees in writing to be bound by the terms of this Agreement; and (b) the assigning party remains liable for obligations incurred under the Agreement prior to the assignment. Any other attempt to transfer or assign is void.

14.3 **Change of Control.** Upon a change of control (for example, through a stock purchase or sale, merger, or other form of corporate transaction): (a) the party experiencing the change of control will provide written notice to the other party within thirty days after the change of control; and (b) the other party may immediately terminate this Agreement any time between the change of control and thirty days after it receives the written notice in subsection (a).

14.4 **Force Majeure.** Neither party will be liable for inadequate performance to the extent caused by a condition (for example, natural disaster, act of war or terrorism, riot, labor condition, governmental action, and Internet disturbance) that was beyond the party's reasonable control.

14.5 **No Waiver.** Failure to enforce any provision of this Agreement will not constitute a waiver.

14.6 **Severability.** If any provision of this Agreement is found unenforceable, the balance of the Agreement will remain in full force and effect.

14.7 **No Agency.** The parties are independent contractors, and this Agreement does not create an agency, partnership or joint venture.

14.8 **No Third-Party Beneficiaries.** There are no third-party beneficiaries to this Agreement.

14.9 **Equitable Relief.** Nothing in this Agreement will limit either party's ability to seek equitable relief.

14.10 **Governing Law.** This Agreement is governed by California law, excluding that state's choice of law rules. FOR ANY DISPUTE ARISING OUT OF OR RELATING TO THIS AGREEMENT, THE PARTIES CONSENT TO PERSONAL JURISDICTION IN, AND THE EXCLUSIVE VENUE OF, THE COURTS IN SANTA CLARA COUNTY, CALIFORNIA.

14.11 **Amendments.** Any amendment must be in writing and expressly state that it is amending this Agreement.

14.12 **Survival.** The following sections will survive expiration or termination of this Agreement: Section 3, 6, 7.1, 11.2, 12, 13, 14, and 15.

14.13 **Entire Agreement.** This Agreement, and all documents referenced herein, is the parties' entire agreement relating to its subject and supersedes any prior or contemporaneous agreements on that subject. The terms located at a URL and referenced in this Agreement are hereby incorporated by this reference.

14.14 **Interpretation of Conflicting Terms.** If there is a conflict between the documents that make up this Agreement, the documents will control in the following order: the Order Page, the Agreement, and the terms located at any URL. If Customer signs a physical agreement with Google to receive the Services, the physical agreement will override this online Agreement.

14.15 **Counterparts.** The parties may enter into this Agreement in counterparts, including facsimile, PDF or other electronic copies, which taken together will constitute one instrument.

15. **Definitions.**

"Acceptable Use Policy" means the acceptable use policy for the Services available at https://www.google.com/a/help/intl/en/admins/use_policy.html or such other URL as Google may provide.

"Account Manager" means the Google business person working with Customer regarding Customer's purchase of the Services.

"Admin Account(s)" means the administrative account(s) provided to Customer by Google for the purpose of administering the Services. The use of the Admin Account(s) requires a password, which Google will provide to Customer.

"Admin Console" means the online tool provided by Google to Customer for use in reporting and certain other administration functions.

"Administrators" mean the Customer-designated technical personnel who administer the Services to End Users on Customer's behalf.

"Ads" means online advertisements displayed by Google to End Users.

"Affiliate" means any entity that directly or indirectly controls, is controlled by, or is under common control with a party.

"Brand Features" means the trade names, trademarks, service marks, logos, domain names, and other distinctive brand features of each party, respectively, as secured by such party from time to time.

"Confidential Information" means information disclosed by a party to the other party under this Agreement 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 Services by Customer or End Users.

"Customer Domain Names" mean the domain names owned or controlled by Customer, which will be used in connection with the Services and specified in the Order Page.

"Emergency Security Issue" means either: (a) Customer's use of the Services in violation of the Acceptable Use Policy, which could disrupt: (i) the Services; (ii) other customer's use of the Services; or (iii) the Google network or servers used to provide the Services; or (b) unauthorized third party access to the Services.

"End Users" means the individuals Customer permits to use the Services.

"End User Account" means a Google-hosted account established by Customer through the Services for an End User.

"Export Control Laws" means all applicable export and reexport control laws and regulations, including the Export Administration Regulations (“EAR”) maintained by the U.S. Department of Commerce, 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.

"Fees" means the amounts invoiced to Customer by Google for the Services as described in an Order Page.

"Help Center" means the Google help center accessible at https://www.google.com/support/, or other such URL as Google may provide.

"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 Services 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.

"Initial Services Term" means the term for the applicable Services beginning on the Service Commencement Date and continuing for the duration set forth on the Order Page.

"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 Google products which are not part of the Services, 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: https://www.google.com/support/a/bin/answer.py?hl=en&answer=181865, or such other URL as Google may provide.

"Non-Google Apps Product Terms" means the terms found at the following URL: https://www.google.com/apps/intl/en/terms/additional_services.html, or such other URL as Google may provide.

"Notification Email Address" means the email address designated by Customer to receive email notifications from Google. Customer may change this email address through the Admin Console.

"Order Page" means the online order page Customer completes in signing up for the Services or attached to this Agreement, and which contains: (i) the Services being ordered; (ii) Fees; (iii) number of, and Initial Services Term for, End User Accounts; (iv) the applicable form of payment; and (v) Customer Domain Names.

"Service Commencement Date" is the date upon which Google makes the Services available to Customer, and will be within one week of Google's receipt of the completed Order Page, unless otherwise agreed by the parties.

"Service Pages" mean the web pages displaying the Services to End Users.

"Services" means the applicable Google Apps Core Services (e.g. Google Apps Premier Edition or Google Apps for Business and Google Apps Vault) provided by Google and used by Customer under this Agreement. The Services are as described here: https://www.google.com/apps/intl/en/terms/user_features.html, or such other URL as Google may provide.

"Services Term" means the Initial Services Term and all renewal terms for the applicable Services.

"SLA" means the Service Level Agreement located here: https://www.google.com/a/help/intl/en/admins/sla.html, or such other URL as Google may provide.

"Suspend" means the immediate disabling of access to the Services, or components of the Services, as applicable, to prevent further use of the Services.

"Taxes" means any duties, customs fees, or taxes (other than Google's income tax) associated with the sale of the Services, including any related penalties or interest.

