

**STATE OF UTAH COOPERATIVE CONTRACT**

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

Zivaro, Inc. f/k/a Global Technology Resources, Inc DBA GTRI

Name

990 S Broadway Suite 300

Street Address

Denver

CO

80209

City

State

Zip

Vendor # VC0000120803 Commodity Code #: 920-05 Legal Status of Contractor: For-Profit Corporation

Contact Name: Gail Springer Phone Number: 720-836-7331 Email: gspringer@gtri.com

2. CONTRACT PORTFOLIO NAME: Cloud Solutions.
3. GENERAL PURPOSE OF CONTRACT: Provide Cloud Solutions under the service models awarded in Attachment B.
4. PROCUREMENT: This contract is entered into as a result of the procurement process on FY2018, Solicitation# SK18008
5. CONTRACT PERIOD: Effective Date: Wednesday, April 17, 2019. Termination Date: Tuesday, September 15, 2026 unless terminated early or extended in accordance with the terms and conditions of this contract.
6. Administrative Fee: 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) of contract sales 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.
7. 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 Schedule
ATTACHMENT D: Contractor's Response to Solicitation # SK18008
ATTACHMENT E: Service Offering EULAs, SLAs
- Any conflicts between Attachment A and the other Attachments will be resolved in favor of Attachment A.**
9. DOCUMENTS INCORPORATED INTO THIS CONTRACT BY REFERENCE BUT NOT ATTACHED:
- All other governmental laws, regulations, or actions applicable to the goods and/or services authorized by this contract.
 - Utah Procurement Code, Procurement Rules, and Contractor's response to solicitation #SK18008.
10. 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. Notwithstanding verbal or other representations by the parties, the "Effective Date" of this Contract shall be the date provided within Section 5 above.

CONTRACTOR



Contractor's signature

4/16/2019

Date

DIVISION OF PURCHASING


Christopher Hughes (Apr 16, 2019)

Director, Division of Purchasing

Apr 16, 2019

Date

Gail Springer, Director Federal Programs

Type or Print Name and Title



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 Addendum¹ ("PA");
- (2) NASPO ValuePoint Master Agreement Terms & Conditions, including the applicable Exhibits² 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 SaaS.

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: Unless otherwise specified as a shorter term in a Participating Addendum, the term of the Master Agreement will run from contract execution to September 15, 2026.

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 claims, damages or causes of action including reasonable attorneys' fees and related costs for any death, injury, or damage to property arising directly or indirectly 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. If the Contractor promptly and

reasonably investigates and defends any Intellectual Property Claim, it shall have control over the defense and settlement of it. However, the Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Contractor's reasonable request and expense, information and assistance necessary for such defense. If the Contractor fails to vigorously pursue the defense or settlement of the Intellectual Property Claim, the Indemnified Party may assume the defense or settlement of it and the Contractor shall be liable for all costs and expenses, including reasonable attorneys' fees and related costs, incurred by the Indemnified Party in the pursuit of the Intellectual Property Claim. Unless otherwise agreed in writing, this section is not subject to any limitations of liability in this Master Agreement or in any other document executed in conjunction with this Master Agreement.

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;

(2) CLOUD MINIMUM INSURANCE COVERAGE:

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

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

(4) Professional Liability. As applicable, Professional Liability Insurance Policy in the minimum amount of \$1,000,000 per occurrence and \$1,000,000 in the aggregate, written on 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) names the Participating States identified in the Request for Proposal as additional insureds, (2) provides that no material alteration, cancellation, non-renewal, or expiration of the coverage contained in such policy shall have effect unless the named Participating State has been given at least thirty (30) days prior written notice, and (3) provides that the Contractor's liability insurance policy shall be primary, 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. 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 Office³.

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. States or other entities permitted to participate may use an informal competitive process to determine which Master Agreements to participate in through execution of a Participating Addendum. 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

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.

requests for such participation through NASPO ValuePoint. Any permission to participate through execution of a Participating Addendum is not a determination that procurement authority exists in the Participating Entity; they must ensure that they have the requisite procurement authority to execute a Participating Addendum.

h. Resale. Subject to any explicit permission in a Participating Addendum, Purchasing Entities may not resell goods, software, or Services obtained under this Master Agreement. This limitation does not prohibit: payments by employees of a Purchasing Entity as explicitly permitted under this agreement; sales of goods to the general public as surplus property; and fees associated with inventory transactions with other 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: Orders under this Master Agreement are fixed-price or fixed-rate orders, not cost reimbursement contracts. 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: Contractor will provide access to Purchasing Entity's Data only to those Contractor employees, contractors and subcontractors ("Contractor Staff") who need to access the Data to fulfill Contractor's obligations under this Agreement. Contractor shall not access a Purchasing Entity's user accounts or Data, except on the course of data center operations, response to service or technical issues, as required by the express terms of this Master Agreement, or at a Purchasing Entity's written request.

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

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

23. Operations Management: 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 Contractor warrants that the Products it provides under this Master Agreement are free of malware. The Contractor must use 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 a separate transition Statement of Work.

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

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

33. Waiver of Breach: Failure of the Lead State, Participating Entity, 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 JAGGAER, formerly SciQuest, whereby JAGGAER 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://calculator.naspovaluepoint.org>. 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 H.

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. NASPO ValuePoint Cooperative Program Marketing, Training, and Performance Review:

- a. Contractor agrees to work cooperatively with NASPO ValuePoint personnel. Contractor agrees to present plans to NASPO ValuePoint for the education of Contractor's contract administrator(s) and sales/marketing workforce regarding the Master Agreement contract, including the competitive nature of NASPO ValuePoint procurements, the Master agreement and participating addendum process, and the manner in which qualifying entities can participate in the Master Agreement.
- b. Contractor agrees, as Participating Addendums become executed, if requested by ValuePoint personnel to provide plans to launch the program within the participating state. Plans will include time frames to launch the agreement and confirmation that the Contractor's website has been updated to properly reflect the contract offer as available in the participating state.
- c. Contractor agrees, absent anything to the contrary outlined in a Participating Addendum, to consider customer proposed terms and conditions, as deemed important to the customer, for possible inclusion into the customer agreement. Contractor will ensure that their sales force is aware of this contracting option.
- d. Contractor agrees to participate in an annual contract performance review at a location selected by the Lead State and NASPO ValuePoint, which may include a discussion of marketing action plans, target strategies, marketing materials, as well as Contractor reporting and timeliness of payment of administration fees.
- e. Contractor acknowledges that the NASPO ValuePoint logos may not be used by Contractor in sales and marketing until a logo use agreement is executed with NASPO ValuePoint.
- f. The Lead State expects to evaluate the utilization of the Master Agreement at the annual performance review. Lead State may, in its discretion, terminate the Master Agreement pursuant to section 6 when Contractor utilization does not warrant further administration of the Master Agreement. The Lead State may exercise its right to not renew the Master Agreement if vendor fails to record or report revenue for three consecutive quarters, upon 60-calendar day written notice to the Contractor. This subsection does not limit the discretionary right of either the Lead State or Contractor to terminate the Master Agreement pursuant to section 7.
- g. Contractor agrees, within 30 days of their effective date, to notify the Lead State and NASPO ValuePoint of any contractual most-favored-customer provisions in third-part contracts or agreements that may affect the promotion of this Master Agreements or whose terms provide for adjustments to future rates or pricing based on rates, pricing in, or Orders from this master agreement. Upon request of the Lead State or NASPO

ValuePoint, Contractor shall provide a copy of any such provisions.

45. NASPO ValuePoint Cloud Offerings Search Tool: In support of the Cloud Offerings Search Tool here: <http://www.naspovaluepoint.org/#!/contract-details/71/search> Contractor shall ensure its Cloud Offerings are accurately reported and updated to the Lead State in the format/template shown in Attachment I.

46. 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.

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.

Attachment B – Scope of Services Awarded to Contractor

1.1 Awarded Service Model(s).

Contractor is awarded the following Service Model:

- Software as a Service (SaaS)

1.2 Risk Categorization.*

Contractor's offered solutions offer the ability to store and secure data under the following risk categories:

Service Model	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered
SaaS	x	x	x	Private, public, community and hybrid

*Contractor may add additional OEM solutions during the life of the contract.

2.1 Deployment Models.

Contractor may provide cloud based services through the following deployment methods:

- **Private cloud.** The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
- **Community cloud.** The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
- **Public cloud.** The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
- **Hybrid cloud.** The cloud infrastructure is a composition of two or more distinct cloud infrastructures (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)

Attachment C - Pricing Discounts and Schedule

Contractor: Zivaro, Inc. f/k/a Global Technology Resources, Inc. DBA GTRI

Pricing Notes

1. % discounts are based on minimum discounts off Contractor's commercially published pricelists versus fixed pricing. Nonetheless, Orders will be fixed-price or fixed-rate and not cost reimbursable contracts. Contractor has the ability to update and refresh its respective price catalog, as long as the agreed-upon discounts are fixed.
2. Minimum guaranteed contract discounts do not preclude an Offeror and/or its authorized resellers from providing deeper or additional, incremental discounts at their sole discretion.
3. Purchasing entities shall benefit from any promotional pricing offered by Contractor to similar customers. Promotional pricing shall not be cause for a permanent price change.
4. Contractor's price catalog include the price structures of the cloud service models, value added services (i.e., Maintenance Services, Professional Services, Etc.), and deployment models that it intends to provide including the types of data it is able to hold under each model. Pricing shall all-inclusive of infrastructure and software costs and management of infrastructure, network, OS, and software.
5. Contractor provides tiered pricing to accompany its named user licensing model, therefore, as user count reaches tier thresholds, unit price decreases.

Cloud Service Model: Software as a Service (SaaS)

Description	Discount
Splunk License and maintenance	1.00%
Average SaaS OEM Discount Off	1.00%

Additional Value Added Services

Item Description	Onsite Hourly Rate		Remote Hourly Rate	
	NVP Price	Catalog Price	NVP Price	Catalog Price
Maintenance Services	please see labor rates below			
Professional Services	please see labor rates below			
Deployment Services	please see labor rates below			
Integration Services)	please see labor rates below			
Consulting/Advisory Services	please see labor rates below			
Architectural Design Services	please see labor rates below			
Statement of Work Services	please see labor rates below			
Partner Services	please see labor rates below			
	please see labor rates below			

Deliverable Rates

	NVP Price	Catalog Price
Enterprise Systems Engineer Level 1	\$ 117.90	\$ 180.00
Enterprise Systems Engineer Level 2	\$ 160.50	\$ 245.00
Enterprise Systems Engineer Level 3	\$ 199.80	\$ 305.00
Enterprise Systems Engineer Level 4	\$ 239.10	\$ 365.00
Network Engineer Level 1	\$ 91.70	\$ 140.00
Network Engineer Level 2	\$ 111.35	\$ 170.00
Network Engineer Level 3	\$ 147.40	\$ 225.00
Network Engineer Level 4	\$ 180.15	\$ 275.00
Network Engineer Level 5	\$ 222.70	\$ 335.00
Wireless Engineer Level 3	\$ 160.50	\$ 245.00
Wireless Engineer Level 4	\$ 186.70	\$ 285.00
Wireless Engineer Level 5	\$ 203.05	\$ 310.00
Project Coordinator Level 1	\$ 88.45	\$ 135.00
Project Coordinator Level 2	\$ 98.25	\$ 150.00
Project Coordinator Level 3	\$ 111.35	\$ 170.00
Project Manager Level 1	\$ 111.35	\$ 170.00
Project Manager Level 2	\$ 137.55	\$ 205.00
Project Manager Level 3	\$ 153.95	\$ 230.00
Network Security Engineer Level 2	\$ 117.90	\$ 180.00
Network Security Engineer Level 3	\$ 147.40	\$ 225.00
Network Security Engineer Level 4	\$ 183.40	\$ 280.00
Network Security Engineer Level 5	\$ 203.05	\$ 310.00
Subject Matter Expert Level 1	\$ 140.85	\$ 210.00
Subject Matter Expert Level 2	\$ 160.50	\$ 245.00
Subject Matter Expert Level 3	\$ 199.80	\$ 305.00
Subject Matter Expert Level 4	\$ 222.70	\$ 335.00
Cloud Architect	\$ 200.16	\$ 305.54

Attachment C - Pricing Discounts and Schedule

Contractor: Zivaro, Inc. f/k/a Global Technology Resources, Inc. DBA GTRI

System Administrator Level 1	\$ 82.43	\$ 125.83
System Administrator Level 2	\$ 100.09	\$ 152.79
System Administrator Level 3	\$ 129.52	\$ 197.71
Cloud Administrator Level 1	\$ 117.75	\$ 179.74
Cloud Administrator Level 2	\$ 141.29	\$ 215.68
Cloud Administrator Level 3	\$ 164.84	\$ 251.62
Splunk Professional Services	\$ 299.00	\$ 312.50

State of Utah and NASPO Value Point

Cloud Solutions

Solicitation SK18008

July 6, 2018

Detailed Product Offering

Prepared for:

State of Utah
State Contract Analyst
Division of Purchasing

Prepared by:

Global Technology Resources, Inc.
990 S. Broadway, Suite 300
Denver, CO 80209
1-877-603-1984
Fax 1-888-803-6520
www.GTRI.com



1. NASPO Pricing Notes and Explanations

GTRI is pleased to provide the State of Utah with its pricing response to Bid #: SK18008. The attached NASPO Pricing offers pricing for the following Cloud Models: IaaS, and PaaS.

GTRI is a leading provider of Cloud Solutions and Services to the Public Sector and a leading delivery and resale partner for Amazon Web Services (AWS). Our distinct capabilities and qualifications include the following:

- AWS Advanced Public Sector Partner
- Exceptional cybersecurity as demonstrated by Cloud Security Alliance (CSA) Assessments
- ISO 9001, certifications
- Agile project delivery

The following consulting services and solutions are available:

- Section 1: Professional Services and access to AWS Professional Services
- Section 2: Managed Service Provider (MSP) offerings backed by a rigorous AWS audit process
- Section 3: Resale of the full AWS IaaS catalog
- Section 4: AWS Marketplace and related offerings
- Section 5: Pricing and rate card for consulting services

1.1. Professional Services

GTRI is one of the leading public sector providers of cloud professional services. Our highly qualified, AWS-certified consultants have extensive experience in cloud solutions, including transition from legacy environments, migration of workloads, and optimization in the cloud. Services can be tailored to individual customer requirements and will be priced according to those requirements using the appropriate labor rates as listed in Section 5. Services can also be delivered through GTRI by AWS-badged consultants (AWS Professional Services).

1.2. AWS Marketplace and Related SaaS Offerings

Service Name	Description	Service Model	Deployment Models	Pricing
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.	SaaS	All (Public, Private, Hybrid, Community)	Per our AWS reseller agreement, there is no standard discount for Marketplace Offerings.
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	SaaS	All (Public, Private, Hybrid, Community)	Minimum 2% discount off list price

	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.	SaaS	All (Public, Private, Hybrid, Community)	Per our AWS reseller agreement, there is no standard discount for Marketplace Offerings.
Splunk Licenses for Logging	As an Elite Splunk Partner, GTRI has the ability to obtain Splunk licenses for a full monitoring platform to be added to any AWS environment.	SaaS	All (Public, Private, Hybrid, Community)	1% off MSRP

1.3. Cloud Solutions Pricing

GTRI offers all AWS IaaS resale offerings at a minimum 2% discount off the AWS list price. This model has several advantages over traditional SKU- based price catalogs:

- **Real Time Discounts.** By providing a discount off list price, we are always able to pass along any price reductions immediately as they occur. This is significant because AWS frequently publishes price reductions and as of 2016 had reduced prices more than 50 times since inception. NASPO ValuePoint Participating Entities get those price reductions immediately without waiting for a formal catalog refresh.
- **Simple and Flexible:** AWS has more than 25,000 SKUs which must be priced individually in a traditional model. We instead use the AWS Calculator to price a complete configuration based on current list price. Our sales engineers can quickly create multiple configurations for comparison based on scenarios and inputs provided by the customer.
- **Transparent:** Customers can always view AWS list pricing in real time. AWS offers a Simple Monthly calculator that is easy to use and includes video tutorials for first-time users. AWS also provides tools that allow customers to evaluate their pricing against other Cloud Service Providers (CSPs).

States may include a wide range of consulting services in their Participating Addendums for all aspects of cloud transition, migration, operations, and optimization.

Rate Card - Hourly

Our hourly ceiling rates for these consulting services are provided below. GTRI does not have different rates for contractor versus customer sites. If necessary, we will negotiate with each purchasing entity for travel costs.

Labor Category	Rate
Enterprise Systems Engineer Level 1	\$117.90
Enterprise Systems Engineer Level 2	\$160.50
Enterprise Systems Engineer Level 3	\$199.80
Enterprise Systems Engineer Level 4	\$239.10
Network Engineer Level 1	\$91.70
Network Engineer Level 2	\$111.35
Network Engineer Level 3	\$147.40
Network Engineer Level 4	\$180.15
Network Engineer Level 5	\$222.70
Wireless Engineer Level 3	\$160.50
Wireless Engineer Level 4	\$186.70
Wireless Engineer Level 5	\$203.05
Project Coordinator Level 1	\$88.45
Project Coordinator Level 2	\$98.25
Project Coordinator Level 3	\$111.35
Project Manager Level 1	\$111.35
Project Manager Level 2	\$137.55
Project Manager Level 3	\$153.95
Network Security Engineer Level 2	\$117.90
Network Security Engineer Level 3	\$147.40
Network Security Engineer Level 4	\$183.40
Network Security Engineer Level 5	\$203.05
Subject Matter Expert Level 1	\$140.85
Subject Matter Expert Level 2	\$160.50

Labor Category	Rate
Subject Matter Expert Level 3	\$199.80
Subject Matter Expert Level 4	\$222.70
Cloud Architect	\$200.16
System Administrator Level 1	\$82.43
System Administrator Level 2	\$100.09
System Administrator Level 3	\$129.52
Cloud Administrator Level 1	\$117.75
Cloud Administrator Level 2	\$141.29
Cloud Administrator Level3	\$164.84
Splunk Professional Services	\$299.00

State of Utah and NASPO Value Point

Cloud Solutions

Solicitation SK18008

July 6, 2018

Executive Summary

Prepared for:
State of Utah
State Contract Analyst
Division of Purchasing

Prepared by:
Global Technology Resources, Inc.
990 S. Broadway, Suite 300
Denver, CO 80209
1-877-603-1984
Fax 1-888-803-6520
www.GTRI.com





TABLE OF CONTENTS

1.	Executive Summary	1
----	-------------------------	---

1. Executive Summary

Global Technology Resources Inc (GTRI) is a strategic IT partner that helps organizations navigate between business needs and technology solutions in order to advance and support the mission, reduce costs, and reduce risks for our clients. GTRI is a strong and growing company who has been in business for two decades, providing information technology infrastructure (IT) products and services. These services are complemented by dedicated account teams, an extensive pre- and post-sales engineering corps, a full-time project management office, enterprise consulting teams and 24x7x365 managed services operations.

GTRI is an experienced, highly skilled solutions integrator. Our team establishes confidence in the early stages of IT transformation by delivering full visibility into legacy systems, utilization and application dependencies. We build trust with a consultative, supplier-neutral approach that matches your global needs with the right mix of services such as cloud, colocation, and IT services. GTRI's end-to-end hybrid IT capabilities enable you to use a single partner for every stage reducing risk, security threats, time to market, downtime, and overall project costs. The result is more business value and agile IT services delivery that propels your business forward.

GTRI operates in five major verticals nationally and international, namely: Federal Government, State and Local Government, Education, Finance and Healthcare. GTRI's solution sets are categorized into the following practice areas: Hybrid IT (including Cloud Solutions), Analytics/Big Data, Collaboration and Security.

GTRI is no stranger to Cloud Solutions. Our early history was in the design and implementation of IT infrastructure related products and services, the natural progression was to expand our business into Cloud Solutions. GTRI began its Cloud Solutions practice in 2012 and has consistently grown that practice in the subsequent years.

GTRI is an Amazon Web Services (AWS) Public Sector Partner and AWS Advanced Technology Partner. Our growing technical staff includes ten certified architects and cloud administrators who design, implement, and provide managed services across a variety of Cloud Solutions within our vertical markets. GTRI employs cloud consultants who possess the Cloud Security Alliance (CSA) Certificate of Cloud Security Knowledge (CCSK) and the (ISC)2 Certified Cloud Security Professional (CCSP) certifications. The NIST cloud specifications are within the body of knowledge contained in these certifications.

As large provider of services to the Federal Government, GTRI currently has obtained a multi-year Authority to Operate (ATO) government workloads in AWS. As part of the maintenance of that ATO, GTRI must adhere to and maintain stringent logical and physical security controls across the entire organization.

GTRI also proudly holds an ISO 9001:2015 quality certification, which reflects our tireless commitment to continual process improvement. These quality control efforts help ensure our customers receive the best services across all verticals we serve.

GTRI appreciates the ability to participate in this Cloud Solutions RFP and are confident you will find our service offerings exceed the business and technical requirements presented.

State of Utah and NASPO Value Point

Cloud Solutions

Solicitation SK18008

July 6, 2018

Mandatory Minimums Response

Prepared for:
State of Utah
State Contract Analyst
Division of Purchasing

Prepared by:
Global Technology Resources, Inc.
990 S. Broadway, Suite 300
Denver, CO 80209
1-877-603-1984
Fax 1-888-803-6520
www.GTRI.com



TABLE OF CONTENTS

1.	Mandatory Minimums	1
1.1.	General Requirements	1
1.2.	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.	1
1.3.	Offeror must at a minimum complete, provide, and maintain a completed CSA STAR Registry Self-Assessment.....	1
1.4.	Sample Service Level Agreements (5.3.4).....	1
1.4.1.	Customer Service	2
1.4.2.	Functional Services SLAs.....	5
1.5.	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. (5.3.5)	7

TABLE OF EXHIBITS

Exhibit 1 –GTRI Reporting Tool.....	1
Exhibit 2 –Sample Customer Service SLAs	4
Exhibit 3 –AWS Customer Service SLAs	5
Exhibit 4 –Functional Services SLAs	7

1. Mandatory Minimums


1.1. General Requirements

GTRI will provide a Usage Report Administrator responsible for the quarterly sales reporting described in the Master Agreement Terms and Conditions. GTRI has a team tasked with submittal of usage reports. This team has experience submitting usage reports to NASPO prime contract holders as well as submitting monthly reports under our federal prime contracts.