"Term" means the term of the Agreement, which will begin on the Effective Date and continue until the earlier of (i) the end of the last Services Term or (ii) the Agreement is terminated as set forth herein.

"Third Party Request" means a request from a third party for records relating to an End User's use of the Services. 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 services provided by Google to the Administrators during the Term pursuant to the TSS Guidelines.

"TSS Guidelines" means Google's technical support services guidelines then in effect for the Services. TSS Guidelines are at the following URL: https://www.google.com/a/help/intl/en/admins/tssg.html or such other URL as Google may provide.

"URL Terms" means the "Acceptable Use Policy," the "SLA," and the "TSS Guidelines."

Version: March 28, 2012

Log Delivery Service User Guide

Notice

As the global leader in Content Delivery Network (CDN) services, Akamai makes the Internet fast, reliable and secure. The company's advanced web performance, cloud security and media delivery solutions are revolutionizing how businesses optimize the consumer, enterprise or entertainment experience to any device, anywhere. To learn how Akamai solutions, and its team of Internet experts, are helping businesses move faster forward, please visit www.akamai.com and follow @Akamai on Twitter.

Table of Contents

Introduction

Akamai understands that you rely on us to deliver your content to give your end users a better overall online experience. We also understand that delivering your content means that we have the server logs that previously existed on your servers – and that this is often mission-critical information to your organization.

The Akamai Log Delivery Service provides you with the server logs from the various services that you are using, and is available for the following services:

- HTTP Content Delivery, HTTP Downloads, HTTP Downloads Client Side Delivery, Site Delivery, Web Application Accelerator, and Site Accelerator
- Streaming (QuickTime Streaming, Real Media Streaming, Windows Media Streaming, and Streaming for Flash)
- NetStorage (FTP)
- Fast File Upload module for Web Application Accelerator and Site Accelerator. Fast File Upload is for HTTP-based uploads using PUT/POST methods as opposed to FTP uploads.
- IPA/SXL, IP Application Accelerator and Session Accelerator
- Fast DNS
- Traffic Management
- AnswerX Managed

Intended Audience

This document is intended for those who receive logs from Akamai and are responsible for processing them.

How does Log Delivery Service work?

Akamai's infrastructure is constantly gathering the log entries from the thousands of edge servers around the world. Log Delivery Service creates a copy of these logs, separates your logs from other customer logs, and then delivers your logs based on a predetermined schedule. Most of the log files will be delivered within a 24-hour period. Due to the distributed nature of the network, some number of log lines can be delayed and be part of a later delivery.

You must be a customer of Log Delivery Service to configure the service and begin receiving logs from that point forward. Logs are not available retroactively, and logs for other customers who do not subscribe to Log Delivery Service are not retrievable.

Configuring Log Delivery Service

Log Delivery Service can be configured through Luna Control Center. Access to this configuration tool is available by default to users who have Technical or Admin role types and who are associated with contracts that include Log Delivery Service.

You can create one LDS configuration for each service associated with a CP code or DNS for Fast DNS or domain for Traffic Management. Services include HTTP Content Delivery, HTTP Downloads, HTTP Downloads for Client Side Delivery, Site Delivery, Streaming for Flash, QuickTime Streaming, Real Media Streaming, Windows Media Streaming, NetStorage (FTP) and the Fast File Upload module for Web Application Accelerator, Site Accelerator, Fast DNS and Traffic Management, and AnswerX Managed.

Once a service is initially configured for LDS, the tool shows whether the configuration is active or suspended, and allows you to view or edit it.

Access Log Delivery Service Configuration

To access Log Delivery Service configuration:

1. Log in to Luna Control Center.
2. Click **Configure** to expand the menu.
3. Click **Log Delivery**.

The Log Delivery Service page appears, showing the objects and service(s) associated with your login. This page indicates whether there is any existing log delivery configuration, or whether the current log delivery configurations are enabled or suspended. You can create new configurations and view or modify current configurations from this page, as well as view as summary of the configurations that are active. You can configure how many objects are included in the view at the bottom of the screen.

If you have configured Log Delivery Service for FTP delivery, you can test the FTP connection for a log delivery configuration under Action > Test FTP Connection.

Create a new Log Delivery Service configuration

When creating a new Log Delivery Service configuration, you can create a configuration by defining all options You can also copy an existing configuration by using the Copy button to configure one object, or the Bulk Copy button to configure multiple objects based on the one chosen configuration.

To create a new log delivery configuration for a service:

1. Find the object ID/service pair to be configured on the Log Delivery Service page.
2. Select the Begin Log Delivery action. Only services that lack a current log delivery configuration will have this action. The Begin Log Delivery page appears.
3. Define the options for your log delivery configuration, delivery method, and contact information. You will click Next three times before you have completed all options. See *Appendix B: Overview of Log Delivery Options*on page 10 for a complete discussion of available options.
4. Click Finish to initiation this Log Delivery Service configuration.

Because it is difficult to time the LDS provisioning process, a best practice is to set the start date to two days before you need the service enabled.

Editing an existing Log Delivery Service configuration

To edit an existing Log Delivery Service configuration:

1. Find the Object ID/service pair whose configuration you want to modify on the Log Delivery Service page.
2. Select the Change Settings action
3. Edit the configuration options as desired.
4. Click Finish.

Because it is difficult to time the LDS provisioning process, a best practice is to introduce changes two days before you need the service to be modified.

Enabling or suspending a Log Delivery Service configuration

To enable or suspend a Log Delivery Service configuration:

1. Find the Object ID/service pair whose configuration you want to suspend or resume on the Log Delivery Service page.
2. The service will have a Resume or a Suspend action, as appropriate, that allows you to perform that function on this configuration. Click the associated Resume or Suspend action.

Because it is difficult to time the LDS provisioning process, a best practice is to suspend or resume delivery two days before you need the service enabled.

Requesting redelivery of a log

The Luna Control Center enables you to request redelivery of logs. You can select the Object IDs for which you wish to redeliver the logs, and the time period of the desired data. Note that log redeliveries are sent at lower priority than normal log deliveries, and you may not request a redelivery until the original log has been completed. Logs are available for redelivery for the log retention period indicated next to each configuration.

Redelivery Intervals

The intervals at which redelivery occurs depends on the type of log aggregation associated with your object. The Types of Aggregation section discusses these options as part of configuring Log Delivery.

Aggregated by log arrival time: For objects configured for aggregation by log arrival time, you can select a date, start time, and end time for the redelivery. Log data will be redelivered for all time intervals in the selected time period, based on the delivery frequency that is configured for that CP code. For example, if you select a 24-hour time period for redelivery and your object is configured for a 24-hour delivery interval, you will receive one redelivery encompassing 24 hours of log data. If your object is configured for a one-hour delivery internal, you will receive 24 redeliveries, each encompassing one hour of log data.

Aggregated by calendar day: For objects configured for aggregation by calendar day (that is, hit time), you can select the date for the redelivery. Calendar day redeliveries can only be requested for entire days, so the start and end time fields are not displayed. All log data that has been collected for that date will be redelivered. Since some hits may have arrived after the original delivery, it is possible that the redelivery may include more log lines than the initial delivery.

Note that if you select multiple objects, the redeliveries for each object will occur according to that object's log configuration as described here.

To request a redelivery:

1. Select the **Redelivery** tab where you will see the list of already requested redeliveries. Select **New Redelivery**.
2. Select the product and configurations to be redelivered.
3. Identify the dates, start time, and end time in GMT. For calendar day deliveries, select the date for redelivery. Please note that all redeliveries will be based on the delivery frequency you're configured for, but will be optimized to include the interval between the selected start time and end time.
4. Click Redeliver.

If you are requesting a redelivery because the original delivery failed, first be sure that the cause of the original failure is corrected, and verify that the redelivery was successful once it is complete.

Overview of Log Delivery Options

This section describes the options you must define when configuring Log Delivery Service.

Log Delivery Options

Message size

The maximum message size applies to both email and FTP. For all encoding options other than MIME encoded email deliveries, Log Delivery Service attempts to send logs in parts no greater than the maximum size when compressed. However, logs are broken into parts before compression, so the estimation of message size is based on the uncompressed file size.

Estimation of compressed size currently assumes a 10-to-1 compression ratio. The same ratio is assumed for all customers and is based on historic log compressibility statistics. Because actual compressibility can vary from customer to customer, delivery parts can be larger or smaller than the configured size. For MIME encoded email deliveries, the logs will be split after compression.

Encoding

Log Delivery Service supports three mail formats: gzipped and uuencoded, MIME with base-64 encoding, and GPG Encrypted. Note that GPG Encrypted requires a key that you must upload. Otherwise, the main difference is support by mail clients. Most customers using email delivery are currently configured for gzipped and uuencoded logs.

There are four types of encoding, two of which are only available in email:

Type	Available via	Description
gz	FTP or email	GNU gzip
gpg	FTP or email	Gzipped and GPG encrypted
gzu	Email only	Gzipped and uuencoded
gzm	Email only	Mime (base64)

Status

Log delivery configurations, once created, will have a status of either "Suspended" or "Active" status. Services that have no log delivery configured will show a Begin Log Delivery action, through which you can set up a configuration. Log delivery configurations that are currently active provide you the option to suspend them.

Start date

The start date affects both the beginning of log collection and the delivery of logs, but is unrelated to the timestamps of the logs collected. Log collection will begin within no more than 6-8 hours of the start date, and the first delivery will be scheduled within no more than two days.

Because it is difficult to time the LDS provisioning process, a best practice is to set the start date to two days before you need the service enabled.

NOTE: LDS relies on data from the Domain Validation Tool. This works the same way in LDS as it does for Akamai's billing and reporting systems. Logs will only be collected if the origin server domain in the logs matches a domain on the valid domains list. The domain list is independent of any other LDS provisioning, and will go live on the LDS system automatically. It may take up to 3 hours for a domain to become live for Log Delivery Service configurations from the time it is entered in the Valid Domains area of Luna Control Center.

End date

The end date can be set to either an actual end date, or "indefinite" by leaving the field blank. If it is set to a specific end date, you might get a few empty deliveries before it ends. Like the start date, the end date is unrelated to the timestamps of the individual log lines, and refers only to log collection and delivery.

Administrative contact information

To complete the configuration, you must provide a name, phone number, and email address. This information will be used both for announcements about changes to the product as well as automated delivery failure notices.

Do not use the same email address to which emailed logs are sent. We strongly recommend that you also avoid using an address in the same domain or hosted on the same mail server as the delivery address; should that domain or server become unavailable, you will then still receive announcements and failure notices.

You can enter multiple email addresses, delimited by commas. Please make certain that all contacts that might need access to delivery failure notices are included in the email. You might find it simpler to give the address of a mailing list to which multiple users are subscribed, rather than updating the LDS configuration each time the email list needs to be changed.

It is essential that the admin email address is checked on a regular basis. Akamai only guarantees availability of logs for seven calendar days after they were originally delivered. If delivery failures are responded to promptly, five business days should be adequate to diagnose and correct the causes of delivery failures, as well as to contact Akamai to arrange a redelivery. For more on delivery failures and requesting redeliveries, see Troubleshooting.

Type of Delivery

Log Delivery Service offers delivery via email or FTP.

Email

If you are using email for your log delivery, be sure that your email server is able to handle the volume of mail that LDS will produce. We do not recommend using a standard corporate email server for handling logs from LDS.

Advantages of email:

- Widely available.
- Incorporates a robust delivery mechanism that queues messages and retries when servers become temporarily unavailable.
- Eliminates the need to configure an FTP server and manage usernames and passwords.

Disadvantages of email:

- Some email servers might not be able to handle high volumes.
- Cannot be used in conjunction with NetStorage.
- If log delivery emails are rejected by your email servers for some reason, Akamai has no way of knowing this or notifying you about the failed delivery.

FTP

Advantages of FTP:

- Can handle higher log volumes than the average mail server.

- Can be used in conjunction with NetStorage .

Disadvantages of FTP:

- Less widely available than email.
- More vulnerable to delivery failures, due to the lack of a built-in redelivery mechanism.
- Requires more initial configuration than email.

Type of Aggregation

Logs can be aggregated by arrival time or by calendar day.

NOTE:

1. All time periods referenced in the Aggregation discussion are GMT.

Aggregate by log arrival time

Aggregating logs by arrival time allows you to receive log data in the order in which it arrives from the edge servers. This is the fastest way to get log data, but it is more difficult to process. For example, a delivery that covers a 24-hour period will usually contain some data from the previous several days.

Logs can be delivered at intervals of 1, 2, 3, 4, 6, 8, 12 and 24 hours, with the default being every 24 hours. Less frequent deliveries are often more convenient to work with. More frequent deliveries will help ensure that you get your logs as soon as possible. Large logs, however, may require more frequent deliveries.