Current NASPO Reporting Process:

- Utilize internal GTRI reporting tool to view all orders placed under a NASPO contract within a selected time period
 - o Open the reporting tool
 - o Select a date range based on the required reporting time period
 - o Select view report
 - o Export to excel and enter required information into the required format for report submission

Please see below for a snapshot of a sample internal NASPO (formerly WSCA) report



Cisco Sales Order#	GTRI PO #	GTRI PO Date	Invoice Number	Invoice Date	Customer PO #	Customer PO Date	Customer Entity Type	Customer Name	Customer Bill To Address	Type of Sale: Product, Service or Basic Install & Config (P, S)	Item #
WSCA Report											

Exhibit 1 –GTRI Reporting Tool

1.2. 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.

GTRI agrees to cooperate with NASPO ValuePoint and SciQuest (and any authorized agent or successor entity to SciQuest) with uploading an Offeror's ordering instructions.

1.3. Offeror must at a minimum complete, provide, and maintain a completed CSA STAR Registry Self-Assessment.

Please see Attachment B for the AWS CSA Consensus Assessments Initiative Questionnaire

1.4. Sample Service Level Agreements (5.3.4)

GTRI has several different Service Level Agreements (SLA) to cover all areas of a IaaS and PaaS solution. Our SLA's cover areas such as customer response, reliability, RPO, and RPO for example.

Below are sample SLA's that GTRI provides. We can negotiate SLAs with each Purchasing Entity as needed.

1.4.1. Customer Service

During the on-boarding process, GTRI will work with the customer to establish support interfaces. For all support requests, the partner or the partner's client may request support using either phone, email, portal, or chat communications as follows:

- o GTRI Support Phone – (855) 290-5488
- o GTRI Support Email – support@GTRI.cloud
- o GTRI Support Portal – <https://support.GTRI.cloud>
- o GTRI Support Chat – Find in the support portal.

When you submit an email to our support address (above), this will automatically create a new ticket within our ticketing system. If you wish to utilize the customer portal for our ticketing system you can do so at the address noted above. When you create a new account here you can track and manage your requests online (add comments, resolve, etc.). When you initiate a chat session with our support agents, the chat transcript will be copied into a new ticket when the chat session ends, which you can track and manage through the portal.

GTRI support technicians will communicate with customers using the same methods noted above, and primarily via email through the ticketing system. To maintain a high level of customer service, our agents will focus all client communications through our ticketing system where possible. This is important so that all agents can view the full history of the issue. GTRI agents will work in the best interest of the partner and client to maximize the support experience for the partner and client.

If the customer needs to escalate the priority of a request with GTRI, they should do so by using one of the following methods:

- Reply to the email thread for the request, indicating the details of your escalation request (priority, critical time frames, etc.)
- Log on to the support portal online and add a note to your ticket with the details of your escalation request (priority, critical time frames, etc.)
- Log on to the support portal online and initiate a chat session with a support agent. Ask the agent to add a note to your ticket with the details of your escalation request (priority, critical time frames, etc.)
- Call the support phone line and when you are connected to an agent, ask the agent to locate your ticket (provide ticket number if possible), and ask them to update the ticket with the details of your escalation request (priority, critical time frames, etc.)

GTRI offers three tiers of support as follows:

- **Bronze** - 8x5 – Customers who require daytime support (9am – 5PM ET) excluding weekends.
- **Silver** - 12x7 – Customers who require daytime support with extended service hours (7am – 7pm ET) and including weekends.
- **Gold** - 24x7 – Customers who require coverage 24x7. This level of service is ideal for customers who have high impact applications or are involved in emergency response.

This SLA covers provision and support of the following services:

- GTRI Services Support Team Access

- Maintenance and Support of dependent services
- Development of new or enhanced services
- Advice and consultancy

The SLA remains valid until superseded by a revised agreement, which has been endorsed by relevant signatories from both parties. The agreement will be reviewed annually and applied to all Service Orders associated with this agreement.

Note that the response and resolution times stated within this SLA are sample times only and final SLAs will be negotiated with each Purchasing Entity. GTRI will provide reports to review actual SLA response times. Upon calling the GTRI Services Support Team the call will normally be answered within **60 seconds**. This may be longer if the lines are busy, but this shouldn't exceed **90 seconds**.

Upon calling the GTRI Services Support Team you will be asked to give basic details of the incident. You will be asked for details of the system - so please have this information ready. Once the incident has been logged you will be given a **ticket number**. This ticket number will be e-mailed to you and must be quoted on any future contact. Anyone wishing to speak directly with a technical expert must contact the general GTRI Services Support Team number first.

GTRI service bands Response and Resolution Times

Upon placing a call, you will be asked to assess the **Severity (Business Impact)** of the incident according to the levels indicated in the Business Impact table below. Because of this the NOC Team will allocate a **Priority** to the incident or request. The priority of the incident or request will determine the target response and resolution times as negotiated in your contract (see "Service Levels" below – all service levels are listed for ease of use and understanding.)

If the GTRI Services Support Team cannot immediately resolve the incident on the phone, you will receive a call back from a member of the GTRI Services Support Team or GTRI Services Support technical escalation team according to the priority level.

1.4.1.1. GTRI Response and Resolution Targets

Below are GTRI's sample SLA's for response and resolution

BRONZE Support Provided 8 hours per day, 5 days per week:

Priority	Response Time	Resolution Target
Urgent (1)	Portal (Immediate) / Phone 5 Min	4 hours
High (2)	Portal (Immediate) / Phone 10 Min	8 Hours
Medium (3)	Portal (Immediate) / Phone 30 Min	1 Business Day
Low (4)	Portal (Immediate) / Phone 30 Min	3 Business Days

SILVER Support Provided 12 hours per day, 7 days per week:

Priority	Response Time	Resolution Target
----------	---------------	-------------------

Urgent (1)	Portal (Immediate) / Phone 5 Min	2 hours
High (2)	Portal (Immediate) / Phone 10 Min	4 Hours
Medium (3)	Portal (Immediate) / Phone 30 Min	8 Hours
Low (4)	Portal (Immediate) / Phone 30 Min	2 Business Days

GOLD Support Provided 24 hours per day, 7 days per week:

Priority	Response Time	Resolution Target
Urgent (1)	Portal (Immediate) / Phone 5 Min	1 hours
High (2)	Portal (Immediate) / Phone 10 Min	2 Hours
Medium (3)	Portal (Immediate) / Phone 30 Min	4 Hours
Low (4)	Portal (Immediate) / Phone 30 Min	1 Business Day

Exhibit 2 –Sample Customer Service SLAs
AWS Support

	Basic	Developer	Business	Enterprise
Customer Service – 24x7x365	✓	✓	✓	✓
Support Forums	✓	✓	✓	✓
Documentation, Whitepapers, Best Practice Guides	✓	✓	✓	✓
Access to Technical Support	Support for Health Checks	Email (local business hours)	Phone, Chat, Email (24/7)	Phone, Chat, Email, Technical Account Manager (TAM) (24/7)
Primary Case Handling	Technical Customer Service Associate	Cloud Support Associate	Cloud Support Engineer	Sr. Cloud Support Engineer
Users Who Can Create Technical Support Cases		1	Unlimited (AWS Identity and Access Management)	Unlimited (IAM supported)

			[IAM] supported)	
Response Time		General guidance: < 24 business hours System impaired: < 12 business hours	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 15 minutes
Architecture Support		General Guidance	Contextual Use Case Guidance	Contextual Application Architecture Guidance
Access to Support API			✓	✓
Third-Party Software Support			✓	✓
AWS Trusted Advisor	4 core checks	4 core checks	Full checks	Full checks
Infrastructure Event Management			Contact Us for Pricing	✓
Direct Access to TAM				✓
Architectural Review				✓
Support Concierge				✓
Training				Access to online self-paced labs
Operations Support				Operational reviews, recommendations, and reporting

Exhibit 3 –AWS Customer Service SLAs

1.4.2. Functional Services SLAs

Amazon provides SLAs for their foundational services. AWS frequently adds new SLA commitments or strengthens existing ones as their services mature, so it is worth checking the AWS Web site for the most recent details. As of this writing, here are a summary of the SLAs that AWS offers along with the service credits if the expected service levels are not met.

AWS Service	Service Type	Service Commitment	SLA Credit	Percentage
RDS	Relational Database	99.95%	Less than 99.95% but equal to or greater than 99.0%	10 %
RDS	Relational Database	99.95%	Less than 99.0%	25 %
S3	Storage	99.9%	Equal to or greater than 99.0% but less than 99.9%	10 %
S3	Storage	99.9%	Less than 99.0%	25 %
EC2	Compute	99.95%	Less than 99.95% but equal to or greater than 99.0%	10 %
EC2	Computer	99.95%	Less than 99.0%	30 %
DynamoDB (Global)	NoSQL Database	99.999%	Less than 99.999% but equal to or greater than 99%	10%
DynamoDB (Global)	NoSQL Database	99.999%	Less than 99.0%	30%
DynamoDB (Local)	NoSQL Database	99.99%	Less than 99.99% but equal to or greater than 99%	10%
DynamoDB (Local)	NoSQL Database	99.99%	Less than 99%	10%
Route 53	DNS + Global Load Balancing	100 %	5 – 30 minutes in a Billing Cycle	1 day Service Credit

Route 53	DNS + Global Load Balancing	100 %	31 minutes – 4 hours in a Billing Cycle	7 days Service Credit
Route 53	DNS + Global Load Balancing	100 %	More than 4 hours in a Billing Cycle	30 days Service Credit
CloudFront	Content Distribution Network	99.9%	Equal to or greater than 99% but less than 99.9%	10%
CloudFront	Content Distribution Network	99.9%	Less than 99%	30%
AWS Shield Advanced	DDos Mitigation	100%	Unavailable during a 24 hour period	1 day service credit

Exhibit 4 –Functional Services SLAs

1.5. 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. (5.3.5)

Table below shows Cloud services available from GTRI including the services from AWS that are implemented by GTRI.

Cloud Services Information	Service Description
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 customers use to build and host their software systems. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing customers to quickly scale capacity, both up and down, as their computing requirements change. Amazon EC2 instances in Amazon Virtual Private Cloud (Amazon VPC) offer native support for the IPv6 protocol.
Amazon ECS	Amazon Elastic Container Service (Amazon ECS) is a highly scalable, high performance container management service that supports Docker containers and allows customers to easily run applications on a managed cluster of Amazon EC2 instances. With simple API calls, customers can launch and stop Docker-enabled applications, query the complete state of a cluster, and access many familiar features. Customers can use Amazon ECS to schedule the placement of

Cloud Services Information	Service Description
	containers across a cluster based on resource needs and availability requirements. Customers can also integrate their own scheduler or third-party schedulers to meet business or application-specific requirements. Amazon ECS can send container instances and task state changes to CloudWatch Events.
Amazon EKS (in Preview)	Amazon Elastic Container Service for Kubernetes (Amazon EKS) is a managed service that makes it easy for customers to run Kubernetes on AWS without needing to install and operate their own Kubernetes clusters. With Amazon EKS, upgrades and high availability are managed by AWS. Amazon EKS runs three Kubernetes masters across three Availability Zones in order to ensure high availability. Amazon EKS automatically detects and replaces unhealthy masters, and it provides automated version upgrades and patching for the masters.
AWS Fargate	AWS Fargate is a technology for deploying and managing containers, which frees customers from having to manage any of the underlying infrastructure. With AWS Fargate, customers no longer have to provision, configure, and scale clusters of virtual machines to run containers. Amazon ECS uses containers provisioned by Fargate to automatically scale, load balance, and manage scheduling of containers for availability, providing an easier way to build and operate containerized applications. AWS Fargate will also support Amazon EKS in 2018.
Amazon ECR	Amazon Elastic Container Registry (Amazon ECR) is a fully managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images. Amazon ECR is integrated with Amazon ECS, simplifying development-to-production workflows. Amazon ECR eliminates the need for customers to operate their own container repositories or worry about scaling the underlying infrastructure. Amazon ECR hosts images in a highly available and scalable architecture, allowing customers to reliably deploy containers for applications.
AWS Elastic Beanstalk	AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS. Customers can simply upload their code and AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto scaling to application health monitoring. At the same time, customers retain full control over the AWS resources powering their application and can access the underlying resources at any time.
AWS Lambda	AWS Lambda lets customers run code without provisioning or managing servers. With AWS Lambda, customers can run code for virtually any type of application or back-end service—all with zero administration. Just upload code and AWS Lambda takes care of everything required to run and scale the code with high availability. Customers can set up their code to automatically trigger from other AWS Cloud services or call it directly from any web or mobile application. Customers can also develop their AWS Lambda functions in C# using the .NET Core 1.0 runtime.

Cloud Services Information	Service Description
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.
Amazon Lightsail	Amazon Lightsail is designed to be the easiest way to launch and manage a virtual private server with AWS. Amazon Lightsail plans include everything customer need to jumpstart their project—a virtual machine, SSD-based storage, data transfer, DNS management, and a static IP—for a low, predictable price. Amazon Lightsail virtual private servers run on the same highly available and reliable AWS cloud infrastructure used by millions of customers.
AWS Batch	AWS Batch is a set of batch management capabilities that enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (e.g., CPU or memory optimized instances) based on the volume and specific resource requirements of the batch jobs submitted. With AWS Batch, there is no need to install and manage batch computing software or server clusters, allowing customers to instead focus on analyzing results and solving problems. AWS Batch plans, schedules, and executes batch computing workloads using Amazon EC2 and Spot Instances.
AWS Serverless Application Repository	The AWS Serverless Application Repository is a collection of serverless applications and serverless application components provided by AWS and other AWS partners and developers. With a growing selection of applications to choose from, the AWS Serverless Application Repository makes it easy to get started with the AWS serverless platform. The AWS Serverless Application Repository includes applications for Alexa Skills, chatbots, data processing, IoT, real time stream processing, web and mobile back-ends, social media trend analysis, image resizing, and more from publishers on AWS.
Storage	
Amazon S3	Amazon Simple Storage Service (Amazon S3) is storage for the Internet. Customers can use Amazon S3 to store and retrieve any amount of data, at any time, from anywhere on the web. Amazon S3 offers a range of storage classes designed for different use cases: Amazon S3 Standard for general-purpose storage of frequently accessed data, Amazon S3 Standard – Infrequent Access (Standard – IA) for long-lived, but less frequently accessed data, and Amazon Glacier for long-term archive.
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, customers can store their data cost effectively for months, years, or even decades. Amazon Glacier enables customers to offload the administrative burdens of operating and scaling storage to AWS, so customers don't have to worry about capacity

Cloud Services Information	Service Description
	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.
Amazon EFS	Amazon Elastic File System (Amazon EFS) provides simple, scalable file storage for use with Amazon EC2 instances in the AWS Cloud. Amazon EFS is easy to use and offers a simple interface that allows customers to create and configure file systems quickly and easily. With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as files are added and removed, so applications have the storage they need, when they need it. Amazon EFS also allows customers to access file data from on-premises data centers. On-premises servers can move file data to and from an Amazon EFS file systems when connected to an Amazon VPC with AWS Direct Connect, allowing customers to migrate data sets to Amazon EFS, enable cloud bursting scenarios, or back up on-premises data to Amazon EFS.
AWS Import/Export Disk	AWS Import/Export Disk accelerates moving large amounts of data into and out of the AWS cloud using portable storage devices for transport. AWS Import/Export Disk transfers 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 Disk is often faster than Internet transfer and more cost effective than upgrading 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 an on-premises IT environment and AWS's storage infrastructure. AWS Storage Gateway also provides a virtual on-premises file server, which enables customers to store and retrieve Amazon S3 objects through standard file storage protocols.
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 2016, PostgreSQL, and MariaDB. Amazon RDS for Oracle Database Instance supports outbound network access and can send emails using an utl_smtp package, plus communicate with external Web Servers and TCP/IP-based servers using utl_http and utl_tcp packages.

Cloud Services Information	Service Description
Amazon Aurora	Amazon Aurora is a MySQL-compatible relational database engine that combines the speed and availability of high-end commercial databases with the simplicity and cost effectiveness of open-source databases. Amazon Aurora provides up to five times better performance than MySQL with the security, availability, and reliability of a commercial database at one-tenth the cost. Amazon Aurora's PostgreSQL support allows customers to get up to several times better performance than the typical PostgreSQL database.
Amazon DynamoDB	Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Customers 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 data using your existing business intelligence tools. Customers can start small for just \$0.25 per hour with no commitments or upfront costs and scale to a petabyte or more for \$1,000 per terabyte per year, less than one-tenth the cost of most other data warehousing solutions. For customer convenience, multibyte (UTF-8) characters can be used in Amazon Redshift table, column, and other database object names.
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. Amazon ElastiCache supports two open-source, in-memory engines: Redis and Memcached.
Migration	
AWS Migration Hub	AWS Migration Hub provides a single place to discover existing servers and track the status of each application migration. It gives customers better visibility into their application portfolio and streamlines migration tracking. If a customer plans to migrate some or all of their workloads to AWS using tools like AWS Server Migration Service, AWS Database Migration Service, or partner tools, then they should consider using AWS Migration Hub. If a customer is rehosting servers and/or replatforming databases in AWS, then the AWS Migration Hub can reduce time spent determining current status and next steps.
AWS Application Discovery Service	AWS Application Discovery Service helps customers quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centers, their associated dependencies, and their performance profile. AWS Application Discovery Service automatically collects configuration and usage data from server, storage, and networking equipment to

Cloud Services Information	Service Description
	develop a list of applications, how they perform, and how they are interdependent.
AWS Database Migration Service	AWS Database Migration Service helps customers migrate databases to AWS easily and securely. The source database remains fully operational during the migration. Customers can migrate data to and from most widely used commercial and open-source databases. The service supports homogenous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora or Microsoft SQL Server to MySQL.
AWS SMS	AWS Server Migration Service (AWS SMS) is an agentless service which makes it easier and faster for customers to migrate thousands of on-premises workloads to AWS. AWS SMS allows customers to automate, schedule, and track incremental replications of live server volumes, making it easier to coordinate large-scale server migrations.
AWS Snowball	AWS Snowball accelerates transferring large amounts of data between the cloud and portable storage devices that a customer mails to us. AWS transfers data directly onto and off of the customer's storage devices using Amazon's high-speed internal network. The data load typically begins the next business day after the storage device arrives at AWS. After the data export or import completes, we return the storage device. For large data sets, AWS Snowball is significantly faster than Internet transfer and more cost effective than upgrading connectivity.
AWS Snowball Edge	AWS Snowball Edge is a 100TB data transfer device with on-board storage and compute. Customers can use AWS Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local data sets, or to support independent local workloads in remote locations. AWS Snowball Edge connects to existing applications and infrastructure using standard storage interfaces. AWS Snowball Edge can cluster together to form a local storage tier and process data on-premises, helping ensure that applications continue to run even when they are not able to access the cloud.

Cloud Services Information	Service Description
AWS Snowmobile	AWS Snowmobile is an Exabyte-scale data transfer service that can move extremely large amounts of data to AWS in a fast, secure, and cost-effective manner. Transfer up to 100PB per AWS Snowmobile—a 45-foot long ruggedized shipping container pulled by a semi-trailer truck—to move massive volumes of data to the cloud. All data is encrypted with 256-bit encryption with managed encryption keys using AWS Key Management Service (AWS KMS). AWS Snowmobile includes GPS tracking, alarm monitoring, 24/7 video surveillance and an optional escort security vehicle while in transit. After the data is loaded, AWS Snowmobile is driven back to AWS where the data is imported into Amazon S3 or Amazon Glacier.
Networking and Content Delivery	
Amazon VPC	Amazon Virtual Private Cloud (Amazon VPC) enables customers to launch AWS resources into a virtual network that they've defined. This virtual network closely resembles a traditional network that a customer would operate in their own data center, with the benefits of using the scalable infrastructure of AWS.
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 (for example, streaming video) to end users with low latency, high data transfer speeds, and no commitments. Amazon CloudFront has a type of edge location called Regional Edge Cache that further improves performance for viewers. Regional Edge Caches, in addition to improving performance, help reduce the load on origin resources, minimizing operational burden associated with scaling origin and reducing origin costs. Regional Edge Caches are turned on by default for Amazon CloudFront distributions; customers do not need to make any changes to their distributions to take advantage of this feature. There are also no additional charges to use this feature.
Amazon Route 53	Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed as an extremely reliable and cost-effective way to route users to Internet applications by translating site names into IP addresses. Amazon Route 53 also connects user requests to infrastructure running in AWS, such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets. Amazon Route 53 health checks support endpoints with IPv6 addresses and DNS queries over IPv6.
Amazon API Gateway	Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. With a few clicks in the AWS Management Console, customers can create an API that acts as a "front door" for applications to access data, business logic, or functionality from back-end services, such as workloads running on Amazon EC2, code running on AWS Lambda, or any Web application. Amazon API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management.

Cloud Services Information	Service Description
	Developers can now use custom authorizers on Amazon API Gateway to return additional fields in their authorization response. The custom authorizers can be used to authorize API requests to their backend using bearer token strategies such as OAuth.
AWS Direct Connect	AWS Direct Connect makes it easy to establish a dedicated network connection from a customer's premises to AWS. Using AWS Direct Connect, a customer can establish private connectivity between AWS and a data center, office, or colocation environment, which in many cases can reduce network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.
Developer Tools	
AWS CodeStar	AWS CodeStar is a cloud-based development service that provides the tools that customers need to quickly develop, build, and deploy applications on AWS. With AWS CodeStar, customers can set up their entire continuous delivery toolchain in minutes, allowing them to start releasing code faster. AWS CodeStar makes it easy for a whole team to work together securely, with built-in role-based policies that allow customers to easily manage access and add owners, contributors, and viewers to projects. Each AWS CodeStar project comes with a unified project dashboard and integration with Atlassian JIRA software, a third-party issue tracking and project management tool. With the AWS CodeStar project dashboard, customers can easily track their entire software development process, from a backlog work item to production code deployment.
AWS CodeBuild	AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. With AWS CodeBuild, customers don't need to provision, manage, and scale their own build servers. AWS CodeBuild scales continuously and processes multiple builds concurrently, so builds are not left waiting in a queue. Customers can get started quickly by using prepackaged build environments or create custom build environments that use their own build tools.
AWS CodeDeploy	AWS CodeDeploy is a service that automates code deployments to Amazon EC2 instances. AWS CodeDeploy makes it easier for customers to rapidly release new features, helps avoid downtime during deployment, and handles the complexity of updating applications. Customers can use AWS CodeDeploy to automate deployments, eliminating the need for error-prone manual operations, and the service scales with infrastructure so that customers 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 customers to operate their own source control system or worry about scaling its infrastructure. Customers can use AWS CodeCommit to store anything from

Cloud Services Information	Service Description
	code to binaries, and it supports the standard functionality of Git, allowing it to work seamlessly with existing Git-based tools.
AWS CodePipeline	AWS CodePipeline is a continuous delivery and release automation service that aids smooth deployments. Customers can design their development workflow for checking in code, building the code, deploying the application into staging, testing it, and releasing it to production. Customers can integrate third-party tools into any step of the release process, or they can use AWS CodePipeline as an end-to-end solution.
AWS X-Ray	AWS X-Ray helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-Ray, customers can understand how their application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. X-Ray provides an end-to-end view of requests as they travel through the application, and shows a map of the application's underlying components. Customers can use X-Ray to analyze both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.
Management Tools	
Amazon CloudWatch	Amazon CloudWatch is a service that enables customers to collect, view, and analyze metrics. Amazon CloudWatch lets customers programmatically retrieve monitoring data, view graphs, and set alarms to help troubleshoot, spot trends, and take automated action based on the state of a cloud environment.
Amazon CloudWatch Events	Monitor and automate action on Amazon EBS snapshots using Amazon CloudWatch Events. Amazon CloudWatch Events provide a stream of events describing changes to AWS resources. Amazon EBS CloudWatch Events fire when a snapshot completes or when a snapshot has been shared with someone, which allows customers to automate and streamline their data backup workflows by eliminating the need to poll the snapshot API to track snapshot status.
AWS CloudTrail	With AWS CloudTrail, customers can get a history of AWS API calls for their account, including API calls made via the AWS Management Console, the AWS Software Development Kits (SDKs), the command line tools, and higher-level AWS Cloud services. Customers can also identify which users and accounts called AWS APIs for services that support AWS CloudTrail, the source IP address that the calls were made from, and when the calls occurred. Customers can integrate AWS CloudTrail into applications using the API, automate trail creation for their organization, check the status of their trails, and control how administrators turn AWS CloudTrail logging on and off. AWS CloudTrail supports Amazon S3 Data Events. Customers can record all API actions on Amazon S3 objects and receive detailed information.
AWS Config	AWS Config is a fully managed service that provides customers with an AWS resource inventory, configuration history, and configuration change notifications

Cloud Services Information	Service Description
	to enable security and governance. With AWS Config customers can discover existing AWS resources, export a complete inventory of their AWS resources with all configuration details, and determine how a resource was configured at any point in time. Customers can use AWS Config to record configuration changes to software on their Amazon EC2 instances as well as Virtual Machines (VMs) or servers in their on-premises environment. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.
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. Customers can use the sample templates for AWS CloudFormation or create their own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run the application.
AWS OpsWorks	AWS OpsWorks provides a simple and flexible way to create and manage stacks and applications. With AWS OpsWorks, customers can provision AWS resources, manage their configuration, deploy applications to those resources, and monitor their health.
AWS Trusted Advisor	AWS Trusted Advisor acts like a customized cloud expert, and it helps customers provision their resources by following best practices. AWS Trusted Advisor inspects an AWS environment and finds opportunities to save money, improve system performance and reliability, or help close security gaps.
AWS Systems Manager	AWS Systems Manager allows customers to centralize operational data from multiple AWS services and automate tasks across AWS resources. Customers can create logical groups of resources such as applications, different layers of an application stack, or production versus development environments. With AWS Systems Manager, customers can select a resource group and view its recent API activity, resource configuration changes, related notifications, operational alerts, software inventory, and patch compliance status. AWS Systems Manager provides a central place to view and manage AWS resources, so customers can have complete visibility and control over their operations.
AWS Service Catalog	AWS Service Catalog allows organizations to create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows customers to centrally manage commonly deployed IT services, and helps them achieve consistent governance and meet compliance requirements, while enabling users to quickly deploy only the approved IT services they need.
AWS Management Console	The AWS Management Console provides customers with the ability to access and manage AWS Cloud services through a simple and intuitive web-based user interface. Customers can also use the AWS Console mobile app to quickly view resources on the go.

Cloud Services Information	Service Description
AWS Command Line Tool	The AWS Command Line Interface (CLI) is a unified tool to manage AWS Cloud services. With just one tool to download and configure, customers can control multiple AWS services 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 S3.
AWS Personal Health Dashboard	Customers can now receive notification and remediation guidance when AWS is experiencing events that may impact them. Available to all AWS customers, AWS Personal Health Dashboard provides a personalized view into the performance and availability of the AWS services a customer is using, as well as alerts that are automatically triggered by changes in the health of those services.
Amazon Elastic Transcoder	Amazon Elastic Transcoder lets customers convert media files that they have stored in Amazon S3 into media files in the formats required by consumer playback devices. For example, customers 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 Kinesis Video Streams	Amazon Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), and other processing. Kinesis Video Streams automatically provisions and elastically scales all the infrastructure needed to ingest streaming video data from millions of devices. It also durably stores, encrypts, and indexes video data in streams, and allows customers to access their data through easy-to-use APIs. Kinesis Video Streams enables customers to quickly build computer vision and ML applications through integration with Amazon Rekognition Video and libraries for ML frameworks such as Apache MxNet, TensorFlow, and OpenCV.
AWS Elemental MediaConvert	AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It allows customers to easily create video-on-demand (VOD) content for broadcast and multiscreen delivery at scale. The service combines advanced video and audio capabilities with a simple web services interface and pay-as-you-go pricing. With AWS Elemental MediaConvert, customers can focus on delivering compelling media experiences without having to worry about the complexity of building and operating their own video processing infrastructure.
AWS Elemental MediaLive	AWS Elemental MediaLive is a broadcast-grade live video processing service. It lets customers create high-quality video streams for delivery to broadcast televisions and internet-connected multiscreen devices, like connected TVs, tablets, smart phones, and set-top boxes. The service works by encoding live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to viewers. With AWS Elemental MediaLive, customers can easily set up streams for both live events and 24x7 channels with advanced broadcasting features, high availability, and pay-as-you-go pricing. AWS Elemental MediaLive lets customers focus on creating compelling live video experiences for their viewers without the

Cloud Services Information	Service Description
	complexity of building and operating broadcast-grade video processing infrastructure.
AWS Elemental MediaPackage	AWS Elemental MediaPackage reliably prepares and protects video for delivery over the Internet. From a single video input, AWS Elemental MediaPackage creates video streams formatted to play on connected TVs, mobile phones, computers, tablets, and game consoles. It makes it easy to implement popular video features for viewers (start-over, pause, rewind, etc.), like those commonly found on DVRs. AWS Elemental MediaPackage can also protect customer content using Digital Rights Management (DRM). AWS Elemental MediaPackage scales automatically in response to load, so viewers will always get a great experience without customers having to accurately predict in advance the capacity they'll need.
AWS Elemental MediaStore	AWS Elemental MediaStore is an AWS storage service optimized for media. It gives customers the performance, consistency, and low latency required to deliver live and on-demand video content. AWS Elemental MediaStore acts as the origin store in a video workflow. Its high performance capabilities meet the needs of the most demanding media delivery workloads, combined with long-term, cost-effective storage.
AWS Elemental MediaTailor	AWS Elemental MediaTailor lets video providers insert individually targeted advertising into their video streams without sacrificing broadcast-level quality-of-service. With AWS Elemental MediaTailor, viewers of live or on-demand video each receive a stream that combines content with ads personalized to them. But unlike other personalized ad solutions, with AWS Elemental MediaTailor entire stream—video and ads—are delivered with broadcast-grade video quality to improve the experience for viewers.
Machine Learning	
Amazon SageMaker	Amazon SageMaker is a fully-managed service that enables developers and data scientists to quickly and easily build, train, and deploy machine learning models at any scale. Amazon SageMaker removes the barriers that typically slow down developers who want to use machine learning. Amazon SageMaker includes modules that can be used together or independently to build, train, and deploy machine learning models.
Amazon Comprehend	Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. Amazon Comprehend identifies the language of the text; extracts key phrases, places, people, brands, or events; understands how positive or negative the text is; and automatically organizes a collection of text files by topic. The service constantly learns and improves from a variety of information sources, including Amazon.com product descriptions and consumer reviews—one of the largest natural language data sets in the world—to keep pace with the evolution of language.

Cloud Services Information	Service Description
Amazon Lex	Amazon Lex is a service for building conversational interfaces into any application using voice and text. Lex provides the advanced deep learning functionalities of Automatic Speech Recognition (ASR) for converting speech to text, and Natural Language Understanding (NLU) to recognize the intent of the text, to enable customers to build applications with highly engaging user experiences and lifelike conversational interactions. With Amazon Lex, the same deep learning technologies that power Amazon Alexa are now available to any developer, enabling customers to quickly and easily build sophisticated, natural language, conversational bots (“chatbots”).
Amazon Polly	Amazon Polly is a service that turns text into lifelike speech. Polly lets customers create applications that talk, enabling customers to build entirely new categories of speech-enabled products. Amazon Polly is an Amazon AI service that uses advanced deep learning technologies to synthesize speech that sounds like a human voice. Amazon Polly includes 47 lifelike voices spread across 24 languages, so customers can select the ideal voice and build speech-enabled applications that work in many different countries.
Amazon Rekognition	Amazon Rekognition is a service that makes it easy to add image analysis to applications. With Rekognition, customers can detect objects, scenes, and faces in images. Customers can also search and compare faces. Amazon Rekognition’s API enables customers to quickly add sophisticated deep learning-based visual search and image classification to their applications. Amazon Rekognition uses deep neural network models to detect and label thousands of objects and scenes in images, and we are continually adding new labels and facial recognition features to the service.
Amazon Machine Learning	Amazon Machine Learning makes it easy for developers to build smart applications, including applications for fraud detection, demand forecasting, targeted marketing, and click prediction. The powerful algorithms of Amazon Machine Learning create Machine Learning (ML) models by finding patterns in existing data. The service uses these models to process new data and generate predictions for an application.
Amazon Translate (in Preview)	Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. Neural machine translation is a form of language translation automation that uses machine learning and deep learning models to deliver more accurate and more natural sounding translation than traditional statistical and rule-based translation algorithms. Amazon Translate allows customers to easily translate large volumes of text efficiently, and to localize websites and applications for international users.
Amazon Transcribe	Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for developers to add speech to text capability to their applications. Using the Amazon Transcribe API, customers can analyze audio files stored in Amazon S3 and have the service return a text file of the transcribed speech. Amazon Transcribe can be used for lots of common applications, including the transcription of customer service calls and generating subtitles on audio and

Cloud Services Information	Service Description
	video content. The service can transcribe audio files stored in common formats, like WAV and MP3, with time stamps for every word so customers can easily locate the audio in the original source by searching for the text. Amazon Transcribe is continually learning and improving to keep pace with the evolution of language.
AWS DeepLens	AWS DeepLens is the world's first deep-learning enabled video camera for developers of all skill levels to grow their machine learning skills through hands-on computer vision tutorials, example code, and pre-built models. Customers can run the included sample projects as is, connect them with other AWS Cloud services, train a model in Amazon SageMaker and deploy it to AWS DeepLens, or extend the functionality by triggering a lambda function when an action takes place. Customers can even apply more advanced analytics on the cloud using Amazon Kinesis Video Streams and Amazon Rekognition video. AWS DeepLens provides the building blocks for machine learning needs.
AWS Deep Learning AMIs	The AWS Deep Learning AMIs provide machine learning practitioners and researchers with the infrastructure and tools to accelerate deep learning in the cloud, at any scale. Customers can quickly launch Amazon EC2 instances pre-installed with popular deep learning frameworks such as Apache MXNet and Gluon, TensorFlow, Microsoft Cognitive Toolkit, Caffe, Caffe2, Theano, Torch, Pytorch, and Keras to train sophisticated, custom AI models, experiment with new algorithms, or to learn new skills and techniques.
Apache MXNet on AWS	Apache MXNet is a fast and scalable training and inference framework with an easy-to-use, concise API for machine learning. MXNet includes the Gluon interface that allows developers of all skill levels to get started with deep learning on the cloud, on edge devices, and on mobile apps. In just a few lines of Gluon code, customers can build linear regression, convolutional networks and recurrent LSTMs for object detection, speech recognition, recommendation, and personalization.
TensorFlow on AWS	TensorFlow™ enables developers to quickly and easily get started with deep learning in the cloud. The framework has broad support in the industry and has become a popular choice for deep learning research and application development, particularly in areas such as computer vision, natural language understanding and speech translation. Customers can get started using TensorFlow on AWS by launching the AWS Deep Learning AMI which comes bundled with TensorFlow, as well as other popular deep learning frameworks such as Apache MXNet and Gluon, Caffe, Caffe2, Theano, Torch, Keras, and the Microsoft Cognitive Toolkit.
Analytics	
Amazon Athena	Amazon Athena, a serverless query service that makes it easy to analyze data directly in Amazon S3 using standard SQL. With a few clicks in the AWS Management Console, customers can point Athena at their data stored in S3 and begin using standard SQL to run ad-hoc queries and get results in seconds. With

Cloud Services Information	Service Description
	Athena, there are no clusters to manage and tune, there is no infrastructure to setup, and customers pay only for the queries they run. Athena scales automatically, executing queries in parallel, so results are fast, even with large datasets and complex queries.
Amazon EMR	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. Customers configure policies to automatically add (scale out) and terminate (scale in) nodes in an Amazon EMR cluster programmatically.
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 a website. Amazon CloudSearch enables customers to search large collections of data such as web pages, document files, forum posts, or product information. With Amazon CloudSearch, customers can quickly add search capabilities to their websites without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As the volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet customers' needs.
Amazon Elasticsearch Service	Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. Customers can set up and configure their Amazon Elasticsearch cluster in minutes from the AWS Management Console. With Amazon Elasticsearch Service, customers get direct access to the Elasticsearch open-source API so that code and applications already being used with existing Elasticsearch environments will work seamlessly.
Amazon Kinesis	Amazon Kinesis makes it easy to collect, process, and analyze real-time, streaming data so that customers can get timely insights and react quickly to new information. Amazon Kinesis offers key capabilities to cost-effectively process streaming data at any scale, along with the flexibility to choose the tools that best suit the requirements of the application. Amazon Kinesis enables customers to process and analyze data as it arrives and respond instantly instead of having to wait until all of the data is collected before the processing can begin. Amazon Kinesis offers four capabilities: Kinesis Video Streams makes it easy to securely stream video from connected devices to AWS for analytics, machine learning (ML), and other processing. Kinesis Data Streams enables customers to build custom, real-time applications that process data streams using popular stream processing frameworks. Kinesis Data Firehose is the easiest way to capture, transform, and load data streams into AWS data stores for near real-time analytics with existing business intelligence tools.

Cloud Services Information	Service Description
	Kinesis Data Analytics is the easiest way to process data streams in real time with SQL without having to learn new programming languages or processing frameworks.
Amazon QuickSight	Amazon QuickSight is a fast, cloud-powered business analytics service that makes it easy to build visualizations, perform ad-hoc analysis, and quickly get business insights from data. Using our cloud-based service customers can easily connect to data, perform advanced analysis, and create stunning visualizations and rich dashboards that can be accessed from any browser or mobile device. QuickSight enables organizations to scale their business analytics capabilities to hundreds of thousands of users, and delivers fast and responsive query performance by using a robust in-memory engine (SPICE).
AWS Data Pipeline	AWS Data Pipeline is a web service that helps customers 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, customers can regularly access their 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 Glue	AWS Glue is a fully managed ETL service that makes it easy to move data stores. AWS Glue simplifies and automates the difficult and time consuming data discovery, conversion, mapping, and job scheduling tasks. AWS Glue is integrated with Amazon S3, Amazon RDS, and Amazon Redshift and can connect to any JDBC-compliant data store. AWS Glue automatically crawls data sources, identifies data formats, and then suggests schemas and transformations, so customers don't have to spend time hand-coding data flows. AWS Glue is coming soon.
Security & Identity	
AWS IAM	AWS Identity and Access Management (IAM) is a web service that enables AWS customers to manage users and user permissions. 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 IAM, customers can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access. AWS IAM roles and the AWS Security Token Service (STS) can be used to set up cross-account access between AWS accounts. When an IAM role in another AWS account is used to obtain cross-account access to services and resources in that account, AWS CloudTrail logs the cross-account activity.
Amazon Cloud Directory	Amazon Cloud Directory makes it easy to create highly flexible, scalable, and cost effective directories. With Cloud Directory, customers can create directories for a variety of use cases, such as organizational charts, course catalogs, and device registries. While traditional directory solutions, such as Active Directory Lightweight Directory Services (AD LDS) and other LDAP-based directories, limit customers to a single hierarchy, Cloud Directory offers customers the

Cloud Services Information	Service Description
	flexibility to create directories with hierarchies that span multiple dimensions. Amazon Cloud Directory offers an extensible schema, designed to be shared across applications, and automatically scales to hundreds of millions of objects.
Amazon Cognito	Amazon Cognito lets customers easily add user sign-up and sign-in to their mobile and web apps. With Amazon Cognito, customers also have the options to authenticate users through social identity providers such as Facebook, Twitter, or Amazon, with SAML identity solutions, or by using their own identity system. In addition, Amazon Cognito enables customers to save data locally on users' devices, allowing applications to work even when the devices are offline. Customers can then synchronize data across users' devices so that their app experience remains consistent regardless of the device they use. With Amazon Cognito, customers can focus on creating great app experiences instead of worrying about building, securing, and scaling a solution to handle user management, authentication, and sync across devices.
Amazon GuardDuty	Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious or unauthorized behavior to help customers protect their AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. GuardDuty also detects potentially compromised instances or reconnaissance by attackers. GuardDuty identifies suspected attackers through integrated threat intelligence feeds and uses machine learning to detect anomalies in account and workload activity. When a potential threat is detected, the service delivers a detailed security alert to the GuardDuty console and AWS CloudWatch Events. This makes alerts actionable and easy to integrate into existing event management and workflow systems.
Amazon Inspector	Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity.
Amazon Macie	Amazon Macie is an AI-powered security service that helps customers prevent data loss by automatically discovering, classifying, and protecting sensitive data stored in AWS. Amazon Macie uses ML to recognize sensitive data such as Personally Identifiable Information (PII) or intellectual property, assigns a business value, and provides visibility into where this data is stored and how it is being used in the organization. Amazon Macie continuously monitors data access activity for anomalies and delivers alerts when it detects risk of unauthorized access or inadvertent data leaks.
AWS Certificate Manager	AWS Certificate Manager is a service that lets customers easily provision, manage, and deploy Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS Cloud services. SSL/TLS certificates are used to

Cloud Services Information	Service Description
	secure network communications and establish the identity of websites over the Internet.
AWS CloudHSM	The AWS CloudHSM service helps customers meet corporate, contractual, and regulatory compliance requirements for data security by using dedicated Hardware Security Module (HSM) appliances within the AWS Cloud. With AWS CloudHSM, customers control the encryption keys and cryptographic operations performed by the HSM.
AWS Directory Service	AWS Directory Service is a managed service that allows customers to connect their AWS resources with an existing on-premises Microsoft Active Directory or to set up a new, standalone directory in the AWS Cloud. Connecting to an on-premises directory is easy, and when this connection is established, all users can access AWS resources and applications with their existing corporate credentials.
AWS KMS	AWS Key Management Service (AWS KMS) is a managed service that makes it easy for customers to create and control the encryption keys used to encrypt their data. AWS KMS uses HSMs to protect the security of 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 customers with logs of all key usage.
AWS Organizations	AWS Organizations offers policy-based management for multiple AWS accounts. With Organizations, customers can create groups of accounts and then apply policies to those groups. Organizations enables customers to centrally manage policies across multiple accounts, without requiring custom scripts and manual processes. Using AWS Organizations, customers can create Service Control Policies (SCPs) that centrally control AWS service use across multiple AWS accounts. Customers can also use Organizations to help automate the creation of new accounts through APIs. Organizations helps simplify the billing for multiple accounts by enabling customers to setup a single payment method for all the accounts in their organization through consolidated billing.
AWS Shield	AWS Shield is a managed Distributed Denial of Service (DDoS) protection service that protects web applications on AWS from DDoS attacks. AWS Shield Standard is available to all AWS customers at no additional cost, and protects applications from the most common, frequently occurring DDoS attacks. AWS Shield Advanced is available for a higher level of protection for applications on Elastic Load Balancing, Amazon CloudFront, and Amazon Route 53. With AWS Shield Advanced, customers also get access to AWS WAF at no additional cost. AWS Shield Advanced also provides 24x7 access to the DDoS Response Team (DRT) and DDoS Cost Protection, to protect their AWS bill against usage fee surges during mitigation of a DDoS attack.
AWS WAF	AWS WAF is a web application firewall that helps protect web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. Also, AWS WAF includes a full-featured Application Programming Interface (API) that customers

Cloud Services Information	Service Description
	can use to automate the creation, deployment, and maintenance of web security rules.
AWS Artifact	AWS Artifact, available in the AWS Management Console, is a self-service audit artifact retrieval portal that provides our customers with on-demand access to AWS's compliance documentation. AWS Artifact provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA).
Mobile Services	
AWS Mobile Hub	AWS Mobile Hub lets customers easily add and configure features for mobile apps, including user authentication, data storage, back-end logic, push notifications, content delivery, and analytics. AWS Mobile Hub gives customers easy access to testing on real devices, as well as analytics dashboards to track app usage, all from a single, integrated console. AWS Mobile Hub includes integration with Amazon Lex, so customers can build mobile apps that use speech and text to a mobile app. Amazon Lex lets customers use the technology that powers Amazon Alexa to create engaging speech- and text-based conversational bots in their own app. AWS Mobile Hub's SaaS Connectors make it easy to securely access data in third-party enterprise SaaS applications from a mobile apps.
AWS Device Farm	Improve the quality of iOS, Android, and web applications by testing against real phones and tablets in the AWS Cloud. AWS Device Farm is an app testing service that enables customers to test their iOS, Android, and Fire OS apps on real, physical phones and tablets that are hosted by AWS. The service allows customers to upload their own tests or use built-in, script-free compatibility tests.
Amazon Mobile Analytics	With Amazon Mobile Analytics, customers can measure app usage and app revenue. By tracking key trends such as new vs. returning users, app revenue, user retention, and custom in-app behavior events, customers can make data-driven decisions to increase engagement and monetization for their app. Customers can view key charts in the Mobile Analytics console and automatically export app event data to Amazon S3 and Amazon Redshift to run custom analysis.
AWS Mobile SDK	The AWS Mobile SDK helps customers build high quality mobile apps quickly and easily. It provides easy access to a range of AWS services, including Amazon Lambda, Amazon S3, Amazon DynamoDB, Amazon Mobile Analytics, Amazon Machine Learning, Elastic Load Balancing, Auto Scaling, and more. The AWS Mobile SDK includes libraries, code samples, and documentation for