Aggregate by calendar day (hit time)

Aggregating logs by calendar day (hit time) allows you to receive most of the hits for each GMT calendar day (a 24-hour period) in a single bulk delivery. You may set the threshold, that is, the percentage of log completeness (described below), upon which you would like the log data sent. Please note that the higher you set the threshold, the longer you will have to wait for your log data. Residual data will contain any hits that were not processed in time to make it into the bulk delivery. This data is sent at regular intervals after each day for a period of four days if you check the "Deliver residual data" checkbox.

Completeness Threshold

Due to the distributed nature of the edge network, it is impossible to determine in advance that all hits that occurred on a particular day have been processed and are ready for delivery by LDS. For calendar day deliveries, you can set a completion threshold that is used by the LDS system to determine when the data is complete enough to send the bulk delivery. The completion threshold determines the percentage of expected logs that must be processed before delivery. Setting the threshold value to a high percentage will ensure that your delivery has data that is complete as possible, but may cause deliveries to be delayed to

meet the threshold. Setting the threshold value to a lower percentage will provide for more timely deliveries, but more of the data may be deferred to a leftover delivery.

Residual Delivery Frequency

You can also choose to receive leftover deliveries containing log lines that were not included in either the bulk delivery or a previous leftover delivery. Most of the log data will be included in the bulk delivery, so the leftover deliveries will often contain no log lines. However, the delivery will be sent, even if the file does not contain any log lines. These residual deliveries are sent once every 24 hours.

Log Identifier String

The log identifier string is a unique label that is added to the file names of log files. If you have multiple LDS configurations, the log file identifier string can make it much easier to distinguish them visually. Log identifier strings are limited to alphanumeric characters (numbers and letters) only, and may be no longer than 80 characters. They are case-sensitive, meaning that the case of names entered in the configuration tool will be preserved in log file names.

See *Appendix B: Filenames and Email Subject Lines* on page 25 for a complete description of log file names and email subject lines.

Email Address

You may only have logs delivered to a single email address. If you need logs to go to multiple email addresses, consider sending logs to a mailing list managed on your local mail server.

The email address is limited to 80 characters. You may see a phone number in previously configured email addresses; this information is pulled from the person's portal login information, if it's available.

FTP Login

You will need the FTP server name, the login, the password, and the directory path (if the account used for LDS does not automatically default to this directory).

Before accepting a new configuration, the portal will run a series of tests against your FTP server. These tests consist of trying to:

1. connect to the customer FTP server (uses active mode)
2. login with the userid and password given
3. cd to the log directory specified in the configuration
4. put a file on the server
5. delete a file (del)
6. do a directory listing on the log directory (dir)

7. overwrite a file (put the same file twice)
8. rename a file (mv)

Note that these tests will come from a machine associated with Luna Control Center, and not an actual LDS delivery machine.

Administrative contact information

To complete the configuration, you must provide a name, phone number, and email address. This information will be used both for announcements about changes to the product as well as automated delivery failure notices.

Do not use the same email address to which emailed logs are sent. We strongly recommend that you also avoid using an address in the same domain or hosted on the same mail server as the delivery address; should that domain or server become unavailable, you will then still receive announcements and failure notices.

You can enter multiple email addresses, delimited by commas. Please make certain that all contacts that might need access to delivery failure notices are included in the email. You might find it simpler to give the address of a mailing list to which multiple users are subscribed, rather than updating the LDS configuration each time the email list needs to be changed.

It is essential that the admin email address is checked on a regular basis. Akamai only guarantees availability of logs for seven calendar days after they were originally delivered. If delivery failures are responded to promptly, five business days should be adequate to diagnose and correct the causes of delivery failures, as well as to contact Akamai to arrange a redelivery. For more on delivery failures and requesting redeliveries, see Troubleshooting.

Formats and Example Log Lines

This chapter describes the available log formats and gives sample log lines for you to review. The "cs" and "sc" in log field names indicate the direction of communication, that is, client-to-server and server-to-client, respectively.

Luna Format Name	Description
dns	Fields: <ID> - <timestamp> <readable timestamp>,<requestor ip>,<requestor port>,<request name>,IN,<request type>,<"E" if EDNS0 request or blank>,<edns0 size if EDNS0 or blank>,<"D" if DNSSEC request or blank>,<"T" if TCP request or blank>,<answers>
Client Side Downloads	Application Version Information: Timestamp GUID ClientIP CPCODE AppVersion AppInfo StreamingFlags Certified Download: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize StreamingFlags StreamEvent StreamId HTTPSourceNum PeerSourceNum PeerOverheadBytes Download Status Update: Timestamp GUID ClientIP CPCODE TransId RL HTTPSourceDeltaBytes PeerDeltaBytes ValidDeltaBytes FailedDeltaBytes ObjectStatus ContentType StreamingFlags InstallCPCODE InstallTimestamp StreamEvent StreamId HTTPSourceNum PeerSourceNum PeerOverheadDeltaBytes EULA Acceptance Status: Timestamp GUID ClientIP CPCODE EulaTimestamp EulaVersion EulaAuthenticator

Luna Format Name	Description
	Failure - Client-Detected: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize Reason ExtraInfo StreamingFlags PeerOverheadBytes Initiated Download: Timestamp GUID ClientIP CPCODE TransId RL APIVersion StreamingFlags Installed a NetSession Interface Client: Timestamp GUID ClientIP CPCODE InstallTimestamp Paused Download: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize PauseReason StreamingFlags PeerOverheadBytes Resumed Download: Timestamp GUID ClientIP CPCODE TransId RL StartTimestamp ActiveMsec SourceBytes PeerBytes ValidBytes FailedBytes FileSize PauseReason StreamingFlags PeerOverheadBytes
Combined **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent"
Combined + Edge tokenization **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "x-edge-tokenization"

Luna Format Name	Description
Combined + HDD Delivery **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" guid bitrate playout-time encrypted-bytes stream-format
Combined + Web App FireWall Log record format:	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" "waf-info"
Common **Log record format:**	client_ip - - [date] "http_method arl HTTP/1.1" status_code total_bytes
W3C + Edge Tokenization **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" "x-edge-tokenization"
W3C + HDD **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \tguid \t bitrate \tplayout-time \tencrypted-bytes \tstream-format
W3C + WebApp Firewall Log record format:	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"x-wafinfo
W3C **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl \t status_code \t total_bytes \t transfer_time
WMS9	Fields: c-ip date time c-dns cs-url c-starttime x-duration c-rate c-status c-playerid c-playerversion c-playerlanguage cs(User-Agent) cs(Referer) c-hostexe c-hostexever c-os c-osversion c-cpu filelength filesize avgbandwidth protocol transport audi-ocodec videocodec channelURL sc-bytes c-bytes s-pkts-sent c-pkts-received c-pkts-lost-client c-pkts-lost-net c-pkts-lost-cont-net c-resendreqs c-pkts-recovered-ECC c-pkts-recovered-resent c-buffercount c-totalbuffertime c-quality s-ip s-dns s-totalclients s-cpu-util cs-user-name s-session-id s-con-tent-path cs-media-name c-max-bandwidth cs-media-role s-proxied
NetStorage	Fields: log_ts region leader_ip (fields[4-12] : Varying values in these fields) server_ip protocol client_ip user_id log_id flags network_id fwd_host session_id start_ts end_ts transfer_ts u_ flags request_id proto_ver host_header error overhead_bytes bytes_read bytes_send
Combined + ARL without domain	client_ip domain - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"

Luna Format Name	Description
Log record format:	
Combined + Cookie **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"
Combined + Cookie + Accept Language **Log record format:**	client_ip - - [date] "http_method arl-stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "accept_language"
Combined + Cookie + ARL without domain **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"
Combined+Cookie+ARL without domain+Host Header+Transfer Time+Custom Field **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" host_header transfer_time "custom"
Combined + Cookie + Custom Field **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "custom"
Combined + Cookie + Host Header **Log record format:**	client_ip - - [date] host_header "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie"
Combined + Cookie + Total Size + Byte Range Log record format:	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" total_object_size byte_range
Combined + Cookie + Total Size + Byte Range + Completion Status **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" total_object_size byte_range last_byte_served_flag_1_or_0
Combined + Referrer + User Agent **Log record format:**	client_ip - - [date] "http_method arl HTTP/1.1" status_code total_bytes "referrer" "user_agent"

Luna Format Name	Description
Combined + Time(ms) + Cookie + Custom Field **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "custom"
Combined + Time(ms) + Cookie + Host Header + Custom Field **Log record format:**	client_ip - - [date] host_header "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "custom"
Combined + UBB **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "UBB_GUID"
Combined + URI-stem Log record format:	Fields: date time c-ip cs-username cs-method cs-uri-stem cs-uri-query sc-status sc-bytes cs(User-Agent) cs(Cookie) cs (Referer)
Combined + Web App FireWall + Custom Field **Log record format**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "waf-info" "customfield"
Combined + Web App FireWall + Host Header **Log record format:**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_code total_bytes "referrer" "user_agent" "cookie" "host header"
Extended **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie"
Extended + Accept Language **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t "accept_language"
Extended + Astraeus	Fields: date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken astraeus-used reliability-mode
Extended + Byte Range **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer" \t "user_agent" \t "cookie" \t total_object_size \t byte_range
Extended + Completion Flag	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time

Luna Format Name	Description
Log record format:	\t "referrer" \t "user_agent" \t "cookie" \t total_object_size \t byte_range \t last_byte_served_flag
Extended + Content-Type + Cache-related	Fields: start-time-unixtimestamp-dot-ms cs-ip cs-method ssl-version cs-uri sc-status sc-content-bytes sc-total-bytes s-object-size s-uncompressed-size http-overhead-bytes cs (Referer) cs(User-Agent) cs(Cookie) cs(Host) sc(Content-Type) x-time-ssloverhead-ms x-time-turnaround-ms x-time-transfer-ms x-request-id x-max-age x-cache-status x-cache-refresh-source x-last-byte-served-flag x-no-store-flag
Extended + CPU Stats	Fields: date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) normalized-cpu
Extended + Custom Field **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie"\t "custom"
Extended + Custom Field + Server Content Bytes **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"-customfield"\t"content-bytes"
Extended + Custom Field + Without Client IP **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t --- \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"customfield"
Extended + EdgeJava	Fields: date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) dca-app-str
Extended + Entity Size **Log record format:**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" "entity-size"
Extended + Hashed Client IP **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \client ip \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"-customfield"\t"content-bytes"
Extended + Host Header **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" "host_header"
Extended + Host Header + ARL without domain **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie"

Luna Format Name	Description
Extended + Origin Retrieval Status	Fields: date time cs-ip cs-method cs-uri-origin cs-uri-stem cs-uri-query sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) cs(Host-Header) cs-request-id cached
Extended + Request End Time	Fields: start-time-unixtimestamp-dot-ms end-time-ms cs-ip cs-method ssl-version cs-uri sc-status sc-content-bytes sc-total-bytes s-object-size s-uncompressed-size http-overhead-bytes cs(Referer) cs(User-Agent) cs(Cookie) cs(Host) sc(Content-Type) x-time-ssloverhead-ms x-time-turnaround-ms x-time-transfer-ms x-request-id x-max-age x-cache-status x-cache-refresh-source x-last-byte-served-flag x-no-store-flag Note: end-time-ms is https://-docs.akamai.com/esp/user/edgesuite/log-format.xml#r4
Extended + UBB **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t x-ubbguid
Extended + URI-Query Log record format	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t cs_stem \t cs_query \t status_code \t "user_agent"
Extended + Without Client IP **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t --- \t http_method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" \t"entity-size"\t"-customfield"\t"content-bytes"
IIS 6.0 **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t client_ip \t http_ method \t arl_stem \t status_code \t total_bytes \t transfer_time \t "referrer"\t "user_agent"\t "cookie" "referer", status, -, S_ TOTAL_BYTES_SERVED, S_CLIENT_REQ_HDR_SIZE, S_ TRANSFER_TIME
Speedera **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)
Speedera + ARL without domain **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)
Speedera + ARL without domain + Last Byte Status **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)
Speedera + Last Byte Status **Log record format**	date_YYYY-MM-DD \t time_HH:MM:SS \t c_ip \t cs-username \t s-sitename \t s-computername \t s-ip \t cs-method \t cs-uri-stem \t cs-uri-query \t sc-status \t sc-win32-status \t sc-bytes \t cs-bytes \t time-taken \t s-port \t cs-protocol \t cs(User-Agent) \t cs (Cookie) \t cs(Referer)