Cloud Services Information	Service Description
	iOS, Android, Fire OS, and Unity so customers can build apps that deliver great experiences across devices and platforms.
Augmented Reality and Virtual Reality	
Amazon Sumerian (in Preview)	Amazon Sumerian lets customers create and run 3D, Augmented Reality (AR) and Virtual Reality (VR) applications. Customers can build immersive and interactive scenes that run on AR and VR, mobile devices, and web browser. Whether a customer is non-technical, a web or mobile developer, or has years of 3D development experience, getting started with Sumerian is easy. Customers can design scenes directly from their browser and, because Sumerian is a web-based application, customers can quickly add connections in their scenes to existing AWS Cloud services.
Application Integration	
Amazon MQ	Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Message brokers allow different software systems—often using different programming languages, and on different platforms—to communicate and exchange information. Amazon MQ manages the administration and maintenance of ActiveMQ, a popular open-source message broker. The underlying infrastructure is automatically provisioned for high availability and message durability to support customers’ reliability of applications. With Amazon MQ, customers get direct access to the ActiveMQ console and industry standard APIs and protocols for messaging, including JMS, NMS, AMQP, STOMP, MQTT, and WebSockets.
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 SNS	Amazon Simple Notification Service (Amazon SNS) is a fast, flexible, fully managed push notification service that lets customers send individual messages or to fan-out messages to large numbers of recipients. Amazon SNS makes it simple and cost effective to send push notifications to mobile device users and email recipients or even send messages to other distributed services. With Amazon SNS, customers can send notifications to Apple, Google, Fire OS, and Windows devices, as well as to Android devices in China with Baidu Cloud Push. Customers can use SNS to send SMS messages to mobile device users worldwide.
AWS AppSync (in Preview)	AWS AppSync automatically updates the data in web and mobile applications in real time, and updates data for offline users as soon as they reconnect. AppSync makes it easy to build collaborative mobile and web applications that deliver responsive, collaborative user experiences. Customers can use AWS AppSync to build native mobile and web apps with iOS, Android, JavaScript and React Native. Get started by going to the AWS AppSync console, specify the data for an app with simple code statements, and

Cloud Services Information	Service Description
	AppSync will manage everything needed to store, process, and retrieve the data for the application.
AWS Step Functions	AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components lets customers scale and change applications quickly. AWS Step Functions is a reliable way to coordinate components and step through the functions of application. AWS Step Functions provides a graphical console to arrange and visualize the components, which makes it simple to build and run multi-step applications. AWS Step Functions automatically triggers and tracks each step, and retries when there are errors, so the application executes in order and as expected. The state of each step is logged, so when things do go wrong, customers can diagnose and debug problems quickly. Customers can change and add steps without even writing code, so they can easily evolve their application and innovate faster.
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 an application. Coordinating tasks across the application involves managing inter-task dependencies, scheduling, and concurrency in accordance with the logical flow of the application. Amazon SWF gives customers full control over implementing tasks and coordinating them without having to manage underlying complexities such as tracking their progress and maintaining their state.
Amazon FPS	Amazon Flexible Payments Service (Amazon FPS) facilitates the digital transfer of money between any two entities (humans or computers).
Customer Engagement	
Amazon Connect	Amazon Connect is a self-service, cloud-based contact center service that makes it easy for any business to deliver better customer service at lower cost. Amazon Connect is based on the same contact center technology used by Amazon customer service associates around the world to power millions of customer conversations. The self-service graphical interface in Amazon Connect makes it easy for non-technical users to design contact flows, manage agents, and track performance metrics—no specialized skills required. Customers pay by the minute for Amazon Connect usage plus any associated telephony services.
Amazon Pinpoint	Amazon Pinpoint makes it easy to run targeted campaigns to drive user engagement in mobile apps. Amazon Pinpoint helps customers understand user behavior, define which users to target, determine which messages to send, schedule the best time to deliver the messages, and then track the results of a campaign. Amazon Pinpoint is built to scale with an app, enabling customers to collect and process billions of events per day, and send millions of targeted push notifications to users.

Cloud Services Information	Service Description
Amazon SES	Amazon Simple Email Service (Amazon SES) is a cost-effective email service built on the reliable and scalable infrastructure that Amazon.com developed to serve its own customer base. With Amazon SES, customers can send and receive email with no required minimum commitments. Amazon SES offers dedicated IP addresses, which enable customers to manage the reputation of the IP addresses that Amazon SES uses to send their email. Dedicated IP addresses are Amazon SES IP addresses that are reserved exclusively for email sending.
Business Productivity	
Alexa for Business	Alexa for Business allows organizations of all sizes to introduce Alexa to their workplace. With Alexa for Business, customers can use the Alexa they know as an intelligent assistant to stay organized and focus on the work that matters. Alexa helps workers be more productive as they move throughout their day, at home and at their desks as enrolled users with personal devices, and in conference rooms, copy rooms or other shared spaces with shared devices. Alexa for Business includes the tools and controls that administrators need to deploy and manage shared Alexa devices, skills, and users at scale.
Amazon Chime	Amazon Chime is a high-quality communications service that transforms online meetings with an easy-to-use app which works seamlessly across all devices. With Amazon Chime, customers can schedule and attend online meetings and video conferences, and chat, call, and collaborate, inside and outside of their organization, all with a single app. Now customers can work productively from wherever they are.
Amazon WorkDocs	Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities to help 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 WorkDocs default storage quota is 1TB.
Amazon WorkMail	Amazon WorkMail is a secure, managed business email and calendar service with support for existing desktop and mobile email clients. Amazon WorkMail gives users the ability to seamlessly access their email, contacts, and calendars using Microsoft Outlook, their web browser, or their native iOS and Android email applications. Amazon WorkMail offers SMTP journaling, which allows customers to record all email communication sent or received by their Amazon WorkMail organization. Journaling allows customers integrate with third-party compliance solutions for email archiving and e-discovery, allowing them to meet data security and information protection policies.
Desktop and Application Streaming	
Amazon WorkSpaces	Amazon WorkSpaces allows customers to easily provision cloud-based desktops that allow users to access the documents, applications, and resources they need with the device of their choice, including laptops, iPad, Kindle Fire, or Android

Cloud Services Information	Service Description
	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 highly competitive cost that is half the cost of most Virtual Desktop Infrastructure (VDI) solutions. Amazon WorkSpaces offers bundles that come with a Windows 10 desktop experience, powered by Windows Server 2016. Amazon WorkSpaces also provides Graphics bundles that offer a virtual cloud desktop with a high-end GPU that supports engineering, design, and architectural applications.
Amazon AppStream	The Amazon AppStream web service deploys application on AWS infrastructure and streams input and output between the application and devices such as personal computers, tablets, and mobile phones. The 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 App Stream 2.0	Amazon AppStream 2.0 is a fully managed, secure application streaming service that allows customers to stream desktop applications from AWS to any device, without rewriting them. AppStream 2.0 provides users with instant-on access to the applications they need, and a responsive, fluid user experience running in an HTML5 web browser.
Internet of Things	
AWS IoT Core	AWS IoT Core is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT Core can support billions of devices and trillions of messages, and can process and route those messages to AWS endpoints and to other devices reliably and securely. With AWS IoT Core, applications can keep track of and communicate with all devices, all the time, even when they aren't connected. AWS IoT Core makes it easy to use AWS Cloud services to build IoT applications that gather, process, analyze, and act on data generated by connected devices, without having to manage any infrastructure.
Amazon FreeRTOS	Amazon FreeRTOS (FreeRTOS) is an operating system that makes microcontroller-based edge devices easy to program, deploy, secure, and maintain. Amazon FreeRTOS is based on the FreeRTOS kernel, the popular open source operating system for microcontrollers, and includes software libraries that make it easy to securely connect devices locally, to the cloud, and update them remotely. The Amazon FreeRTOS console enables customers to easily select and download the software components relevant to their use case.
AWS Greengrass	AWS Greengrass is software that lets customers run local compute, messaging, and data caching for connected devices in a secure way. With AWS Greengrass, connected devices can run AWS Lambda functions, keep device data in sync, and communicate with other devices securely, even when not connected to the Internet. Using AWS Lambda, AWS Greengrass ensures that IoT devices can

Cloud Services Information	Service Description
	respond quickly to local events, operate with intermittent connections, and minimize the cost of transmitting IoT data to the cloud.
AWS IoT 1-Click	AWS IoT 1-Click is a new service that makes it easy for simple devices to trigger actions like Lambda functions. With AWS IoT 1-Click, simple devices are ready to securely connect to AWS IoT Core right out of the box. This makes it easy for developers to deploy these devices into their IoT applications and for device manufacturers to register their simple devices with AWS IoT Core at the factory.
AWS IoT Analytics	AWS IoT Analytics is a fully-managed IoT analytics service that collects, pre-processes, enriches, stores, and analyzes IoT device data at scale. IoT Analytics can perform simple ad hoc queries as well as complex analysis, and is a simpler way to run IoT analytics for use cases such as understanding the performance of devices, predicting device failures, and machine learning. It is designed specifically for IoT and automatically captures and stores the message timestamp so it is easy to perform time-series analysis. IoT Analytics can also enrich the data with device-specific metadata such as device type and location using the AWS IoT registry. IoT Analytics stores data in an IoT-optimized data store so customers can run queries on large datasets.
AWS IoT Button	The AWS IoT Button is a programmable button based on the Amazon Dash Button hardware. This simple Wi-Fi device is easy to configure and designed for developers to get started with AWS IoT Core, AWS Lambda, Amazon DynamoDB, Amazon SNS, and many other Amazon Web Services without writing device-specific code.
AWS IoT Device Defender	AWS IoT Device Defender is a fully managed service that helps customers secure their fleet of IoT devices. AWS IoT Device Defender continuously audits the security policies associated with devices to make sure that they aren't deviating from security best practices. AWS IoT Device Defender makes it easy to maintain and enforce security policies, such as ensuring device identity, authenticating and authorizing devices, and encrypting device data.
AWS IoT Device Management	AWS IoT Device Management is a service that makes it easy to securely onboard, organize, monitor, and remotely manage IoT devices at scale throughout their lifecycle. Customers can use IoT Device Management to upload and view device information and configuration, organize their device inventory, monitor their fleet of devices, and remotely manage devices deployed across many locations including updating device software Over-The-Air (OTA). With IoT Device Management, customers can scale their device fleets and reduce the cost and effort of managing large IoT device deployments.
Game Development	
Amazon Lumberyard (Beta)	Amazon Lumberyard is a free, cross-platform, 3D game engine for customers to create the highest-quality games, connect their games to the vast compute and storage of the AWS Cloud, and engage fans on Twitch. By starting game

Cloud Services Information	Service Description
	projects with Lumberyard, customers can spend more of their time creating great gameplay and building communities of fans and less time on the undifferentiated heavy lifting of building a game engine and managing server infrastructure.
Amazon GameLift	Amazon GameLift, a managed service for deploying, operating, and scaling session-based multiplayer games, reduces the time required to build a multiplayer backend from thousands of hours to just minutes. Available for developers using Amazon Lumberyard, Amazon GameLift is built on AWS's highly available cloud infrastructure and allows customers to quickly scale high-performance game servers up and down to meet player demand—without any additional engineering effort or upfront costs.
Software	
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.
Splunk	For environments that required additional logging and analytics for a SIEM and alerting environment, Splunk software can be implemented and all logs can be pushed into Splunk. Customized dashboards can be developed for security teams along with alerts.
AWS Support	
AWS Support	AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced technical support engineers to help customers of all sizes and technical abilities successfully use the products and features provided by AWS.

State of Utah and NASPO Value Point

Cloud Solutions

Solicitation SK18008

July 6, 2018

Organization and Staffing

Prepared for:

State of Utah
State Contract Analyst
Division of Purchasing

Prepared by:

Global Technology Resources, Inc.
990 S. Broadway, Suite 300
Denver, CO 80209
1-877-603-1984
Fax 1-888-803-6520
www.GTRI.com



TABLE OF CONTENTS

1.	Organization Profile.....	1
1.1.	Contract Manager.....	1
1.2.	Roles and Responsibilities of Contract Manager	4

1. Organization Profile

GTRI is a strategic IT partner that helps organizations navigate between business needs and technology solutions in order to advance and support the mission, reduce costs, and reduce risks. We leverage our team of highly certified engineers and architects that use proven methodologies to design and implement innovative solutions. Founded in 1998, GTRI is a Denver-based small business that maintains an average of 135 full-time employees. The company is a financially sound firm with net revenues that exceeded \$117M in 2017. We have a history of being ranked on the top of the Denver Business Journal's Fastest Growing Private Companies as well as Inc.'s 500|5000 list of the fastest growing private companies in America, and routinely accepts excellence awards from companies such as Boeing and Lockheed Martin. GTRI is ISO 9001:2008 certified for our processes and procedures in the procurement and delivery of IT systems to the Government.

1.1. Contract Manager

GTRI has identified a Contract Manager as the single point of contact for the management of the NASPO ValuePoint Master Agreement.

Gail Springer, Director Federal Programs
720-836-7331
GSpringer@gtri.com
Work hours M-F 7:00 am – 4:00 pm MT

We selected Ms. Springer as the Contract Manager due to her experience managing the Department of Interior's Foundation Cloud Hosting Services, ten-year Indefinite Delivery, Indefinite Quantity (IDIQ) contract. Ms. Springer began managing large government contracts for GTRI in May 2009. Since then she has managed over ten multiple year contracts. These contracts described in the table below:

Contract Name	Scope	Period of Performance
US Navy SPAWAR 8(a) PBX COTS contract	Provide PBX and associated equipment as competed per Request for Quote (RFQ) for over five (5) years.	May 2009-January 2013
US Navy SPAWAR Communications and Networking COTS contract	Provide Communication and Network COTS equipment and associated services as competed per Request for Quote (RFQ) for over five (5) years.	March 2012-March 2017
US Navy SPAWAR Intelligence Surveillance and Recognizance (ISR) COTS contract	Provide ISR COTS equipment and associated services as competed per Request for Quote (RFQ) for over five (5) years.	March 2012-March 2017
US Navy SPAWAR Multi Function Equipment COTS contract	Provide Multifunctional COTS equipment and associated services as competed per Request for Quote (RFQ) for over five (5) years.	June 2012 – June 2017
US Navy SPAWAR Command and Control COTS contract	Provide Command and Control COTS equipment and associated services as	April 2017-April 2022

	competed per Request for Quote (RFQ) for over five (5) years.	
US Air Force Network Centric -2 Products Contract	Provide COTS items the implement and integrate into the Air Force's central network for six (6) years.	November 2013- November 2019
US Army Information Technology Enterprise Solutions- 3 Hardware contract	Support the Army's requirements, within CONUS and OCONUS to include remote OCONUS, covering a full range of Information Technology (IT) equipment for client, server, storage, and network environments; for related incidental services and software; for maintenance/warranty of legacy IT equipment; and for warranty variations.	February 2016- February 2021
General Services Administration, Schedule 70 contract	Provide Professional Services through pre-negotiated labor categories and rates for the Federal government.	February 2007- April 2018
National Aeronautics and Space Administration, Solutions for Enterprise Wide Procurement IV contract	Solutions for Enterprise-Wide Procurement (SEWP, pronounced 'soup'), is a multi-award Government-Wide Acquisition Contract (GWAC) vehicle focused on commercial IT products and product-based services. With over 140 pre-competed Prime Contract Holders, SEWP offers a wide range of commercial advanced technology including tablets, desktops and servers; IT peripherals; network equipment; storage systems; security tools; software products; cloud-based services; telecommunications; Health IT; sensors; video conferencing systems and other IT, Communication and Audio-Visual products. Product based Services such as installation, training, maintenance and warranty and a full range of product-based services are also available through SEWP.	June 2007 – April 2015
National Aeronautics and Space Administration, Solutions for Enterprise Wide Procurement V Groups A and Group D contracts		May 2015- October 2025
Department of Interior, Foundation Cloud Hosting Services contract	Contractor shall utilize pricing for subsequent task orders as established within this contract as reflected in the base period and each of the option periods in accordance with Fixed Price (FP) Unit of Service (UoS). Pricing for this task order includes each of the Service Line(s) identified below. Storage Services	May 2013- April 2023

	• Secure File Transfer Services • Virtual Machine Services • Database Hosting Services • Web Hosting Services • Development and Test Environment Services	
--	---	--

Below is a detailed resume for Ms. Springer.

Gail Springer

Director, Federal Programs

Leads teams responsible for responding for contract management, Requests for Proposals responses, day to day operations of large government contracts, compliance management, and Facility Security Officer.

Summary:

- **Manage GTRI's prime government contract:**
 - o Implemented process to meet asset tagging requirements for contract and rolled out throughout the team. Manage team members for locating new vendors and improving process to register customer assets labels in appropriate databases.
 - o Built relationship with local Defense Contract Management Agency (DCMA) for first contract managed and audited by DCMA
 - o Coordinated with finance team to improve procedures and invoicing to decrease payment cycle by 20% and decrease invoicing errors by 60%.
 - o Lead implementation planning teams for preparation of 3 contracts anticipating award in 2012.
 - Set up several team to include all business units involved in business development task order response, execution of task orders, purchasing, payments invoicing, and reporting for anticipated. Advised each team of requirements to complete gap analysis, implementation plan, growth plan, and budget.
 - o **5-year \$5B Army services contract:** conducted internal award kickoff meeting and training with sales team. Developed and implemented OCONUS deployment procedures as required by contract.
 - o **7-year FBI services contract:** provided prime response sections and completed conditional offers to win first task order released on contract. GTRI's portion of task order includes \$200,000 for services in 8 years for 1 FTE. Indirect manager for FTEs on contract.
 - o **9-month task order for VOIP implementation for BRAC relocation:** recruited team lead and engineers on site to complete engineering installation plan and installation. Coordinated schedules and resources with prime project manager to meet customers' requirements.
 - o **2-Year Voice Proof of Concept implementation task order** with Veteran's Affairs under GSA Schedule 70 contract.
 - o **10-year Foundation Cloud Hosting Services contract:** includes 8-year task order for design, architect, implement, migrate applications and manage cloud hosted solution.

Certifications and Training:

PMP Certified #1497109
MBA

- Lead Cyber Security Team to implement new cyber security policies to meet NIST 800-175 requirements to protect customer data.

1.2. Roles and Responsibilities of Contract Manager

GTRI has a dedicated Program Management Office (PMO), assigned Program Manager (PM), automated tools, and processes to meet the requirements of the NASPO ValuePoint acquisition contract. Our PMO includes several business unit performers within GTRI necessary to manage an similar type multi agency use contract. Exhibit 1, below, shows the PMO team responsible for managing the contract.




Exhibit 1 - GTRI Program Management Office Structure.

Ms. Springer will oversee the contract as the dedicated Contract Manager (CM). She will utilize our Global Business Processing Suite (GBPS) of applications and tools based upon Microsoft Suite of Tools. Our GBPS has customizable dashboards and views to track the status of all quotes, deliveries, installations, reports, invoices, and payments. GBPS allows for automation of contract catalogs, pricing discounts and flexible reporting structures. These dashboards and reporting tools, clear reporting structures, clear chain of command, and sound processes allow our CM to manage and oversee all activity for immediate communication, including:

- Execute Master Agreement with Lead State,
- Execute Participating Addendums with Participating Entities,
- Maintain administrative, physical, technical, and procedural infrastructure associated with the provision of Cloud,
- Maintain records pertaining to the Master Agreement and orders Placed by Purchasing Entities,
- Report and pay administrative fees,
- Complete ValuePoint Summary and Detailed Usage Reports, and
- Participate cooperatively with NASPO ValuePoint personnel for marketing, training and performance reviews.

State of Utah and NASPO Value Point

Cloud Solutions

Solicitation SK18008

July 6, 2018

Technical Response

Prepared for:
State of Utah
State Contract Analyst
Division of Purchasing

Prepared by:
Global Technology Resources, Inc.
990 S. Broadway, Suite 300
Denver, CO 80209
1-877-603-1984
Fax 1-888-803-6520
www.GTRI.com



TABLE OF CONTENTS

Technical Response (M) (E) (8.1).....	1
1. AWS IaaS/PaaS Technical Requirements (M) (E) (8.1.1).....	1
1.1. How GTRI's Solution Meets the Essential Characteristics of Cloud Computing.....	1
1.1.1. How GTRI's Solution Complies with the Requirements in Attachment C	3
1.1.2. How GTRI's Solution Adheres to the Services, Definitions, and Deployment Models Identified in the Scope of Services, in Attachment D.....	4
1.1.3. AWS Support for Different Data Impact Levels	4
1.1.4. Support for the Essential Characteristics of Cloud Computing	4
1.1.5. AWS Infrastructure as a Service (IaaS).....	4
1.1.6. AWS Platform as a Service (PaaS).....	5
1.1.7. AWS Support for Cloud Deployment Methods.....	5
1.2. Subcontractors (E)(8.2)	5
1.3. Working with Purchasing Entities (E)(8.3).....	5
1.3.1. Describe How GTRI will Work with Purchasing Entities Before, During, and After a Data Breach (8.3.1).....	5
1.3.2. GTRI's Approach to Not Engage and Not Permit Agents to Push Adware, Software, or Marketing Not Explicitly Authorized (8.3.2)	12
1.3.3. Describe How Application-Hosting Environment Supports Test/Staging Environment Identical to Production (8.3.3)	13
1.3.4. Describe How Applications and Web Sites are Accessible to People With Disabilities (8.3.4).....	13
1.3.5. Describe How Content Delivered Through Web Browsers Are Accessible Using Current Versions of Multiple Browser Platforms (8.3.5)	13
1.3.6. How GTRI will Meet with Purchasing Entity and Cooperate to Determine if Any Sensitive or Personal Information Subject to Any Law, Rule or Regulation (8.3.6).....	14
1.3.7. Project Plans or Work Plans Used to Implement Solutions (8.3.7)	14
1.3.8. Updating Services Periodically (8.3.8)	16
1.4. Customer Service (E) (8.4).....	17
1.4.1. Ensure Excellent Customer Services is Provided (8.4.1).....	17
1.4.2. Ability to Comply with Customer Service Requirements (8.4.2).....	19
1.5. Security of Information (E) (8.5)	22
1.5.1. Measures GTRI Takes to Protect Data (8.5.1).....	23
1.5.2. How GTRI Complies with all Applicable Laws Related to Data Privacy and Security (8.5.2).....	24
1.5.3. GTRI's Approach to Not Access Purchasing Entity's User Accounts or Data (8.5.3).....	25
1.6. Privacy and Security (E) (8.6).....	26

1.6.1.	Commitment to Comply with NIST 800-145 (8.6.1).....	26
1.6.2.	All Government Standards or Standard Organization Security Certifications (8.6.2).....	27
1.6.3.	Security Practices in Place to Secure Data and Applications (8.6.3).....	28
1.6.4.	Data Confidentiality Standards and Practices in Place to Ensure Data Confidentiality (8.6.4).....	30
1.6.5.	Third Party Attestations, Reports, Security Credentials, and Certifications Relating to Data Security, Integrity, and Other Controls (8.6.5).....	30
1.6.6.	Logging Process Including Types of Services and Devices Logged (8.6.6)	30
1.6.7.	Restricting Visibility of Cloud Hosted Data and Documents to Specific User Groups (8.6.7).....	33
1.6.8.	Notification Process in the Event of a Security Incident (8.6.8).....	35
1.6.9.	Security Controls, Both Physical and Virtual, Used to Isolate Hosted Servers (8.6.9)	35
1.6.10.	Technical Reference Architectures that Support IaaS, SaaS, and PaaS (8.6.10)	36
1.6.11.	Security Procedures in Place Regarding Employees Who Have Access to Sensitive Data (8.6.11)	37
1.6.12.	Security Measures and Standards in Place to Secure Confidentiality of Data and Rest and in Transit (8.6.12).....	37
1.6.13.	Policies and Procedures Regarding Notifications of Data Breach (8.6.13).....	38
1.7.	Migration and Redeployment Plan (E) (8.7)	38
1.7.1.	Managing End of Life Activities Closing Down a Service and Safely Deprovisioning (8.7.1).....	38
1.7.2.	Orderly Return of Data Back to Purchasing Agency (8.7.2)	39
1.8.	Service or Data Recovery (E) (8.8)	40
1.8.1.	Contingency Plan to Respond to Certain Situations (8.8.1)	42
1.8.2.	Methodologies for Backup and Restore Services (8.8.2).....	44
1.9.	Data Protection (E) (8.9)	47
1.9.1.	Standard Encryption Technologies and Options to Protect Sensitive Data (8.9.1).....	47
1.9.2.	Willingness to Sign Relevant and Applicable Business Associate Agreements (8.9.2)...	47
1.9.3.	Only Use Data for Purpose Defined in Master Agreement, Addendum or Related Service Level Agreement (8.9.3).....	47
1.10.	Service Level Agreements (E) (8.10).....	47
1.10.1.	Whether Sample Service Level Agreement is Negotiable (8.10.1)	47
1.10.2.	Sample Service Level Agreement (8.10.2)	48
1.11.	Data Disposal (E) (8.11).....	49
1.12.	Performance Measures and Reporting (E) (8.12).....	50
1.12.1.	Ability to Guarantee Reliability and Uptime Greater than 99.5% (8.12.1)	50
1.12.2.	Standard Uptime Service and Related SLA Criteria (8.12.2)	51

1.12.3. Process Used for Customer to Call/Contact for Support (8.12.3).....	51
1.12.4. Consequences/SLA Remedies if Fail to Meet Incident Response and Incident Fix time (8.12.4).....	51
1.12.5. Procedures and Schedules for Planned Downtime (8.12.5).....	51
1.12.6. Consequences/SLA Remedies if Disaster Recovery Metrics are Not Met (8.12.6)	52
1.12.7. Sample Performance Reports (8.12.07)	52
1.12.8. Ability to Print Historical, Statistical, and Usage Reports Locally (8.12.8).....	55
1.12.9. On-Demand Deployment Services Supported 24x365 (8.12.9)	55
1.12.10. Scale-up and Scale-down Availability 24x365 (8.12.10)	55
1.13. Cloud Security Alliance (E) (8.13).....	57
1.14. Service Provisioning (E)	57
1.14.1. Processing Emergency or Rush Service Implementations (8.14.1).....	57
1.14.2. Standard Lead-time for Provisioning Solutions (8.14.2).....	58
1.15. Back Up and Disaster Plan (E).....	59
1.15.1. Ability to Apply Legal Retention Periods and Disposition By Agency (8.15.1).....	59
1.15.2. Known Inherent Disaster Recovery Risks and Potential Mitigation Strategies (8.15.2) ..	60
1.15.3. Infrastructure Supports Multiple Data Centers in US for Failover Capability (8.15.3)....	60
1.16. Hosting and Provisioning (E) (8.16)	61
1.16.1. Documented Cloud Hosting Provisioning Process and Defined/Standard Cloud Provisioning Stack (8.16.1).....	61
1.16.2. Tools Sets Used (8.16.2).....	62
1.17. Trial and Testing Periods (Pre and Post Purchase) (E) (8.17).....	65
1.17.1. Testing and Training Periods (8.17.1)	65
1.17.2. Providing Test and/or Proof of Concept Environment for Evaluation (8.17.2).....	66
1.17.3. Training and support Provided at No Additional Cost (8.17.3).....	67
1.18. Integration and Customization (E) (8.18).....	67
1.18.1. Solutions Can Be Integrated into Complementary Applications (8.18.1)	67
1.18.2. Ways to Customize and Personalize Solutions for Purchasing Entities (8.18.2).....	67
2. Marketing Plan (E) (8.19).....	68
3. Related Value-Added Services to Cloud Solutions (E) (8.20).....	68
3.1. Architecture Workshops.....	68
3.2. Landing Zone Implementation	69
3.3. Professional Services.....	69
3.3.1. Hybrid IT	69
3.3.2. Collaboration	71

3.3.3. Security	71
3.3.4. Big Data Analytics.....	71
3.4. Managed Services.....	72
3.5. Application Assessments.....	73
3.6. GIS Application Services	73
4. Supporting Infrastructure (E) (8.22)	73

List of Tables

Table 1 –Incident Response Roles	7
Table 2 – Browsers Supported.....	13
Table 3 – GTRI’s Approach to Updating Services	17
Table 4 – Sample Customer Service SLAs	19
Table 5 – AWS Support Features	22
Table 6 –Return of Backup Data.....	40
Table 7 –Foundational Services Sample SLAs.....	49
Table 8 –High Availability SLA’s	50
Table 9 –Standard Setup SLA’s.....	59
Table 10 –DevOps Automation Tools	62
Table 11 –Proof of Concept SLA’s.....	66

List of Figures

Figure 1 – GTRI’s Global Services Framework	15
Figure 2 – AWS Isolation and Deployment Options	24
Figure 3 – AWS Isolation and Deployment Options	26
Figure 4 – AWS Isolation and Deployment Options	27
Figure 5 – AWS Isolation and Deployment Options	36
Figure 6 - CloudCheckr Summary Performance Report.....	53
Figure 7 - CloudCheckr Detailed Utilization Report	54
Figure 8 - CloudCheckr Heat Map Report.....	54
Figure 9 - Export of historical performance data	55
Figure 10 – Auto Scaling with Elastic Load Balancing and Amazon CloudWatch alarms.....	56
Figure 11 – Global Map of AWS Regions and Availability Zones	61
Figure 12 – Regions and Availability Zones	61
Figure 13 – GTRI’s Service Offerings.....	69
Figure 14 – Roles in Security and Operations	72
Figure 15 – GTRI’s Responsibilities for Routine Tasks.....	72

Technical Response (M) (E) (8.1)

GTRI is pleased to provide NASPO an Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Value-added services to support each.

The IaaS, and PaaS offering is based upon Amazon Web Services (AWS) cloud hosting and GTRI's managed services. AWS offering can be a basic IaaS without any managed services, or it can be a PaaS, or it can add additional value-added services. Therefore, we are providing the technical point by point details for the most technical requirements with PaaS and pricing out both PaaS and IaaS with AWS as the Cloud Hosting Provider (CSP).

Lastly, GTRI is providing Value-added services to support both options and assessment to prepare for these options.

1. AWS IaaS/PaaS Technical Requirements (M) (E) (8.1.1)

GTRI is an AWS Advanced Partner that allows us to resell, implement and maintain AWS services. The sections below identify how AWS cloud solution meets the technical requirements and GTRI's Managed Services deliver a secure, reliable, cloud PaaS solution.

1.1. How GTRI's Solution Meets the Essential Characteristics of Cloud Computing

According to the official NIST definition, "cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

The NIST definition lists five essential characteristics of cloud computing: on-demand self-service, broad network access, resource pooling, rapid elasticity or expansion, and measured service. The definition is intended to serve as a means for broad comparisons of cloud services and deployment strategies, and to provide a baseline for discussion from what is cloud computing to how to best use cloud computing. The following sections detail how GTRI meets the NIST definition of cloud.

GTRI's solution meets all five on NIST essential characteristics of cloud computing:

- o *On Demand Self-Service*
- o *Broad Network Access*
- o *Resource pooling,*
- o *Rapid elasticity or expansion, and*
- o *Measured service.*

On-Demand Self-Service

GTRI's Cloud Service Providers (CSP) 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 with the equivalent of 1, 10, 100, or 1,000 servers. Whether an organization needs to prototype an application or host a production solution, AWS makes it simple for customers to get started and be productive.

Broad Network Access

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.

Resource Pooling

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 Security Standards 3.2 as of July 2016.

Rapid Elasticity

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.

Measured Service

AWS uses 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.

Hybrid Model (Extend IT Services)

A hybrid cloud environment allows organizations to address immediate IT needs though utilizing the benefits of cloud computing, while also retaining on-premises infrastructure. A hybrid model is a prudent approach to cloud adoption for organizations that require the immediate use of scalable cloud services but are not ready to fully migrate all application and workloads to the cloud.

AWS provides the tools and solutions to integrate existing on-premises resources with the AWS cloud. By using AWS to enhance and extend your capabilities, without giving up the investments you have already made, you can accelerate your adoption of cloud computing.

General Hybrid Cloud Requirements and Issues: Some of the common requirements and issues associated with hybrid cloud are:

- On-demand, scalable compute resources.
- Flexible, secure, and reliable network connectivity.
- Automated backup and recovery.
- A highly secure and controlled platform, with a wide array of additional security features.
- Integrated access control.
- Easy-to-use management tools that integrate with on-premises management resources.

AWS Capabilities for Hybrid Cloud Solutions: AWS provides all the capabilities required for a dynamic, reliable, and secure hybrid cloud solution:

- **Extend Network Configuration:** Flexible network connectivity is a cornerstone of integrating distributed environments, including AWS and your existing on-premises equipment. With Amazon VPC, you can extend your on-premises network configuration into your virtual private networks on the AWS cloud. AWS resources can operate as if they are part of your existing corporate network. Amazon VPC lets you provision a logically isolated section of the 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.
- **Integrated Cloud Backups:** AWS helps simplify the backup and recovery environment for the enterprise. You can leverage the on-demand nature of the cloud and automate your backup and recovery processes, so they are not only less complex and lightweight, but also easy to manage and maintain. Storage services with AWS are designed to provide 99.999999999% durability, so you can feel confident your backups are protected.
- **Integrated Network Connection:** On-premises connection with AWS is best accomplished with AWS Storage Gateway, a software appliance installed in your data center with cloud-based storage to provide seamless and secure integration between an organization's existing IT environment and the AWS storage infrastructure. Using industry-standard storage protocols, the service allows you to store data in the AWS cloud for scalable and cost-effective storage. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all your data encrypted in the Amazon Simple Storage Service (Amazon S3) or Amazon Glacier.
- **Integrated Resource Management and Workload Migration:** All AWS cloud services are driven by robust APIs that allow for a wide variety of monitoring and management tools that integrate easily with your AWS cloud resources. It's likely that many of the tools that your organization is using to manage your on-premises environments can be extended to include AWS as well. Integrating your AWS environment can provide a simpler and quicker path for cloud adoption, because your operations team does not need to learn new tools or develop completely new processes.

Storage services with AWS are designed to provide 99.999999999% durability, so you can feel confident your backups are protected.

Solution Use Cases: Use cases for AWS hybrid solutions include:

- Migrating workloads and data that are “cloud ready” (i.e., applications that do not need significant rearchitecting for a cloud migration).
- Retaining data on-premises to meet regulatory and compliance needs.

Hybrid Cloud Resources: AWS provides the tools, information, and guidance to build a hybrid cloud environment that can offer an immediate impact to customers.

1.1.1.1. How GTRI's Solution Complies with the Requirements in Attachment C

The leading cloud service providers constantly innovating to deliver new services to market. Also, as prices fall, and cloud providers achieve economies of scale, each unit of compute, storage and networking also drops over time and cloud service providers, being in a highly competitive environment, typically pass these discounts on to customers.

GTRI believes that the simplest, most transparent and most efficient way to pass along savings to customers is by providing a fixed minimum discount of AWS' published pricing. That way, customers know they are getting a superior price, and they know that their prices will drop as cloud service providers drop their prices.

The fixed discount also allows customers to take advantage of cost estimation and optimization tools that are on the market to optimize their spending in the cloud.

1.1.2. How GTRI's Solution Adheres to the Services, Definitions, and Deployment Models Identified in the Scope of Services, in Attachment D

GTRI can support applications with low risk data, moderate risk data, or high-risk data, as defined in Attachment D. The subsections below explain how we meet these requirements.

1.1.3. AWS Support for Different Data Impact Levels

Numerous Federal Civilian and Department of Defense (DoD) organizations have successfully achieved security authorizations for low, moderate and high impact systems hosted on AWS. The AWS infrastructure has been evaluated by independent assessors for a variety of government systems as part of their system owners' approval process.

AWS offers the FedRAMP compliant systems that have been granted authorizations, have addressed the FedRAMP security controls (based on NIST SP 800-53), and have been assessed by an accredited independent third-party assessor (3PAO) and maintains continuous monitoring requirements of FedRAMP.

AWS GovCloud (US), has been granted a Joint Authorization Board Provisional Authority-To-Operate (JAB P-ATO) and multiple Agency Authorizations (A-ATO) for high impact level, which would be suitable for *high risk data*. The services in scope of the AWS GovCloud (US) JAB P-ATO boundary at high baseline security is continuously updated on the FedRAMP services section of the AWS Web site.

AWS US East-West, has been granted a Joint Authorization Board Provisional Authority-To-Operate (JAB P-ATO) and multiple Agency Authorizations (A-ATO) for moderate impact level, which would be suitable for *low and moderate risk data*. The services in scope of the AWS US East-West JAB P-ATO boundary at Moderate baseline security categorization is continuously updated on the FedRAMP services section of the AWS Web site.

GTRI has experience designing, building and assisting in the Authority to Operate (ATO) process for multiple low, moderate and high impact systems.

1.1.4. Support for the Essential Characteristics of Cloud Computing

GTRI's cloud solution fully implements the five essential characteristics of cloud computing. Section 1.1.1 above describes this in more detail.

GTRI's cloud solution meets the criteria for Infrastructure as a Service, and Platform as a Service, depending on the components of the AWS service catalog that a customer wants to use.

1.1.5. AWS Infrastructure as a Service (IaaS)

AWS IaaS is the baseline for GTRI's cloud solutions. AWS provides many basic services that comprise a fully-featured IaaS. Here are AWS' basic IaaS building blocks:

Compute: Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud. EC2 provides a simple web service interface to obtain and configure capacity easily. It provides complete control of computing resources and can scale capacity up and down quickly as computing needs change.

Amazon EC2 also has security groups, which are a fully managed firewall service to provide fine-grained control over what network traffic can reach a host.

Storage: Amazon S3 is object storage, allowing users to store and retrieve any amount of data from anywhere. Amazon Elastic Block Store (EBS) provides persistent block storage volumes for use by EC2 instances.

Networking: Amazon Virtual Private Cloud (Amazon VPC) lets users provision a logically isolated section of the AWS Cloud where they can launch AWS resources in a virtual network that they define. Users have complete control over their virtual networking environment, including selection of their own IP address range, creation of subnets, and configuration of route tables and network gateways. Users can use both IPv4 and IPv6 in your VPC for secure and easy access to resources and applications

AWS has a number of higher-level platform services that build on the power of these basic infrastructure services.

1.1.6. AWS Platform as a Service (PaaS)

On AWS, it is also possible for the consumer not to manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but retain control over the deployed applications and possibly application hosting environment configurations. A few examples of AWS services that meet these criteria are:

Elastic Beanstalk: AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services. Consumers can simply upload application code and Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring. At the same time, users retain full control over the AWS resources powering your application and can access the underlying resources at any time.

AWS Lambda: AWS Lambda lets users run code without provisioning or managing servers. Users pay only for the compute time you consume - there is no charge when application code is not running. Unlike Elastic Beanstalk, there is no access to the underlying servers. Users can also set up code to automatically trigger from other AWS services or call it directly from any web or mobile app.

Elastic Container Service: Amazon Elastic Container Service (Amazon ECS) is a highly scalable, high-performance container orchestration service that supports Docker containers and allows users to easily run and scale containerized applications on AWS. Amazon ECS eliminates the need to install and operate container orchestration software, manage and scale a cluster of virtual machines, or schedule containers on those virtual machines.

1.1.7. AWS Support for Cloud Deployment Methods

As the control plane for the AWS cloud is accessed over the Internet, it is a public cloud provider. That said, it is possible to create infrastructure that is not accessible over the internet, that is private or accessible only to a community of interest.

AWS is also well-suited for a hybrid cloud deployment model. More information on how that is accomplished is available above in Section 1.1.1

1.2. Subcontractors (E) (8.2)

GTRI does not intend to use subcontractors for our AWS solutions.

1.3. Working with Purchasing Entities (E) (8.3)

1.3.1. Describe How GTRI will Work with Purchasing Entities Before, During, and After a Data Breach (8.3.1)

GTRI will follow a shared responsibility plan and rely on AWS to identify data breaches that are in their control. GTRI will utilize monitoring tools described later in this proposal to closely monitor and track the customer's cloud environment. If a breach should occur, GTRI will follow our Security Incident Procedures if monitoring or notification from any party identifies a breach. Table 1 below identifies the personnel and responsibilities for a breach.

Role	Responsibilities
Contract Manager (IPM)	- Assess all reports of suspicious cyber activity from users or System Administrator/Network Administrator and report IT security incidents to infrastructure - Advise the Help Desk/System Administrator or Network Administrator on any immediate mitigation actions to be taken - Immediately initiate an Incident/Event Log when a cyber event/incident is suspected, and ensure the documentation of all incidents/events supports technical analysis and legal evidentiary requirements - Develop, implement, and maintain the Incident Response Plan (IRP) and coordinating the IRP with customer - Ensure that the Incident Response Team (IRT) supports and participates in incident response test exercises coordinated by the customer - Exercise the IRP on an annual basis, at a minimum - Update the IRP to incorporate lessons learned from the annual exercise
IRT Members	- Verify and identify cyber incidents and events - Develop and approve triage and incident management mitigation strategies and actions - Assess the operational impacts of incidents - Provide recommendations to the IPM based on the impact on mission or readiness caused by the incident/event - Provide subject matter expert guidance and recommendations in the area of individual expertise on specific mitigation and response actions
Help Desk	- Receive reports of security events/incidents - Report all events/incidents to the IRT Program Manager - Provide the persons reporting events/incident with guidance on as pre-instructed by the IRT - Record all events/incidents reported in a log file
System Administrators	Coordinate and cooperate with the Incident Response Team on mitigation actions impacting applications, systems and networks with which the System Administrator interfaces.
Users	Report ALL suspicious computer events/incidents to the Help Desk or security officer

Role	Responsibilities
	- Provide input to an Incident/Event Report Log when suspicious activity is detected, or as directed by the security officer or Systems Administrator - Promptly perform mitigation actions directed by the Help Desk or IRT - Take only those mitigation actions directed by the Help Desk or IRT - Coordinate and cooperate with the Help Desk and IRT

Table 1 –Incident Response Roles

Pre-Incident Actions

The following actions will be taken as part of this plan to attain the best possible defensive posture in advance of future cyber probes or attacks:

- The provisions of the infrastructure will be incorporated into the security awareness training program to familiarize all personnel with the plan and their responsibilities in the event of a cyber incident occurrence or a request from the customer is received.
- The Incident Response Team will review current administrative, operational and support SOPs, policies and procedures, standards, Service Level Agreements (SLA), Memoranda of Agreement or Understanding (MOA/MOUs) to ensure consistency with the provision of the IRP and to identify any new arrangements needed; e.g. with local law enforcement, legal or intelligence organizations.
- The IPM will verify and test to ensure that all customer tasking affecting applications, systems or networks have been applied.
- The CISO will plan for the augmentation of the IRT and other key offices to cover incidents of extended duration.
- Legal counsel will be consulted before an incident to facilitate protecting the chain of custody of evidence in the event of an incident.

Incident Recognition

The most important first step is to recognize that an incident may be unfolding. Early recognition allows a rapid and informed identification of the nature and scope of the incident. The earlier mitigation strategies can be applied, the more likely they are to be successful. While incidents and events can present a variety of “symptoms”, some of the most common are:

- Website Defacement
- Denial of Service or unwanted disruption of service caused by activities such as an Email-blast, more commonly called “Spamming.”
- Execution of Malicious Code (Virus, Worm, Trojan Horse)
- Access Compromises, such as:
- Attempts (either failed or successful) to gain unauthorized access to a system, or its data, e.g. unsuccessful log-on attempts

- o Unauthorized use of a system for the processing or storage of data; e.g. unexplained output on a screen or a printer
- o Unauthorized use of another user's account
- o Unauthorized use of system privileges
- o Suspicious entries in a system/network account
- o Unexplained new user accounts
- o Unexplained changes in system files
- o Unexplained attempts to write to system files
- o Unexplained modifications to, or deletions of, data, file lengths, file dates, especially in system .EXE files
- o Unexplained file names
- o Unexplained new files
- o Unusual usage patterns or time of use profiles

Incident Reporting

1. Provide a verbal or written initial report to customer within one (1) hour after discovery/detection of a Priority Level 1 incident. The following are examples of Level 1 incidents:
 - Root Compromise
 - User Compromise
 - Denial of Service/Distributed Denial of Service attacks (no matter how successful or unsuccessful)
 - Website defacements
2. Provide an initial report to customer within one (1) hour after discovery/detection of Priority Level 2 incidents, including:
 - User Compromise
 - Successful virus/worm infection
 - Successful introduction of a virus/worm into a network
 - Scanning of classified or critical systems
3. For Priority Level 1 and Priority Level 2 incidents, provide a written preliminary incident report to Customer within 24 hours containing as much information as possible.
4. For Priority Level 1 and Priority Level 2 incidents, provide a written final report to Customer within ten working days of the resolution of an incident.
5. For Priority Level 1 and Priority Level 2 incidents, in cases where incident resolution is expected to take more than thirty (30) days, provide a status report to Customer every ten (10) days.
6. Report Priority Level 3 incidents on a weekly basis. The following are examples of Level 3 incidents:
 - Scanning of unclassified, non-critical systems
 - Detection and elimination of malicious logic before infestation

7. Report Priority Level 4 incidents on a monthly basis. The following are examples of Level 4 incidents:
 - Misuse of resources
 - Spam e-mail
 - Fraudulent e-mail
 - Social Engineering
8. Provide reports to Customer in response to vulnerability management and patch installation data calls by the due date and time.
9. Report incidents involving loss or compromise of classified information to the Department Security Officer, in addition to Customer.