Luna Format Name	Description
W3C + Web App Firewall + Host Header **Log record format`**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-wafinfo cs-host
Extended + Custom Field + Country code For Ghost IP + Sub Customer ID **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom country_code sub_customer_id
Extended + Custom Field + Country code For Ghost IP + Sub Customer ID + Host Header **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom country_code sub_customer_id host_header
Combined + Web App Firewall + Custom Field Without Query string **Log record format**	client_ip - - [date] "http_method arl_stem HTTP/1.1" status_ code total_bytes "referrer" "user_agent" "cookie" "waf-info" "customfield"
Extended + Custom field + GhostIP + Cache Status **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom ghostip cache_status
Extended + Error Message **Log record format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) err-msg
W3C + Custom Field + GeoCode + Bot **Log record format**	starttime(msec) endtime(msec) cs-method cs-host cs-path cs-status cs-query cs(User-Agent) client-ip cs-ref x-custom
Extended + Custom field + GhostIP + CacheStatus + WAF **Log Record Format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom ghostip cache_status x-wafinfo
W3C + Custom Field + GeoCode + Bot + Full Filename in cs-path	starttime(msec) endtime(msec) cs-method cs-host cs-path cs-status cs-query cs(User-Agent) client-ip cs-ref x-custom

Luna Format Name	Description
Log Record Format	
Extended + Custom field + GhostIP + Cache Status + Host Header **Log Record Format**	date time cs-ip cs-method cs-uri sc-status sc-bytes time-taken cs(Referer) cs(User-Agent) cs(Cookie) x-custom ghostip cache_status cs(host-header)
BIND	<date> <time> "client" <requester IP>#<requester port> (<question name>): <query or response> <question name> <class> <question type> <flags> Flags include: • +/- (whether recursion desired or not) • "T" if TCP • "C" if CD set

Filenames and Email Subject Lines

Overview

Log filenames in both email and FTP are the same. Email subject lines are the same as filenames.

The format is:

[Identifier String]_[CP Code].[Format]_[Sorting Status].[Date and Time]-[Part Number].[Encoding]

For example:

customer_1234.esclf_S.200401250000-2400-0.gz

Log Identifier String

The log identifier string is a unique label that is added to the file names of log files. If you have multiple LDS configurations, the log file identifier string can make it much easier to distinguish them visually. Log identifier strings are limited to alphanumeric characters (numbers and letters) only, and may be no longer than 80 characters. They are case-sensitive, meaning that the case of names entered in the configuration tool will be preserved in log file names.

See *Appendix B: Filenames and Email Subject Lines*on page 25 for a complete description of log file names and email subject lines.

CP code

LDS can be per CP code. If you have multiple services associated with a single CP code (for example, multiple streaming formats), you might have several Log Delivery Service configurations using the same CP code (but only one per service per CP code).

Sorting status

Log files are labeled either "S" for sorted or "U" for unsorted. See *Appendix B: Sorting of Logs*on page 28 for more details on sorting.

Date and time

NOTE:

1. Times referenced in Log Delivery file names use GMT.

The date and time field is in the format YYYYMMDDHHMM-HHMM.

For example:

customer_1234.esclf_S.200401241200-2400-0.gz

will be followed by

customer_1234.esclf_S.200401250000-1200-0.gz

Part number

Log files may be split into parts to keep the individual parts below the configured maximum file size. Part numbering begins at 0.

Encoding

Log Delivery Service supports three mail formats: gzipped and uuencoded, MIME with base-64 encoding, and GPG Encrypted. Note that GPG Encrypted requires a key that you must upload. Otherwise, the main difference is support by mail clients. Most customers using email delivery are currently configured for gzipped and uuencoded logs.

There are four types of encoding, two of which are only available in email:

Type	Available via	Description
gz	FTP or email	GNU gzip
gpg	FTP or email	Gzipped and GPG encrypted
gzu	Email only	Gzipped and uuencoded
gzm	Email only	Mime (base64)

File names for deliveries aggregated by hit time

NOTE:

1. Times referenced in Log Delivery file names use GMT.

File names for deliveries using aggregation by hit time also include information about the time range that was examined for hits that occurred on the specified day and the completion percentage that was met when the delivery was scheduled.

The format is:

[Identifier String]_[CP Code].[Format]_[Sorting Status].[Time Range Start]-[Time Range End]-[Date]-[CompletionPercentage]-[PartNumber].[Encoding]

For example:

customer_1234.esclf_S.200401250000-200401260500-20040125-99.5c-0.gz

This file name indicates the bulk delivery of the data for 2004.01.25, using a completion threshold of 99.5 percent, and checking logs that were processed between 2004.01.25 00:00 GMT and 2004.01.26 06:00 GMT for hits that occurred on 2004.01.25.

Sorting of Logs

Within a single delivery, logs are sorted one of two ways:

- The logs are sorted before being broken into parts, which is indicated by an "S" in the filename. Reassembling the parts produces a complete, sorted log.
- The logs are broken into parts, and then each part is sorted separately, which is indicated by a "U" in the filename. Reassembling the parts does not produce a complete, sorted log.

Only deliveries where the number of log lines is above a certain threshold will be delivered with parts sorted separately; all others will be completely sorted. The current threshold is based on uncompressed log file size, and usually approximates 20 million log lines per delivery. This might change in the future as log volumes increase or more resources become available for sorting.

Note that despite sorting within deliveries, logs are processed and delivered in the order in which they are received. Nearly all logs arrive in the LDS system in a timely fashion and are processed very quickly. However, it is normal to see considerable overlap of a few hours between deliveries accompanied by a much smaller trickle of logs for up to a few days. Small variations in log arrival time guarantee that there will always be some overlap between the time periods covered by each delivery, and arrival of some logs to the LDS system can be delayed for longer due to network-related factors outside of Akamai's control. The exact pattern will vary depending on the geographic distribution of your end users.

Any automated processes or log-analysis software you are using will need to be configured to take this potential time overlap into account. See *Appendix B: Using 3rd-Party Log Analysis Tools* on page 29 for some further recommendations on working with LDS and log-analysis software.

Using 3rd-Party Log Analysis Tools

The following are general suggestions or workarounds to ensure the best possible compatibility between Akamai logs and 3rd-party log analysis tools.

Format Recognition

Akamai log formats generally follow widely accepted standards. Occasionally, log-analysis software has trouble with W3C or Combined Log Format. Switching to another format often solves the problem.