Incident Response Procedures

1. Delegate one person to lead the response. The lead person should have as broad a view as possible of the environment in which the incident occurred and should be trained in incident response procedures.
2. Identify that an incident has occurred or is occurring. Verify that the situation is not the result of a simple mistake. Do not rush through the procedures. Instead, follow all the steps of the procedures. Take precautions to prevent destruction or corruption of evidence that may be needed to support criminal prosecution.
3. Notify the appropriate individual in your Component who is responsible for verbally reporting the incident to customer. Limit the notifications to those who have a legitimate need to know.
4. Determine the nature and scope of the incident. Ask questions (who, what, when, where) and take good notes. Timestamp all notes. If the notes are recorded in a stitched binder, it is easy to verify that no notes have been lost or misplaced.
 - a. Look for modifications to system software and configuration files
 - b. Look for tools installed by the intruder
 - c. Check other local systems for modification
 - d. Check remote systems for modifications
 - e. Notify customer to check for systems at other sites that may be involved
5. Maintain all evidence. Contact your agency legal counsel to determine how to properly handle evidence. Identify and properly secure all evidence to maintain its validity in court. Keep a log of everyone who has access to the evidence.
6. Maintain a low profile. Send a small on-site team to secure the area, ask questions, and review the information from the identification phase. Avoid tipping off the attacker. Maintain standard procedures and avoid looking for the attacker with obvious methods.
7. Isolate the system. In some cases, taking the system offline may be a serious step because of the importance of its services. The need to continue important services should be weighed against the potential harm that can arise from the security incident. Therefore, the decision to isolate the system must be made at the appropriate management level. The decision should come after careful examination of system logs to determine if the attack is external or internal and to evaluate the risks of continuing operation. Whether the system is isolated or not, appropriate actions shall be taken to contain the attack and prevent further attacks.

8. Backup the system. If possible, use two different backup methods. The information obtained from backing up the system may be used as evidence. Therefore, the backup media should be previously unused in order to avoid suggestions that it could be faulty. Log the original backup media properly as evidence. Include the following:
 - a. Registers, cache contents
 - b. Memory contents
 - c. State of network connections
 - d. State of running processes
 - e. Contents of the storage media
 - f. Contents of removable and backup media
9. Protect the chain of custody of the backup data. Store the data in a secure location. Keep a record of the individuals who have touched each piece of evidence. The record should include the date, time, and locations of where the evidence is stored.
10. Resolve the problem and return the system to normal operating status. The following are steps that may be taken to eradicate the problem:
 - a. Understand the cause of the incident
 - b. Determine the vulnerability that the intruder used to enter the system
 - c. Load a clean, uncontaminated operating system
 - d. Install appropriate software to ensure the incident will not reoccur
 - e. Apply all the latest patches
 - f. Remove any unnecessary services
 - g. Install a file integrity assessment tool
 - h. Assess neighboring computers to ensure the problem did not spread to other computers
 - i. Verify that the system has returned to its normal operating condition
 - j. Monitor the system to ensure no "back doors" were left undetected

If the problem cannot be solved by the internal incident response personnel, contact customer for technical assistance.

Vulnerability Management

The Component IRT Program Manager will implement the following procedures to ensure proper actions are taken for the success of the Customer's Vulnerability Management Program.

Upon receiving Customer Security Alerts, Product Security Bulletins, Virus Bulletins, and Requests for Validation, the IRT Program Manager will review and take actions where needed to ensure that the Component's Systems Administrators implement the recommended safeguards to protect systems against serious vulnerabilities.

Customer Alerts provide notification about critical new vulnerabilities that pose an immediate threat and contain instructions on how to mitigate these vulnerabilities. Customer Product Security Bulletins and Security Broadcasts provide notification of vulnerabilities or potential vulnerability issues of lesser impact than a Security Alert. Virus Bulletins provide warning about serious viruses that have either been found within the Customer's systems or are in the wild and could infect Customer computer systems if protections are not implemented.

The actions to be taken are:

Customer Security Alerts: The IPM will review and evaluate the information in the Alert and will consult with the Systems Administrators to determine if action needs to be taken to protect the Component's systems. The IPM will disseminate appropriate information and direct all System Administrators to apply the corrective actions as soon as possible. System Administrators will report compliance or justify why the corrective action could not be applied within thirty (30) days or less

Customer Virus Bulletins: The IPM will review and evaluate the information in the bulletin and will consult with the Systems Administrators to determine if action needs to be taken to protect the Component's systems. The IPM will disseminate information to all System Administrators for evaluation and implementation of preventative actions if needed. System Administrators will report any action taken or explain why the preventative action was not applicable

Customer Product Alerts: The IPM will review and evaluate the Product vulnerability information in the bulletin and, for those products used by the Component, will provide the vulnerability information to the Systems Administrators to determine if action needs to be taken to protect the Component's systems

Customer Requests for Validation: The IPM will immediately acknowledge receipt of the request to the customer. The IPM will consult with the System Administrators to determine compliance. The IPM will fill out the validation report with the required information in the format requested. The IPM will track compliance efforts for all systems that do not satisfy requirements and will report to the customer on the progress or problems, including dates and times at which the system(s) are expected to be fixed.

INFORMATION DISSEMINATION CONTROL

Since incident reports reveal sensitive information about the vulnerabilities, capacity to respond and operational readiness, rules for dissemination and handling controls are necessary.

All incident-related materials and documents are to be marked LIMITED OFFICIAL USE ONLY (LOUO), at a minimum. If an incident report is classified, all storage, transmission and communication of NSI shall be done in accordance with the SPOM. If possible, for a classified incident, sanitize and declassify the incident report, and report through normal channels.

The CIO is the Department representative responsible for approving the release of incident reports and associated documentation outside the customer community. The IPM is responsible for releasing incident response and associated information within the customer community.

INCIDENT RESPONSE PLAN COMPLIANCE REQUIREMENTS

Given the extraordinary importance that proper implementation of this incident response plan holds for mission performance and readiness, failure to execute the provisions of this plan through negligence or willful disregard may result in adverse administrative or disciplinary action.

AWS Response Plan

For the services that AWS controls, they have a response plan that notifies GTRI or an incident or critical vulnerability. After GTRI receives the notifications from AWS, we follow our incident response plan identified above.

AWS services are content agnostic, in that they offer the same high level of security to all customers, regardless of the type of content being stored, or the geographical region in which they store their content. Customers retain ownership and control of their content when using AWS services. Customers, rather than AWS, determine what content they store or process using AWS services. Because it is the customer who decides what content to place in the AWS cloud, only the customer can determine what level of security is appropriate for the content they store and process using AWS. Given that customers maintain control of their content when using AWS, customers retain the responsibility to monitor their own

environment for privacy breaches and to notify regulators and affected individuals as required under applicable law.

AWS has implemented a formal, documented incident response policy and program. The policy addresses purpose, scope, roles, responsibilities, and management commitment and has been developed in alignment with the ISO 27001 standards, system utilities are appropriately restricted and monitored. Below is an outline of the three-phased approach AWS has implemented to manage incidents:

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

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

The AWS incident management program is reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance. Additionally, the AWS incident response playbooks are maintained and updated to reflect emerging risks and lessons learned from past incidents. Plans are tested and updated through the due course of business (at least monthly).

1.3.2. GTRI's Approach to Not Engage and Not Permit Agents to Push Adware, Software, or Marketing Not Explicitly Authorized (8.3.2)

GTRI's business approach does not include the use of agents to service the contract. Therefore, we will not have any agents engage or push adware, software or market items not explicitly authorized by the contract.

Internally, GTRI has a sales team that focuses specifically on State and Local customers. This sales team has in depth knowledge of NASPO contracts from over a decade of being agents on other contracts. This experience brings and understanding of the scope requirements for contracts to meet this requirement.

1.3.3. Describe How Application-Hosting Environment Supports Test/Staging Environment Identical to Production (8.3.3)

With GTRI's CSPs, production systems may be easily cloned for use as development and test environments. Staging environments may be easily promoted to production. GTRI's CSPs also implement a snapshot capability to create consistent point-in-time snapshots of block storage devices and virtual machine images. In an AWS environment, GTRI will utilize AWS' Elastic Block Storage snapshot capability to create Amazon Machine Images (AMIs) that are identical to the production environment at a user-specified point in time.

Using that snapshot capability and orchestration software such as Ansible, GTRI can create fully-automated processes to take a snapshot of a development build and adjust its configuration to operate in a UAT or production environment. To diagnose production issues, it is possible to run the process in reverse – to “clone” a mis-behaving production server in a testing environment.

1.3.4. Describe How Applications and Web Sites are Accessible to People With Disabilities (8.3.4)

In 1998, the U.S. Congress amended the Rehabilitation Act of 1973 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.

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 offers the 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. In addition, AWS provides two graphical user interfaces, the AWS Management Console and the AWS ElasticWolf Client Console.

1.3.5. Describe How Content Delivered Through Web Browsers Are Accessible Using Current Versions of Multiple Browser Platforms (8.3.5)

The following browsers are supported for the GTRI's CSPs:

Browser	Version
Google Chrome	Latest three versions
Mozilla Firefox	Latest three versions
Microsoft Edge	Latest three versions
Apple Safari for MacOS	Latest two versions
Microsoft Internet Explorer	11

Table 2 – Browsers Supported

1.3.6. How GTRI will Meet with Purchasing Entity and Cooperate to Determine if Any Sensitive or Personal Information Subject to Any Law, Rule or Regulation (8.3.6)

Getting the security confirmed up front is critical for implementing a solution to properly protect and sensitive or personal information subject to any law, rule or regulation. As part of our Global Solutions Framework (GSF), GTRI initiates each new implementation or addition to an implementation with a kickoff meeting to confirm understanding of scoping. In these kickoff meetings, GTRI has a list of critical questions to confirm with any Purchasing Entity.

As part of the discovery process, GTRI will confirm the impact level of the application. GTRI will also confirm what the customer's compliance needs are because that often drives architectural decisions. As part of the discovery process around governance, risk, security and compliance, GTRI will also confirm any specific security needs, any other laws the system is subject to, and if any sensitive or Personal information may be stored in the solution.

1.3.7. Project Plans or Work Plans Used to Implement Solutions (8.3.7)

GTRI adheres to a very meticulous Global Services Framework (GSF) methodology that is a documented project delivery framework based on and customized around best practices such as ITIL and ISO. It is an environment that creates scalable best practices to meet stringent client needs and expectations, and provides a flexible, highly adaptable foundation for successfully envisioning, designing, deploying, and evolving the next phase of your solution.

This is a part of the strong value GTRI brings to our customer's environment in every project we engage in. It helps provide solutions faster, with fewer people, lower cost, less risk and greater business value. This methodology helps teams address the most common causes of project failure to improve success rates, solution quality, and business impact. The GSF fosters the ability to adapt to continuous change within dynamic projects.

GTRI's GSF separates projects into distinct and logical phases that reduce the likelihood of losing project control. The completion of one phase creates a solid base for the success of the next. Planning and integration issues are identified at the start of the project cycle, so difficulties can be discovered and resolved before client productivity is impacted, project costs multiply and project success is threatened because of unforeseen problems during implementation.

A clearly defined framework with roles, responsibilities, objectives, milestones and success measurements minimizes confusion and delay while maximizing efficiency and cost reduction.

Each customer engagement GTRI participates in goes through a life cycle, a process that includes all the activities in the project that take place up to completion and transition to an operational status. The main function of the GSF is to establish order around which project activities are performed and what the final deliverable is.

Below is a simple view of the GSF process model life cycle.




Figure 1 – GTRI's Global Services Framework

Envision: We work with the client to clearly define your business requirements and objectives to identify a solution best suited to your needs. We discuss key project variables and pinpoint success criteria for the project. During this phase, we also uncover how new solutions can benefit your organization.

Design: We assign a dedicated project manager and a technical consulting team to the project. The team performs onsite assessments / interviews, creates designs and project plans, and documenting all inclusions and exclusions for the project. All this information is compiled into GTRI's formal design document which defines the deployment for all team members.

Deploy: The information compiled in the design phase is used to guide and execute the project. We conduct step-by-step project execution including: physical deployment, testing, cut-over, and training. Each step is followed by an acknowledgement or sign-off by the client to ensure the outcome of each step is met and complete.

Evolve: We validate that the work performed meets the business requirements and the success criteria defined in the envision phase. All deliverables defined in the design phase are submitted for final review and acceptance. During this phase, GTRI can assess future needs, and so begins a second envision phase. GTRI's process comes "full circle" and sets in motion next steps for monitoring, managing, modifying roadmaps, and creating a lasting partnership.

GTRI can assist any client with the implementation of a best practice cloud governance operation that will enable clients to effectively use and benefit from Amazon Web Services ("AWS"). GTRI performs the following six principal tasks across project phases, customized for individual clients, which are generally outlined below and described in further detail herein:

- Phase 1: Migration Planning
 - Task 1 - Develop Recommendation and Implementation Plan
 - Task 2 - Develop Final Implementation Plan & WBS
- Phase 2: Design
 - Task 3 - Develop Security and Application Architecture

- o Task 4 - Conduct Service Catalog Portfolio Analysis and Develop Portfolio Product Roadmap
 - Phase 3: Implementation
 - o Task 5 - Implement operational systems
 - o Task 6 - Project Management

GTRI is a certified **Amazon Web Services Advanced Public-Sector Partner** and **Cisco Gold Partner**. GTRI's clear experience in the software sector and cloud markets are matched by our experience in the on premise, data center, and network worlds. The company is therefore uniquely positioned to rapidly engage and execute on behalf of you on several fronts simultaneously. We collaborate with our valued network of partners, resellers, licensors, third-party contractors and choose our alliances carefully to ensure we provide best-in-breed solutions to our customers. Our culture and relationships are critical assets that you will come to appreciate as much as we do as it sets forth the foundation to attract, develop, and retain great talent to better serve you, as you adopt the cloud into your IT environment.



1.3.8. Updating Services Periodically (8.3.8)

Cloud computing is a highly competitive market, and service providers must constantly add and improve services while reducing end-user prices to survive. GTRI is committed to ensuring that its customers benefit from the improvements and the price cuts that cloud service providers are constantly making.

Requirement	GTRI's Capability
How Offeror's services during Service Line Additions and Updates pursuant to section 2.12 will continue to meet the requirements outlined therein. How Offeror will maintain discounts at the levels set forth in the contract.	GTRI implements its discount to customers in the form of a discount from the cloud service provider's published pricing for all services. Therefore, as the service provider adds or updates services, and they reduce their published pricing (which AWS does several times per year), the customer will enjoy a discount and be assured of competitive pricing.
How Offeror will report to the Purchasing Entities, as needed, regarding changes in technology and make recommendations for service updates.	For changes in AWS technologies, customers can subscribe to receive information about the latest products, services, and feature announcements from AWS. GTRI also provides workshops that are offered free of charge to the public for more detailed information. For customers of our optional managed services, GTRI provides written reports on a regular basis analyzing the customer's service utilization and providing cost and performance optimization suggestions. Even outside of our managed services programs, we have found that helping our customers save money in the cloud is good for business in the long-term and it is in our best interests to do so.
How Offeror will provide transition support to any Purchasing Entity whose	GTRI provides professional services to assist customers with any service migrations that may need to happen. In such a case, we would perform an assessment and provide

operations may be negatively impacted by the service change.	an end-state architecture diagram, narrative and cost estimate to perform the work necessary. It is worth noting that in practice, cloud service providers, especially AWS, typically do not remove services from their catalog or increase their pricing. The scenario we see among our customers is that the cloud service provider will release a new generation of virtual machines that provide superior performance at a lower price point or release a new managed service that performs the same task at a lower price point. In these cases where there are new service offerings, or new virtual machine types that will provide a customer superior performance at a lower price point, we will inform customers when we provide them utilization reports and optimization suggestions.
--	--

Table 3 – GTRI's Approach to Updating Services

1.4. Customer Service (E) (8.4)

1.4.1. Ensure Excellent Customer Services is Provided (8.4.1)

During the on-boarding process, GTRI will work with the customer to establish support interfaces. For all support requests, the partner or the partner's client may request support using either phone, email, portal, or chat communications as follows:

- o GTRI Support Phone – (855) 290-5488
- o GTRI Support Email – support@GTRI.cloud
- o GTRI Support Portal – <https://support.GTRI.cloud>
- o GTRI Support Chat – Find in the support portal.

When you submit an email to our support address (above), this will automatically create a new ticket within our ticketing system. If you wish to utilize the customer portal for our ticketing system you can do so at the address noted above. When you create a new account here you can track and manage your requests online (add comments, resolve, etc.). When you initiate a chat session with our support agents, the chat transcript will be copied into a new ticket when the chat session ends, which you can track and manage through the portal.

GTRI support technicians will communicate with customers using the same methods noted above, and primarily via email through the ticketing system. To maintain a high level of customer service, our agents will focus all client communications through our ticketing system where possible. This is important so that all agents can view the full history of the issue. GTRI agents will work in the best interest of the partner and client to maximize the support experience for the partner and client.

If the customer needs to escalate the priority of a request with GTRI, they should do so by using one of the following methods:

- Reply to the email thread for the request, indicating the details of your escalation request (priority, critical time frames, etc.)
- Log on to the support portal online and add a note to your ticket with the details of your escalation request (priority, critical time frames, etc.)

- Log on to the support portal online and initiate a chat session with a support agent. Ask the agent to add a note to your ticket with the details of your escalation request (priority, critical time frames, etc.)
- Call the support phone line and when you are connected to an agent, ask the agent to locate your ticket (provide ticket number if possible), and ask them to update the ticket with the details of your escalation request (priority, critical time frames, etc.)

GTRI offers three tiers of support as follows:

- **Bronze** - 8x5 – Customers who require daytime support (9am – 5PM ET) excluding weekends.
- **Silver** - 12x7 – Customers who require daytime support with extended service hours (7am – 7pm ET) and including weekends.
- **Gold** - 24x7 – Customers who require coverage 24x7. This level of service is ideal for customers who have high impact applications or are involved in emergency response.

This SLA covers provision and support of the following services:

- GTRI Services Support Team Access
- Maintenance and Support of dependent services
- Development of new or enhanced services
- Advice and consultancy

The SLA remains valid until superseded by a revised agreement, which has been endorsed by relevant signatories from both parties. The agreement will be reviewed annually and applied to all Service Orders associated with this agreement.

Note that the response and resolution times stated within this SLA are sample times only and final SLAs will be negotiated with each Purchasing Entity. GTRI will provide reports to review actual SLA response times. Upon calling the GTRI Services Support Team the call will normally be answered within **60 seconds**. This may be longer if the lines are busy, but this shouldn't exceed **90 seconds**.

Upon calling the GTRI Services Support Team you will be asked to give basic details of the incident. You will be asked for details of the system - so please have this information ready. Once the incident has been logged you will be given a **ticket number**. This ticket number will be e-mailed to you and must be quoted on any future contact. Anyone wishing to speak directly with a technical expert must contact the general GTRI Services Support Team number first.

GTRI service bands Response and Resolution Times

Upon placing a call, you will be asked to assess the **Severity (Business Impact)** of the incident according to the levels indicated in the Business Impact table below. Because of this the NOC Team will allocate a **Priority** to the incident or request. The priority of the incident or request will determine the target response and resolution times as negotiated in your contract (see "Service Levels" below – all service levels are listed for ease of use and understanding.)

If the GTRI Services Support Team cannot immediately resolve the incident on the phone, you will receive a call back from a member of the GTRI Services Support Team or GTRI Services Support technical escalation team according to the priority level.

1.4.1.1. GTRI Response and Resolution Targets

Below are GTRI's sample SLA's for response and resolution

BRONZE Support Provided 8 hours per day, 5 days per week:

Priority	Response Time	Resolution Target
Urgent (1)	Portal (Immediate) / Phone 5 Min	4 hours
High (2)	Portal (Immediate) / Phone 10 Min	8 Hours
Medium (3)	Portal (Immediate) / Phone 30 Min	1 Business Day
Low (4)	Portal (Immediate) / Phone 30 Min	3 Business Days

SILVER Support Provided 12 hours per day, 7 days per week:

Priority	Response Time	Resolution Target
Urgent (1)	Portal (Immediate) / Phone 5 Min	2 hours
High (2)	Portal (Immediate) / Phone 10 Min	4 Hours
Medium (3)	Portal (Immediate) / Phone 30 Min	8 Hours
Low (4)	Portal (Immediate) / Phone 30 Min	2 Business Days

GOLD Support Provided 24 hours per day, 7 days per week:

Priority	Response Time	Resolution Target
Urgent (1)	Portal (Immediate) / Phone 5 Min	1 hours
High (2)	Portal (Immediate) / Phone 10 Min	2 Hours
Medium (3)	Portal (Immediate) / Phone 30 Min	4 Hours
Low (4)	Portal (Immediate) / Phone 30 Min	1 Business Day

Table 4 – Sample Customer Service SLAs

1.4.2. Ability to Comply with Customer Service Requirements (8.4.2)

GTRI will assign our Contract Manager to be the Lead Representative that Executes Participating Addendums. GTRI will keep the contact information for our Contract Manager current for all Purchasing Entities.

GTRI's customer service plan and sample response and resolution times are described in the Section 1.1.4 above. We will follow our help desk SLA's to meet the response time to inquiries within one business day.

GTRI provides design and installation services for cloud application environments. Available services include:

- **Architecture Workshops** - If a purchasing entity is seeking guidance on how to get started in the cloud, we provide tailored architecture workshops that can cover the service catalog, cloud economics and cost optimization, and application architecture in the cloud.
- **Landing Zone Implementation** – Starting out with solid design principles, an understanding of how the environment will grow over time, and how to build in security from day one is vital to operating successfully in the cloud. GTRI can help build your cloud “landing zone”.
- **Professional Services** – GTRI has a deep bench of AWS certified engineers who can build out environments on behalf of purchasing entities or augment their existing development, operations and security teams.
- **Managed Services** – GTRI can help purchasing entities focus on their missions by handling the day-to-day business of keeping a cloud environment secure, performant and operational.

Additional information on GTRI's design and installation services are provided and described in more detail in Section 3.3 below.

In addition to GTRI's Customer Service and Help Desk, we also provide Support through AWS to augment our teams for support in the physical infrastructure or service provided by AWS. AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced technical support engineers. The service helps customers of all sizes and technical abilities to successfully use the products and features provided by AWS. GTRI provides “Business” level support to customers by default, but we can also provide “Enterprise” level support if a customer wishes. Customers are welcome to work through GTRI or reach out to AWS directly if they wish.

AWS Support Features

AWS Support provides a highly personalized level of service for customers seeking technical help. All plans provide 24x7 access to customer service, AWS Documentation, Resource Center, Product FAQs, Discussion Forums, and support for Health Checks. A comparison between Business, and Enterprise support is available in the Table 5 – AWS Support Features below.

AWS Support Pricing

All AWS Support tiers include an unlimited number of support cases, with no long-term contracts. Also, with the Business and Enterprise-level tiers, as your AWS charges grow, you earn volume discounts on your AWS Support costs. Please note that these calculations are performed on a per-account basis. Check out the Simple Monthly Calculator for a customized estimate of your deployment's AWS Support cost. For additional information on AWS Support pricing, visit the AWS Support Plan Pricing webpage.

Contacting AWS Support

Customers can contact AWS Support via the Support Center. Business- and Enterprise-level customers may also “Click to Call” to have AWS contact them at any convenient phone number or strike up a conversation with an engineer via Chat. Enterprise-level customers also have a direct access to their dedicated TAM.

Chat is another way to contact AWS Support. By clicking on the chat support icon in the Support Center, a chat session will be initiated through the browser. This provides a real-time, one-on-one interaction with our support engineers and allows additional information and links to be shared for faster issue resolution.

	Business	Enterprise
Customer Service – 24x7x365	✓	✓
Support Forums	✓	✓
Documentation, Whitepapers, Best Practice Guides	✓	✓
Access to Technical Support	Phone, Chat, Email (24/7)	Phone, Chat, Email, Technical Account Manager (TAM) (24/7)
Primary Case Handling	Cloud Support Engineer	Sr. Cloud Support Engineer
Users Who Can Create Technical Support Cases	Unlimited (AWS Identity and Access Management [IAM] supported)	Unlimited (IAM supported)
Response Time	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour	General guidance: < 24 hours System impaired: < 12 hours Production system impaired: < 4 hours Production system down: < 1 hour Business-critical system down: < 15 minutes
Architecture Support	Contextual Use Case Guidance	Contextual Application Architecture Guidance
Access to Support API	✓	✓
Third-Party Software Support	✓	✓
AWS Trusted Advisor	Full checks	Full checks
Infrastructure Event Management	Contact Us for Pricing	✓

	Business	Enterprise
Direct Access to TAM		✓
Architectural Review		✓
Support Concierge		✓
Training		Access to online self-paced labs
Operations Support		Operational reviews, recommendations, and reporting

Table 5 – AWS Support Features

1.5. Security of Information (E) (8.5)

The National Institute of Standards and Technology (NIST) 800-53 security controls are generally applicable to Federal Information Systems. These are typically systems that must go through a formal assessment and authorization process to ensure sufficient protection of confidentiality, integrity, and availability of information and information systems, based on the security category and impact level of the system (low, moderate, or high), and a risk determination.

AWS' National Institute of Standards and Technology (NIST) compliant cloud infrastructure services have been validated by third-party testing performed against the NIST 800-53 Rev. 4 controls plus FedRAMP requirements. AWS has received FedRAMP Authorizations to Operate (ATO) from multiple authorizing agencies for both the AWS GovCloud (US) Region and the AWS US East/West regions. AWS' whitepaper *NIST Cybersecurity Framework (CSF) Aligning to the NIST CSF in the AWS Cloud* evaluates the NIST CSF and the many AWS Cloud offerings public and commercial sector customers can use to align to the NIST CSF to improve your cybersecurity posture. It also provides a third-party auditor letter validating attestation confirming AWS services' conformance to the NIST CSF risk management practices, allowing you to properly protect your data across AWS.

As an AWS customer, AWS' alignment with common NIST frameworks means that as you build systems and applications on AWS some controls are specifically inherited from AWS and many of the controls are shared inheritance between you and AWS. Under NDA, AWS provides an AWS FedRAMP SSP template based upon NIST 800-53 Rev. 4, which is prepopulated with the applicable NIST 800-5 Rev. 4 low/moderate/high control baseline. Control responsibility is as follows:

- **Shared Responsibility:** You will provide security and configurations of your software components and AWS will provide security for its infrastructure.
- **Customer-Only Responsibility:** You are fully responsible for guest operating systems, deployed applications, and select networking resources (for example, firewalls). More specifically, you are solely responsible for configuring and managing security "in" the cloud.
- **AWS-Only Responsibility:** AWS manages the cloud infrastructure, including the network, data storage, system resources, data centers, physical security, reliability, and supporting hardware and software. Applications built on top of the AWS system inherit the features and configurable options that AWS provides. AWS is solely responsible for configuring and managing security "of" the cloud.

For security authorization purposes, compliance with the FedRAMP requirements (based on NIST 800-53 rev 4 Low/Moderate/High control baseline) is contingent upon AWS fully implementing AWS-Only and Shared controls, and you implementing Customer-Only and Shared controls. A FedRAMP accredited 3PAO (Third Party Assessor Organization) has assessed and authorized AWS' implementation of their control responsibility. The portion of shared controls that you are responsibility for, and controls related to applications you implement on top of the AWS infrastructure, must be separately assessed and authorized by you in agreement with NIST 800-37 and customer-specific security authorization policies and procedures.

1.5.1. Measures GTRI Takes to Protect Data (8.5.1)

GTRI takes protecting its customers data extremely seriously and takes a "defense in depth" approach to data protection. GTRI starts by ensuring that data is encrypted at rest and in transit. We then harden all the customer's virtual servers. We build secure computing enclaves with highly segmented network environments to ensure that application servers are not accessible directly from the internet. We also continuously monitor and audit the secure configuration of the cloud environment.

Data at Rest

Encryption services are provided by AWS, but it is up to the customer to activate them and operate them securely. AWS customers can manage and control encryption keys themselves or let AWS manage keys for them. For example, AWS provides Server-Side Encryption (SSE) which is freely available for storage such as S3, EBS, and image snapshots.

GTRI recommends that our AWS cloud customers use AWS Key Management Service (KMS) whenever possible. AWS Key Management Service (KMS) is a FIPS compliant managed symmetric key service that allows customers to retain control of their regional Customer Master Key (CMK). The AWS S3 SSE-KMS encryption of objects leverages KMS and your CMK. Every customer gets a master key in the KMS and they can use this master key to create sub-keys on the keychain (hierarchy). It is then easy to perform a one-click encryption of server and database storage (RDS). Customers have centralized key management (create, delete, view, set policies) and enforced, automatic key rotation. All KMS usage is logged into the CloudTrail audit logging service for full transparency.

To destroy encrypted data, you simply delete the encryption key that was used to encrypt it and the data is "crypto-shredded" and impossible to recover.

AWS Cloud Hardware Security Modules (CloudHSMs) are AWS managed physical hardware appliances available for storing your self-managed symmetric and asymmetric keys. This may be appropriate for customers who already manage HSM infrastructure.

Data in Motion

GTRI uses both network micro-segmentation and encryption to protect data in motion. We use HTTPS/TLS/SSL for all Internet-facing web services. We use validated x.509 public certificates for these systems to facilitate end-to-end encryption of the communications over the Internet.

AWS offers certificate management and we use AWS Certificate Manager (ACM) to easily create and manage public certs. If the AWS Certificate Manager cannot work, we can also setup the customer's own certificates or automate the process using Let's Encrypt free and public CA service.

Secure Enclaves

GTRI builds out secure enclaves using the AWS QuickStarts for NIST-Based Assurance Frameworks. These are a collection of CloudFormation templates that create a secure cloud environment, ensuring that audit controls are configured properly, least-privilege access controls are implemented, groups and user roles are defined securely, and control-plane and data-plane traffic are properly isolated from each other. These templates have been assessed by a third-party assessment organization (3PAO) and come with a

completed Customer Responsibility Matrix and Security Controls Matrix to form the basis of the necessary security documentation to meet compliance needs in a cloud environment.

These templates also create a highly-segmented network environment with a DMZ tier to filter inbound internet traffic and highly-isolated database and application tiers that do not have direct access to the Internet. Firewalls, Network Access Control Lists and security groups are all configured to “deny by default” to drastically reduce the attack surface of the application and database hosts.

Instance Hardening

GTRI further protects data by controlling the OS configuration used by the EC2 instances. The method that GTRI helps our AWS cloud customers achieve this level of security is to use compute instance images that are hardened to Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standards. In the cloud, you are only as secure as your EC2 Instances base Amazon Machine Images (AMIs) (manage systems with templates not individually), so we help our customers build their own controlled and managed secure images off a default AMIs or import your own hardened AMI based on your preferred STIG. Then we save the AMIs and reuse them for other applications and services. However, we do not store security keys within your stored or shared (community) images.

Continuous Monitoring

In addition to ensuring an application environment is developed in the most secure manner possible continuous monitoring ensures that the environment stays that way. GTRI leverages both the AWS service catalog tools such as CloudWatch, AWS Config and Trusted Advisor and third-party security monitoring platforms like CloudCheckr to ensure that environments remain in a secure state. For customers with additional security needs, we leverage Splunk’s Enterprise Security platform, a market-leading IDS/IPS/SIEM solution.

1.5.2. How GTRI Complies with all Applicable Laws Related to Data Privacy and Security (8.5.2)

When it comes to performing security in cloud infrastructure, the shared security responsibility model is applicable. In an on-premises environment, the organization bears the full responsibility for security. With the cloud, the cloud customer bears more responsibility for security with IaaS environments rather than SaaS. Below is a picture that represents this concept.

	On Premises	IaaS	PaaS	SaaS
Security GRC				
Data Security				
App Security				
Platform Security				
Infrastructure Security				
Physical Security				

Enterprise Responsibility (diagonal line from top-left to bottom-right)

Shared Responsibility (diagonal line from top-right to bottom-left)

Provider Responsibility (diagonal line from bottom-left to top-right)

Figure 2 – AWS Isolation and Deployment Options

The method that GTRI clarifies the roles and responsibilities around security and documents the data privacy and security of cloud environments is by creating a cloud controls matrix. We will create a cloud security controls matrix that documents all the applicable laws and compliance requirements. This matrix

will document the responsible parties for each security control and the security measures taken to address that control item. An example of how we have done this in the past is using the FISMA documents for our U.S. federal cloud customers or using the Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM).

GTRI helps our AWS cloud customers maintain compliance with local government laws by encrypting the data in the cloud and by controlling access to the data. We create secure AWS environments based on the government law regulations, such as secure enclaves using AWS VPCs. For example, if the data is PCI information, then we use the AWS VPC structure that meets the PCI DSS regulations and we create a controls matrix that shows how that AWS environment configuration meets those regulations. We also encrypt the data and secure the encryption keys, as previously mentioned in this Section 1.5.1 of this document. We also control which AWS regions and availability zones are used by customers to keep the data confined to those geographies have specific data privacy laws, such as U.S. International Traffic in Arms Regulations (ITAR) data or data that falls under the EU General Data Protection Regulation (GDPR) regulations.

1.5.3. GTRI's Approach to Not Access Purchasing Entity's User Accounts or Data (8.5.3)

GTRI's staff does not access the data of our customers' cloud systems. We may be helping our clients to manage their IT infrastructure in AWS, but we do not access, view, or make copies of any customer data.

AWS does not access or use customer content for any purpose other than as legally required and to provide the AWS services selected by each customer, to that customer and its end users. AWS never uses customer content or derives information from it for other purposes such as marketing or advertising.

Furthermore, to make it impossible for AWS or GTRI to access the customer's data, our standard operating procedure is to encrypt all data at rest using customer-managed encryption keys.

Because AWS customers retain ownership and control over their content within the AWS environment, they also retain responsibilities relating to the security of that content as part of the AWS "shared responsibility" model. While AWS manages security of the cloud, security in the cloud is the responsibility of the customer, as customers retain control of what security they choose to implement to protect their own content, platform, applications, systems and networks – no differently than they would for applications in an on-site data center.

Following is a picture of the "shared responsibility model" for security as it relates to AWS and its services. It is also important to note that the responsibility varies based on the AWS service being used. For example, but the AWS customer bears more responsibility for securing and configuring their EC2 compute instances. Whereas the AWS customer bears little responsibility for S3 storage, other than configuring the bucket access policy and choosing the encryption method and enabling bucket versioning.




Figure 3 – AWS Isolation and Deployment Options

To assist customers in designing, implementing and operating their own secure AWS environment, AWS provides a wide selection of security tools and features customers can use. Customers can also use their own security tools and controls, including a wide variety of third party security solutions. Customers can configure their AWS services to leverage a range of such security features, tools and controls to protect their content, including sophisticated identity and access management tools, security capabilities, encryption and network security.

1.6. Privacy and Security (E) (8.6)

1.6.1. Commitment to Comply with NIST 800-145 (8.6.1)

Please see Section 1.1.1 above in this response for a detailed overview of how GTRI's cloud service providers provide all the essential capabilities of cloud computing through documented Application Programming Interfaces (APIs).

As for other relevant cloud and security standards, GTRI is an AWS Public Sector Partner and an AWS Advanced Technology Partner. Therefore, we are intimately familiar with NIST and FISMA and FedRAMP security standards and how to implement these security controls within an AWS account. GTRI currently administers AWS GovCloud accounts for U.S. Federal organizations and GTRI bears the responsibility for the customer "shared responsibility model" security controls. We have obtained a multi-year Authority to Operate (ATO) government workloads in AWS.

GTRI employs cloud consultants who possess the Cloud Security Alliance (CSA) Certificate of Cloud Security Knowledge (CCSK) and the (ISC)2 Certified Cloud Security Professional (CCSP) certifications. The NIST cloud specifications are within the body of knowledge covered in those tests.

The NIST SP 800-145, The NIST Definition of Cloud Computing, covers the high-level concepts of cloud computing. While we are familiar with this publication, we are also familiar with other documents published by the NIST Cloud Computing Public Security Working Group, such as the NIST SP 800-144, Guidelines on Security and Privacy in Public Cloud Computing and the NIST SP 500-299, NIST Cloud Computing Security Reference Architecture which cover the security recommendations.

1.6.2. All Government Standards or Standard Organization Security Certifications (8.6.2)

GTRI has chosen AWS as our preferred Infrastructure as a Service (IaaS) Cloud Service Provider because of their extensive security accreditations and compliance certifications. Because GTRI services so many U.S. federal government departments and agencies, it is important for us to work with a CSP that meets the government's cloud security requirements. AWS GovCloud (US) is an isolated region that meets FedRAMP, FIPS, ITAR, FISMA, and NIST 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 picture below shows some of the popular AWS compliance certifications and accreditations.



Figure 4 – AWS Isolation and Deployment Options

AWS Compliance enables customers to understand the robust controls in place at AWS that facilitate security and data protection in the cloud. 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)
- System and Organization Controls (SOC) 1, SOC 2, and SOC 3
- Payment Card Industry Data Security Standard (PCI DSS)
- International Organization for Standardization (ISO) 27001, 27017, 27018, and 9001
- Department of Defense (DoD) Security Requirements Guide (SRG) Impact Levels 2, 4, 5, and 6
- 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)

For information on all the security regulations and standards with which AWS complies, visit the AWS Compliance page.

1.6.3. Security Practices in Place to Secure Data and Applications (8.6.3)

The techniques that GTRI uses to secure data and applications are described in detail in Section 1.5.1 The AWS platform is the bases of the data protection techniques that we employ. AWS provides many additional capabilities to secure systems and protect data.

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.

AWS provides standard DDoS mitigation, MITM, IP spoofing, port-scanning, and packet sniffing for Virtual Private Clouds (VPCs). AWS offers both its free AWS Shield Standard and it's for-a-fee Shield Advanced service to assist customers with mitigating DDoS attacks. It is also possible to use auto-scaling policies to absorb a DDoS attack by rescaling the instance size with "Enhanced Networking" or scaling the pool of EC2 instances with Elastic Load Balancer (ELB). Furthermore, the AWS ELB can only forward sane TCP connections, whereby, SYN floods and other DDoS packets (UDP reflection, ICMP flood) are dropped. The AWS CloudFront CDN service with AWS WAF can block attacks from AWS edge locations. AWS Route 53 can absorb DNS flooding attacks through shuffle sharding and anycast striping.

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. 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.

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.

Another method of providing security for web-based applications is to use the AWS Web Application Firewall (WAF) service. The AWS WAF service integrates with the Application Load Balancer (ALB) or CloudFront CDN. The AWS WAF can be configured using the Managed WAF service to protect your web applications (e.g. OWASP Top 10). The WAF can be configured by rules that inspect CF traffic using Conditions which match IPs, strings, regex, etc., and configured with Rules which contain conditions with Boolean logic and culminate in the configuration of Web ACLs which contain rules and actions (allow, block, count) (default rule = allow all). AWS WAS Managed Rules configure your WAF with threat intel from vendors such as: Trend Micro, Imperva, Fortinet, Alert Logic, F5, TrustWave. AWS WAF Security Automations offer CloudFormation Templates to launch Lambda functions that monitor and dynamically configure the WAF service. AWS Firewall Manager can be used to managed WAF configurations across accounts/regions.

AWS GuardDuty is a managed threat detection service that operates continuously and monitors for malicious or unauthorized behavior to help customers protect their AWS accounts and workloads. GuardDuty monitors for suspicious activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. GuardDuty can also detect potentially compromised instances or reconnaissance by attackers. GTRI will commonly enable this on the AWS accounts for our customers.

1.6.4. Data Confidentiality Standards and Practices in Place to Ensure Data Confidentiality (8.6.4)

GTRI maintains the confidentiality and integrity of the information we store regarding our customer's AWS environments. GTRI secures this information on U.S. based systems and restricts the GTRI employees that have access to this information. GTRI has a System Security Plan that meets NIST 800-171 requirements for all customer data. GTRI also has its own IT practices of encrypting data at rest and on our employee laptops and mobile devices.

GTRI has a set of corporate security policies that define how GTRI employees will use the GTRI IT systems and how they treat the confidentiality of customer information. GTRI has a documented Acceptable Use Policy that defines how employees are allowed to use GTRI IT assets. GTRI's Information Security Policy defines how data is classified and categorized. GTRI marks data as either Public, Internal, GTRI Restricted, or Customer Restricted. This Information Security Policy also details GTRI's access control policies and how passwords, user accounts, and access is controlled. GTRI has a Data Classification Policy that further defines how data is classified, data ownership and how each of these 4 data categories are used and restricted. GTRI also has a Data Classification, Handling and Disposal policy that defines how, based on the same data categories data is created, stored, used, backed up, archived, and eventually deleted. Our internal encryption practices meet FIPS 140.2 encryption requirements.

GTRI's cloud architects and cloud administrators must follow all GTRI security policies and guidelines as a condition of their employment.

AWS does not access customer data, and customers are given the choice as to how they store, manage and protect their data.

Above in this document, in Section 1.5.1 of this response we detailed our methods for using encryption to protect data in transit and at rest to preserve the confidentiality of the information.

Below in Section 1.6.7 we discuss how user authentication is performed and how access controls and authorization is controlled using IAM access controls

GTRI's own IT practices of encrypting data at rest and on our employee laptops. Our internal encryption practices meet FIPS 140.2 encryption requirements. Additionally, we have a System Security Plan that meets NIST 800-171 requirements for all customer data.

1.6.5. Third Party Attestations, Reports, Security Credentials, and Certifications Relating to Data Security, Integrity, and Other Controls (8.6.5)

GTRI employs cloud consultants who possess the Cloud Security Alliance (CSA) Certificate of Cloud Security Knowledge (CCSK) and the (ISC)2 Certified Cloud Security Professional (CCSP) certifications. GTRI's cloud security experts also possess the AWS Certified Security - Specialty certification. These advanced cloud security certifications are validation that our cloud architects know the proper cloud security controls and how to apply these best practices in an AWS IaaS environment.

Earlier in this document in Section 1.6.2 above in this document we described our methods of meeting compliance in AWS cloud environments.

See information provided in response to Section 1.6.1 for a detailed list of the compliance programs we meet. Also, available documentation can be downloaded from AWS website

1.6.6. Logging Process Including Types of Services and Devices Logged (8.6.6)

The logging and monitoring of Application Program Interface (API) calls are key components in security and operational best practices, as well as requirements for industry and regulatory compliance. AWS

customers can leverage multiple AWS features and capabilities, along with third-party tools, to monitor their instances and manage/analyze log files.

AWS CloudTrail

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

With AWS CloudTrail, customers can get a history of AWS API calls for their account, including API calls made 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.

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

AWS CloudTrail: Features and Benefits

Some of the many features of AWS CloudTrail include:

- **Increased Visibility:** AWS CloudTrail provides increased visibility into user activity by recording AWS API calls. Customers can answer questions such as, what actions did a given user take over a given time period? For a given resource, which user has taken actions on it over a given time period? What is the source IP address of a given activity? Which activities failed due to inadequate permissions?
- **Durable and Inexpensive Log File Storage:** AWS CloudTrail uses Amazon S3 for log file storage and delivery, so log files are stored durably and inexpensively. Customers can use Amazon S3 lifecycle configuration rules to further reduce storage costs. For example, customers can define rules to automatically delete old log files or archive them to Amazon Glacier for additional savings.
- **Easy Administration:** AWS CloudTrail is a fully managed service; customers simply turn on AWS CloudTrail for their account using the AWS Management Console, the Command Line Interface, or the AWS CloudTrail SDK and start receiving AWS CloudTrail log files in the specified Amazon S3 bucket.
- **Notifications for Log File Delivery:** AWS CloudTrail can be configured to publish a notification for each log file delivered, thus enabling customers to automatically take action upon log file delivery. AWS CloudTrail uses the Amazon Simple Notification Service (Amazon SNS) for notifications.
- **Choice of Partner Solutions:** Multiple partners including AlertLogic, Boundary, Loggly, Splunk, and Sumologic offer integrated solutions to analyze AWS CloudTrail log files. These solutions include features like change tracking, troubleshooting, and security analysis. For more information, see the AWS CloudTrail partners section.