Chronological Ordering of Logs

Most log analysis software is designed to have all logs processed in chronological order, and will throw out log lines that are not in order. However, LDS is designed in such a way that there will always be overlap between deliveries (Akamai logs are sorted within deliveries – see Sorting of Logs for details).

Your software configuration will involve a trade-off between completeness of data and availability of reports. The longer you wait for all remaining logs to arrive so that you can merge and sort them with previous deliveries, the longer you have to wait for your reports. Because time-range overlap will vary from customer to customer, we recommend some initial examination of your logs from LDS before choosing a software configuration scheme.

WebTrends

Akamai has a close relationship with NetIQ, makers of WebTrends Reporting Center and WebTrends Analysis Suite. We have done some testing to make sure that our logs are as compatible with their software as possible.

To feed deliveries sorted in parts into WebTrends software, customers can use the following workaround. Define a WRC "profile" that considers the parts of an LDS delivery as coming from (in NetIQ terminology) a "server cluster", a web site hosted on multiple machines. Each part can be considered as the data from one server in the cluster. WRC has support to merge these files and to produce aggregate statistics on the entire collection of log files.

We have tested this workaround at Akamai and have verified with NetIQ that this "cluster support" is available in WebTrends Analysis Suite (advanced edition) and in WebTrends Reporting Center v2.0+. WebTrends Analysis Suite (standard) does not have this support.

Log Delivery Times

Log delivery times can vary depending on the load placed on the LDS infrastructure. The times of daily deliveries are especially vulnerable to load-induced variation due to the size and number of logs being delivered. It is usually possible to increase the predictability of delivery times by reducing the frequency of delivery. In any case, automated processes set up to manage incoming logs should allow for varying delivery times.

Troubleshooting

The following section provides recommendations and instructions for avoiding or fixing problems with Log Delivery Service.

Failure Notification Emails

If Akamai is unable to deliver your logs, an email will be sent to the admin contact you have specified in the configuration tool on Luna Control Center. It will give you the object (CP code or domain), date, and time range, some explanatory text about the error, and a transcript of the delivery session (either FTP or SMTP). It will be your responsibility to determine the cause of the failure and correct it if necessary before requesting a redelivery.

The help messages in the emails are based primarily upon the response codes returned, and are fairly general. Because of the variety of FTP and mail server configurations available, it is not always possible to tell what has gone wrong from the help message alone. In most cases it will be necessary to read the transcript.

To help you with this, we have provided some troubleshooting examples from delivery transcripts along with likely interpretations. All are taken from real delivery failures that we have seen, modified to make them anonymous. Remember that not all email and FTP servers produce the same error messages, and that response codes don't always correspond logically to the events that trigger them.

After reading the transcript, you might find it helpful to replicate the behavior of the LDS system. Try sending an email with a similar subject line and size, or delivering a file to the same FTP account. Note that many common problems affecting LDS are transient or intermittent, and might not be detectable during troubleshooting.

Example email errors

In many cases the symptoms of an email problem are several failed attempts ("bounces") concluding in failure after several days. LDS will send a failure notice on the initial bounce, even though the delivery might later succeed.

The mail server is rejecting the deliveries because there is no user with that address.

h80.67.64.10 does not like recipient.
Remote host said: 550 <user@lds-customer.com>: User unknown
DGiving up on 80.67.64.10

There is no known user with that address.

Either create the user, or change the LDS configuration.

The mail server is unreachable.

ZSorry, I wasn't able to establish an SMTP connection. (#4.4.1)

ZConnected to 80.67.64.10 but connection died. (#4.4.2)

This could have a variety of potential causes, including excessive packet loss, routing problems, or the email server being down.

Try sending mail to the address and checking connectivity to and from the mail server.

There are too many active connections to the mail server and it is rejecting new ones.

ZConnected to 80.67.64.10 but greeting failed.
Remote host said: 421 lds-customer.com connection limit reached
DConnected to 80.67.64.10 but sender was rejected.
Remote host said: 505 Authentication required

This is unusual, but can happen when multiple large deliveries are still in progress and overloading the email server. Anything that causes deliveries to take a long time could be the cause.

Note that in this case several previous deliveries might be effectively "hung", and might need to be terminated and redelivered.

LDS can't determine the IP address of the mail server due to DNS problems.

DSorry, I couldn't find any host named lds-customer.com. (#5.1.2)

This could be due to a single incident affecting your DNS server, intermittent DNS performance problems, connectivity problems to the DNS server, or a configuration problem with DNS records for the domain.

Try querying the DNS directly (possibly from a location outside of your corporate network), or checking connectivity to and from it.

Example FTP errors

DNS times out when trying to resolve the FTP server.

open: host name resolve timeout

This could be due to connectivity problems between LDS and the DNS server, or simply a DNS server that fails to respond.

Try querying the DNS directly (possibly from a location outside of your corporate network), or running a traceroute to or from it.

The FTP server can't upload a file of the same name.

Error 553 in FTP put: ftp_test: Permission denied on server. (Overwrite)

Please make sure the user has permission to upload a file into the destination

directory.

The FTP server might be unable to upload a file if the filename already exists. Rename the file and try to upload the file again.

The FTP server is full, or a quota is exceeded.

[8526] FTP.lds-customer.com <--- 452 Error writing file: No space left on device.

[22595] FTP.lds-customer.com <--- 425 Data connection error: No space left on device

[25814] FTP.lds-customer.com <--- 452 Transfer aborted. No space left on device

[26026] FTP.lds-customer.com <--- 452-Maximum quota exceeded. Transfer aborted.

The FTP server is full, or a quota is exceeded.

It is important that you have some method for ensuring that there will always be space for log deliveries. For many customers, that takes the form of automated removal of logs to another server. If this is being done on Akamai NetStorage, you can set up age-based deletion in Luna Control Center.

The FTP server is refusing connections.

[12620] FTP.lds-customer.com **** Socket error (Connection refused) – reconnecting

This error can occur because of a misconfigured access control list.

Check the FTP server. Try connecting to it yourself. Look for a misconfigured access control list.

The FTP server is unexpectedly closing the connection.

[9036] FTP1.lds-customer.com **** Peer closed connection

[22397] FTP2.lds-customer.com <--- 421
[22397] FTP2.lds-customer.com **** remote end closes connection
[22397] FTP2.lds-customer.com ---- Closing control socket

[32620] FTP3.lds-customer.com <--- 426 Connection closed; transfer aborted.

LDS can establish a connection, but it consistently times out before logs can be delivered.

[28859] FTP.lds-customer.com <--- 150 Opening BINARY mode data connection for customer_1234.esclf.200401250000-2400-0.gz.
[28859] FTP.lds-customer.com **** Timeout – reconnecting
[28859] FTP.lds-customer.com ---> ABOR
[28859] FTP.lds-customer.com ---- Closing control socket

[5019] 80.67.64.10 <--- 150 Opening BINARY mode data connection for customer_1234.esclf.200401250000-2400-0.gz.
[5019] 80.67.64.10 **** Broken pipe
[5019] 80.67.64.10 ---> ABOR
[5019] 80.67.64.10 ---- Closing control socket

The connection to the FTP server consistently times out before logs can be delivered. This error tends to be intermittent.

Check connectivity to your FTP server over time.

LDS Public GPG Key

The following is the public GPG key for Log Delivery Service:

-----BEGIN PGP PUBLIC KEY BLOCK-----

Version: GnuPG v1.4.11 (GNU/Linux)

mQGiBFTJAgsRBAD4RqtHK6AJjzenMk6JwaA6a7f09ZKObQ6mc61LblFml/BzqeIA

CZTvSGv31dPnEddyLnSFlAoVJCWIWl6DJ3rkanVqmX/kxCioGT3l9rmJWcEigHtj

GjGKW0YNWeWhz0WlzjuYP0yvGrLSxn6g8PRiBa6VT04T5MBgf4rivxSnhwCgxfgn

5lhVmHP/XNA+bWGnMCLkYYUD/29FIJEkw1IuNyBmf3RWaJzwsxHrz504G0phPqf/

D0k2RSDvIKp6rDrYDHHqRmRl7dHNOAz5qcGjQjEXP4SOFFQ1wxsVV+XJtMs27XWm

2E0Ge3AiC8ctQNhPca8uyyPOHbw4NrYoJgRGib1ZdiXdq9VB5PTxZsaaCniQ/tRm

16tgA/9If7k6ifJSTCL1X06MOONcuCnlwMEIHbyIpY6uZUSeJfMUuJp1MFBzoEDQ

fd1ph00e61IsosPOUqgfu7wuIIkpfvEKzlX3DLjHmgY/YBDDgqySAngtNy9mkS9I

MBs7zEn1q8JbIQrAAKvhNNZKGO/zoYs5hMui5tLoQ/JNR1bG2rRcQWthbWFpIExv

ZyBEZWxpdmVyeSBTZXJ2aWNlIChTZXJ2aWNlIGFkZHJlc3Mgb25seS4gIERvIG5v

dCByZXBseS4pIDxsZHNAbGRzLWxvZ3MuYWthbWFpLmNvbT6IYAQTEQIAIAUCVMkC

CwIbAwYLCQgHAwIEFQIIAwQWAgMBAh4BAheAAAoJEJ2+a/Qkpjjomt IAn132fwvu

kgBp6aupXSTNt52xQ095AJ93BAdwoW6Bq5Cm5QrbNUg954ouLrkCDQRUyQIMEAgA

pa7RIAjxtsCV/82Byi1Gv90jdIDEpD22XYo3kZ5TyPQqgG4rnx7yD6NrLeUPtHsY

XJhalqCNEPqn6+0u8sh4gwh9HghrooweY3Ly8CV/peW0DazvsHhYoYQIM4IxbUmY

zBCIJS5Pf7rHV9dFf+foAx0G80ZWyXxga/bBvTnZ1NWFJ56dkQBtTORng061mlUY

dbmyjqHWiQ+w0YDvC/+6pQn2n3DJCQ7kiazWpCQcab3I/4pv4t7bt7JmBhbPSDP5

M6piC8RBdDdWut1Kf8eboEwXrAxw0EaSal7vF981KjWGuo70i9EZWav++JIWl5uM

so2nIav1f3FBtYcGOUyoswADBQgAj3Githc66p8/hm1gGkYhyaZ9AHgqoHlWHDyu

WqTaNAwOyXqfg3fKCaD49kACFc/FM440NS24rJf11Js++UtBznzbr+1sxuRAHpT3

/+ZrdGvpS8ZwvTUpPdootqtqvC0simXh6Q0456HMnuOfE077BHXUe8OdY2tkAjSP

rZ7taZzNNCN5zf2+JzWMn7Obk6slKI/bV3eQpKZgqIbRC14FkuzdlTlfOW/WXcPL

/ekIXQ3T16LXjut1+SRUz+b/Zy3AS+TJ/ldCCSNqLnG2LtXtd52VmfXv2keoR7PT

5lcWFwMVxaoSh0XyJXNcIVoILN54D+auJpfLKk9COo3z/xCgvohJBBgRAgAJBQJU

yQIMAhsMAAoJEJ2+a/QkpjjohqUAnikbkUTbwkqp0ttFyRO+kHS4MyeSAJ9LoEKT

BWwW1DvqzY3Hka7RJkb5kA==

=WY8v

-----END PGP PUBLIC KEY BLOCK-----

Procmail Configuration Examples

Procmail can be used to automatically put incoming logs from email in a certain directory. Postfix uses procmail and sendmail as its default mail delivery agents. Therefore, you can use either Postfix or Sendmail and the following instructions will work.

(http://www.redhat.com/support/docs/faqs/RH-postfix-FAQ/x69.html)

Suppose that mail fromLog Delivery Service had a header that looks like this:

From: lds@<*>.akamai.com

Subject: Customer_100.wm_S.200010080400-0100-0.log.gz

Where the subject contains a filename, which is unique for the date. Suppose that you have a "virtual" user set up to receive the email called

incoming-logs@yoursite

Suppose you want the incoming log to be put into

/home/ftp/available/Customer_100.wm_S.200010080400-0100-0.log.gz

so that your software can grab it off of your in-house FTP server customers can automatically put incoming logs from email to a certain directory.

Then you would set up a procmailrc file in ~incoming-logs/.procmailrc as follows:[1]

---------- cut here for .procmailrc ----------

LOGFILE=procmail.log

SUBJECT=`formail -xSubject:`

[1]Disclaimer: Akamai provides these examples to help you set up your own systems to receive and handle the Log Delivery Service email. The above code is not guaranteed to function and should only serve as a guideline. See man procmailex for additional useful examples.

```
:0:

* ^From:.* lds@<*>.akamai.com

| cat > /home/ftp/available/${SUBJECT}

----------- cut here ------------
```