- **Log File Aggregation:** AWS CloudTrail can be configured to aggregate log files across multiple accounts and regions so that log files are delivered to a single bucket. For detailed instructions, refer to the Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket section of the user guide.

Amazon CloudWatch

Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications 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. Customers 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.

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

LogAnalyzer for Amazon CloudFront

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

Amazon CloudFront's Access Logs provide detailed information 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.

Reports Generated

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

VPC Flow Logs

VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data is stored using Amazon CloudWatch Logs. GTRI enables VPC Flow Logs in every AWS account VPC to be able to gain visibility to the connections leaving the VPC virtual networks. This flow data can be collected and analyzed and reported on with a variety of utilities.

Macie

Another useful AWS security monitoring service is Amazon Macie. Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

Splunk Enterprise Security

For customers who need more from their continuous monitoring system than the functionality available in the AWS service catalog, we recommend that they leverage Splunk to correlate metrics across their application infrastructure and provide real-time intelligence. Splunk Enterprise Security (ES) enables security teams to use all data to gain organization-wide visibility and security intelligence. Splunk ES can be used for continuous monitoring, incident response, running a security operations center or for providing executives a window into business risk.

Splunk ES provides organizations the ability to:

- Improve security operations with faster response times
- Improve security posture by getting end-to-end visibility across all machine data
- Increase detection and investigation capabilities using advanced analytics
- Make better informed decisions by leveraging threat intelligence

GTRI, as a Splunk Elite Partner can install, configure and manage a customer's Splunk installation, ensuring they receive the maximum value.

1.6.7. Restricting Visibility of Cloud Hosted Data and Documents to Specific User Groups (8.6.7)

AWS Identity and Access Management (IAM) enables customers to securely control access to AWS services and resources for their users. Using IAM, customers can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources. IAM allows you to:

- **Manage IAM users and their access** – You can create users in IAM, assign them individual security credentials (i.e., access keys, passwords, and multi-factor authentication devices) or request temporary security credentials to provide users access to AWS services and resources. You can manage permissions in order to control which operations a user can perform.
- **Manage IAM roles and their permissions** – You can create roles in IAM and manage permissions to control which operations can be performed by the entity, or AWS service, that assumes the role. You can also define which entity is allowed to assume the role.
- **Manage federated users and their permissions** – You can enable identity federation to allow existing identities (e.g., users) in your enterprise to access the AWS Management Console, to call AWS APIs, and to access resources, without the need to create an IAM user for each identity.

GTRI knows that stringent Identity and Access Management (IAM) practices are a MUST when working in an AWS environment. We use IAM policies to control users, groups, permissions, and accounts that run services on AWS resources and rotate Access Keys & Secrets used for API calls. We enable Security Token Service (STS) as a web service that grants requests for temporary, limited privilege credentials for IAM users/roles. GTRI restricts AWS accounts so that no one could use the master payer account, root privileges will not be used, developer accounts need only specific privileges, create general use accounts for each sys-admin or service accounts, and we configure a password policy. We use Multi-Factor

Authentication (MFA) for master account & admin accounts. We can also configure federated identity access for management console and APIs such as with: SAML 2.0, OpenID Connect (OIDC), AWS Microsoft AD Connector (ADFS).

Recommended Best Practices:

- Avoid the use of the "root" account
- Ensure multi-factor authentication (MFA) is enabled for all IAM users that have a console password

Forced IAM User Self-Service Remediation

- Ensure credentials unused for 90 days or greater are disabled
- Ensure access keys are rotated every 90 days or less
- Ensure IAM password policy requires at least one uppercase letter, one lower case letter, one symbol and one number. Minimum password length set to 14 or greater
- Ensure IAM password policy prevents password reuse
- Ensure IAM password policy expires passwords within 90 days or less
- Ensure no root account access key exists
- Ensure MFA is enabled for the "root" account
- Ensure security questions are registered in the AWS account
- Ensure IAM policies are attached only to groups or roles
- Enable detailed billing
- Ensure IAM Master and IAM Manager roles are active
- Ensure security contact information is registered
- Ensure a support role has been created to manage incidents with AWS Support
- Ensure IAM policies that allow full "*" administrative privileges are not created
- Ensure CloudTrail is enabled in all regions
- Ensure CloudTrail log file validation is enabled
- Ensure the S3 bucket CloudTrail logs to is not publicly accessible
- Ensure CloudTrail trails are integrated with CloudWatch Logs
- Ensure AWS Config is enabled in all regions
- Ensure S3 bucket access logging is enabled on the CloudTrail S3 bucket
- Ensure CloudTrail logs are encrypted at rest using KMS CMKs
- Ensure rotation for customer created CMKs is enabled
- Ensure a log metric filter and alarm exist for unauthorized API calls
- Ensure a log metric filter and alarm exist for Management Console sign-in without MFA
- Ensure a log metric filter and alarm exist for usage of "root" account
- Ensure a log metric filter and alarm exist for IAM policy changes
- Ensure a log metric filter and alarm exist for CloudTrail configuration changes

- Ensure a log metric filter and alarm exist for AWS Management Console authentication failures
- Ensure a log metric filter and alarm exist for disabling or scheduled deletion of customer created CMKs
- Ensure a log metric filter and alarm exist for AWS Config configuration changes
- Ensure a log metric filter and alarm exist for changes to Network Access Control Lists (NACL)
- Ensure a log metric filter and alarm exist for changes to network gateways
- Ensure a log metric filter and alarm exist for route table changes
- Ensure a log metric filter and alarm exist for VPC changes
- Ensure appropriate subscribers to each SNS topic
- Ensure no security groups allow ingress from 0.0.0.0/0 to port 22
- Ensure no security groups allow ingress from 0.0.0.0/0 to port 3389
- Ensure VPC flow logging is enabled in all VPCs
- Ensure the default security group of every VPC restricts all traffic
- Ensure routing tables for VPC peering are "least access"

1.6.8. Notification Process in the Event of a Security Incident (8.6.8)

GTRI has described our notification process as part of our Incident Response plan in Section 1.3.1 above. In addition to our Incident Response plan, AWS has 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.

GTRI also implements what we call "event-driven security", whereby we use automation functions such as AWS Lambda to observe the configuration of the AWS environment or observe events taking place in logs and automatically alert and possibly even correct the misconfiguration. We use AWS Simple Notification Service (SNS) to send messages via e-mail or mobile messages to notify security issues.

1.6.9. Security Controls, Both Physical and Virtual, Used to Isolate Hosted Servers (8.6.9)

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 3.2 published in April 2016. More information on AWS's multi-tenant architecture is found in the *AWS Risk and Compliance* whitepaper.

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. AWS isolation and deployment options are illustrated below in Figure 5.




Figure 5 – AWS Isolation and Deployment Options

At GTRI, we use Virtual Private Clouds (VPC) and the arrangement of VPCs within an AWS account to create a secure enclave environment. This environment places the valuable application components in between the separated management VPC and separated from the Internet-facing DMZ VPC. In this way, we create an orbital security model that protects the applications within the AWS virtual networking environment.

One example of how we do this is by using the AWS NIST QuickStart Templates. These are a set of freely-available AWS CloudFormation Templates that help rapidly deploy AWS infrastructure while meeting the security compliance requirements. When working with SLED or federal government organizations, we use the Standardized Architecture for NIST-based Assurance Frameworks on the AWS Cloud: Quick Start Reference Deployment. These templates help rapidly deploy baseline VPCs based on 800-53/171 and FedRAMP/FISMA standards (DOD SRG).

1.6.10. Technical Reference Architectures that Support IaaS, SaaS, and PaaS (8.6.10)

GTRI follows several architectures and frameworks for cloud deployments.

One resource and model for cloud architectures is AWS's Cloud Adoption Framework (CAF). The AWS Cloud Adoption Framework (AWS CAF) is created to help organizations develop efficient and effective plans for their cloud deployments. AWS Cloud Adoption Framework (AWS CAF) organizes guidance into six areas of focus, called perspectives. The Business, People, and Governance Perspectives focus on business capabilities; while the Platform, Security, and Operations Perspectives focus on technical capabilities.

The other resource we utilize is the AWS Well-Architected Framework. The Well-Architected Framework has been developed to help cloud architects build the most secure, high-performing, resilient, and efficient infrastructure possible for their applications. This framework provides a consistent approach for customers and partners to evaluate architectures and provides guidance to help implement designs that will scale with your application needs over time. The Well-Architected Framework is broken down into five pillars: Operational Excellence, Security, Reliability, Performance Efficiency, and Cost Optimization. We follow the recommendations and cloud architectures documented in each of these pillars when constructing AWS cloud environments for our customers.

Another resource is the AWS Architecture Center which provides cloud architectures for various types of cloud environments such as IaaS, PaaS, and SaaS, as well as container and micro-services architectures.

GTRI also follows the best practices when it comes to cloud design patterns. We utilize these design pattern when deploying applications of various types into the cloud.

GTRI considers all these cloud architecture best practices and assists our customers with selecting the best architecture for their applications. Then GTRI implements the cloud environments and configures the appropriate security controls based on the architecture. GTRI uses methods like VPCs and secure enclaves to help protect application virtual server instances from the public cloud.

1.6.11. Security Procedures in Place Regarding Employees Who Have Access to Sensitive Data (8.6.11)

GTRI conducts background investigations on every employee as a condition of hire. Additionally, GTRI's employees who access customer sensitive data has updated background investigations every five (5) years by government investigators.

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

1.6.12. Security Measures and Standards in Place to Secure Confidentiality of Data and Rest and in Transit (8.6.12)

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. AWS offers the ability to add an additional layer of security to data at rest in the cloud by providing scalable and efficient encryption features. This includes:

- Data encryption capabilities in AWS storage and database services, such as Amazon EBS, Amazon S3, Amazon Glacier, Oracle RDS, SQL Server RDS, and Amazon Redshift.
- Flexible key management options, including AWS Key Management Service, that allow customers to choose whether to have AWS manage the encryption keys or keep complete control over their keys.
- Dedicated, hardware-based cryptographic key storage using AWS CloudHSM, allowing customers to satisfy compliance requirements.

Cloud applications often communicate over public links, such as the Internet, so it is important to protect data in transit when you run applications in the cloud.

This involves protecting network traffic between clients and servers, and network traffic between servers.

Services from AWS provide support for both IPsec and 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. AWS encryption in transit features:

- Support for both Internet Protocol Security (IPsec) and Secure Sockets Layer/Transport Layer Security (SSL/TLS) for protection of data in transit.

- APIs for customers to integrate encryption and data protection with any of the services developed or deployed in an AWS environment.

The *AWS Security Best Practices* whitepaper provides greater detail on how to protect data in transit and at rest in the AWS Cloud. Other security resources are also available on our Cloud Security Resources page.

Please refer to the Section 1.5.1 above in this document where this topic of data confidentiality and encryption was discussed extensively.

1.6.13. Policies and Procedures Regarding Notifications of Data Breach (8.6.13)

GTRI follows our Incident response plan for notifications of a Data Breach. Our incident response plan can be found in Section 1.3.1 above.

AWS is certified as a PCI DSS 3.2 Level 1 Service Provider, the highest level of assessment available. The compliance assessment was conducted by Coalfire Systems Inc., an independent Qualified Security Assessor (QSA). The PCI DSS Attestation of Compliance (AOC) and Responsibility Summary are available to customers by using AWS Artifact, a self-service portal for on-demand access to AWS compliance reports.

When it comes to PCI DSS, GTRI intends to use the PCI DSS QuickStart CloudFormation Templates to help ensure that the security controls of the underlying AWS environment are applied.

1.7. Migration and Redeployment Plan (E) (8.7)

1.7.1. Managing End of Life Activities Closing Down a Service and Safely Deprovisioning (8.7.1)

Because AWS customers retain ownership and control over their content within the AWS environment, they also retain responsibilities relating to the security of that content as part of the AWS “shared responsibility” model. While AWS manages security of the cloud, security in the cloud is the responsibility of the customer, as customers retain control of what security they choose to implement to protect their own content, platform, applications, systems and networks – no differently than they would for applications in an on-site data center.

To assist customers in designing, implementing and operating their own secure AWS environment, AWS provides a wide selection of security tools and features customers can use. Customers can also use their own security tools and controls, including a wide variety of third party security solutions. Customers can configure their AWS services to leverage a range of such security features, tools and controls to protect their content, including sophisticated identity and access management tools, security capabilities, encryption and network security.

AWS does not access or use customer content for any purpose other than as legally required and to provide the AWS services selected by each customer, to that customer and its end users. AWS never uses customer content or derives information from it for other purposes such as marketing or advertising.

AWS Customers manage the creation and deletion of their data on AWS, as well as maintain control of access permissions. Customers are responsible for maintaining appropriate data retention policies and procedures. 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, and FedRAMP audits.

One of the best ways to protect data at rest is with encryption. Encryption services are provided by AWS, but it is up to the customer to operate them securely and protect the keys. AWS customers can manage

and control the keys yourself or let AWS handle it for them. For example, AWS provides Server-Side Encryption (SSE) which is freely available for storage such as S3, EBS, and image snapshots.

AWS Cloud Hardware Security Modules (HSMs) are AWS managed physical hardware appliances available for storing your self-managed symmetric and asymmetric keys.

GTRI often recommends that our AWS cloud customers use AWS Key Management Service (KMS). AWS Key Management Service (KMS) is a managed symmetric key service that allows customers to retain control of your regional Customer Master Key (CMK). The AWS S3 SSE-KMS encryption of objects leverages KMS and your CMK. Every customer gets a master key in the KMS and they can use this master key to create sub-keys on the keychain (hierarchy). It is then easy to perform a one-click encryption of server and database storage (RDS). Customers have centralized key management (create, delete, view, set policies) to enforced, automatic key rotation. AWS KMS has visibility into any changes via CloudTrail.

Furthermore, the best way to protect data in transit is with encryption. We use HTTPS/TLS/SSL for all Internet-facing web services. We use validated x.509 public certificates for these systems to facilitate end-to-end encryption of the communications over the Internet.

AWS offers certificate management and we use AWS Certificate Manager (ACM) helps easily create and manage public certs. We can also setup the customer's own certificates or automate the process using Let's Encrypt free and public CA service.

When it comes to destroying the encrypted data, you simply delete the encryption key that was used to encrypt it and the data is essentially "crypto-shredded".

1.7.2. Orderly Return of Data Back to Purchasing Agency (8.7.2)

GTRI can suggest a variety of options to return data back which are described in the table below.

Available Method	Description
AWS VPN	The VPN connection lets you bridge your AWS VPC and IT infrastructure. Data can be returned over IPSeC VPN tunnel with AES256 encryption. After transfer, data will be "crypto-shredded". This solution can be used for small amount of data like file servers less than 2 TB.
S3	Data from EFS, VM, DB backups can be archived and placed on encrypted S3. After that archives can be transferred to customer over HTTPS with IAM authorization. After transfer, data will be "crypto-shredded". This method could be used for small infrastructures.
Storage Gateway	With Storage Gateway your cloud data could be synchronized with local virtual appliance. After synchronization data can be transferred from local virtual appliance and "crypto-shredded". In the cloud. This solution can be used for datasets up to 30 TB.

Available Method	Description
Snowball and Snowball EDGE	AWS Snowball is a service that accelerates transferring large amounts of data into and out of AWS using physical storage appliances, bypassing the internet. AWS Snowball Edge is a 100 TB data transfer device with on-board storage and compute power for select AWS capabilities. Encryption is enforced, protecting your data at rest and in physical transit. Can return to customer all pieces of data like, volumes, VM images, DBs, EFS etc. After transfer, data will be “crypto-shredded”. This solution can be used for datasets from 50TB, 80 TB, and 100 TB.

Table 6 –Return of Backup Data

1.8. Service or Data Recovery (E) (8.8)

GTRI understands that purchasing entities will host applications with differing levels of criticality, from low-impact applications to mission critical applications. During the planning phases of a cloud deployment, GTRI will gather information regarding the criticality levels of the applications to be deployed. If that information does not exist, GTRI can help the purchasing entity determine the criticality of an application. GTRI can support any RTO/RPO that a customer requires using the tools available in the AWS environment, and will work with the customer to define the cloud architecture that meets the customer’s need at the lowest cost.

In general, GTRI encourages customers to design for resilience and high-availability, especially in a cloud environment. In a legacy datacenter environment, it can be prohibitively expensive to design a highly-available architecture because it relies on expensive physical hardware and load balancing equipment. AWS provides technologies that can make highly-available architectures available at very little extra cost. Load balancers are available as a managed service for a minimal charge. It is possible to use auto-scaling technology to ensure that the purchasing entity does not pay for excess server capacity while maintaining multiple servers in disparate geographical locations. AWS refers to applications that use multiple Availability Zones (datacenters), load balancing, and automated fail-over as “Well-Architected” for a cloud environment.

In addition to multiple datacenters, AWS also provides technologies such as cross-region replication and global load balancing if a customer has compliance needs, or exceptionally high availability needs. This allows an application to operate in multiple, geographically distinct (West Coast/Midwest/East Coast) regions.

Ability to recover and restore data within 4 business hours in the event of a severe system outage

The exact definition of what a *severe system outage* is will depend on the customer and the application. During the design phase of a project, GTRI will work with the customer to determine exactly what constitutes a severe system outage and deploy monitoring to ensure that GTRI is notified automatically when such an outage occurs.

GTRI encourages customers to design for High Availability (HA) in the cloud. In the case that an application is operating in multiple Regions or Availability Zones simultaneously, all data is automatically replicated to an alternate site, so there is no need to recover data. All data will be served from the operational servers and datacenters without downtime.

In the “worst-case scenario” - a legacy application running in a cloud environment that cannot be run in an HA architecture relying on snapshots for recovery, GTRI will do the following at a minimum:

- Create automated file system snapshots to support customer’s RPO/RTO objectives

- Schedule automated database backups to support customers RPO/RTO objectives
- Script the creation of the customer's network and server infrastructure using automation technologies such as AWS CloudFormation and Ansible

GTRI can, immediately after learning of an outage, recover data create machine images (AMIs) from the latest backups and snapshots. This process typically takes minutes to complete. At this point, it is possible to launch a new set of virtual servers that are current as of the latest snapshot time.

If a customer is running a mission-critical legacy application, GTRI can use third-party host-level replication technology to provide real-time replication of an application to a standby environment. This approach enables recovery times in minutes at an additional cost.

Describe your RPO and RTO objectives

In general, the shorter the RPO and RTO objectives are, the greater the cost to the customer. GTRI will work with the customer during the design phase of a project to determine what objectives are best for their application and budget. However, GTRI has some standardized architectures that support different RPO/RTO categories.

Cloud-Native Highly-Available Architecture (RPO/RTO = less than 5 minutes)

This is GTRI's preferred architecture because it provides the best RTO/RPO objectives at the lowest price point, but it also requires work to ensure that the application is architected to be "Cloud Native". We achieve these RPO and RTO objectives by leveraging AWS platform services that are highly available and following AWS' Well Architected Framework. By storing web content on Amazon S3, database content using a highly-available relational database configuration (Multi-AZ Relational Database Service or DynamoDB), and leveraging auto-scaling for virtual servers, we can eliminate single points of failure for our applications.

Legacy application with snapshots (RPO/RTO = 4 hours)

Many legacy applications are difficult to deploy in a redundant manner. If an application stores state on a server's local file system, has hard-coded configuration or cannot be load-balanced for some technical reason, then the application will need rework to be optimized for a cloud environment. If that cannot happen prior to deployment, the most economical option is to perform periodic snapshots and use those to restore the application in the event of downtime. This is often referred to as having a "cold standby"

Legacy application with host-level replication (RPO = less than 5 minutes, RTO = less than 1 hour)

Through the AWS marketplace, customers can purchase software to replicate data in real time to another location, either within the same AWS region or in a different AWS region, and automate the recovery of the application in the event of a failure. The exact recovery point depends on the rate of change of the data being replicated, but it is typically in the range of a few seconds. The exact recovery time depends on the scale of the application, and whether the customer wishes to pay for the standby application to be running (hot standby – recovery time = seconds) or wishes to keep the standby application shut down and launch it only in the event of an emergency (warm standby or pilot light – recovery time = less than 1 hour).

Cloud-Native Highly-Available Architecture (RPO/RTO = less than 5 minutes).

Legacy application with snapshots (RPO/RTO = 4 hours).

Legacy application with host-level replication (RPO = less than 5 minutes, RTO = less than 1 hour).

1.8.1. Contingency Plan to Respond to Certain Situations (8.8.1)

Contingency planning is vital when hosting applications with business impact. Since organizations, workloads and criticality levels differ, GTRI expects the details of contingency planning to differ on a case-by-case basis. However, we do have a standard approach that we follow, and have included a basis for customization with customers below.

1.8.1.1. Template Contingency Plan Overview

GTRI, serving as the managed service provider for the purchasing entity will be the customer's primary point of contact throughout the execution of the contingency plan. When GTRI or the customer is made aware of downtime and have decided to invoke the contingency plan, GTRI and the customer will communicate using the helpdesk system as described in Section 1.4.1. GTRI will also engage AWS Customer Service to assist in the execution of the contingency plan or facilitate a return to normal service. AWS Staff operators provide 24x7x365 coverage to detect incidents and to manage the impact and resolution. Also, the Amazon Incident Management team employs industry-standard diagnostic procedures to drive resolution during business-impacting events.

A disaster is any event resulting in damage, loss, or destruction of property, processing resources, or processing services that significantly limits the ability of an application system or IT assets to carry on their usual business. In this plan, disaster classifications are described in terms of the effect that an event would have on this system's operations. The contingency plan consists of the following steps:

1. Notify stakeholders of a disaster and intent to invoke the contingency plan
2. Assess the disaster and its impact
3. Recovery Operations
4. Return to normal service

1.8.1.2. Notify Stakeholders of a disaster

The notification sequence is listed below:

- Notify Shareholders that an incident has occurred.
- Contact GTRI / AWS to gather impact analysis.
- Provide Shareholder with assessment details.

The Contingency Plan is to be activated if one or more of the following criteria are met:

- Stakeholders provide directive to activate contingency plan.

Upon directive to activate the contingency plan the following actions shall take place:

- If the plan is to be activated, the Contingency Planning Coordinator is to notify all Team Leaders and inform them of the details of the event and if relocation is required.
- Upon notification from the Contingency Planning Coordinator, Team Leaders are to notify their respective teams. Team members are to be informed of all applicable information and prepared to respond and relocate if necessary.
- The Contingency Planning Coordinator is to notify the off-site storage facility that a contingency event has been declared and to ship the necessary materials (as determined by damage assessment) to the alternate site.
- The Contingency Planning Coordinator is to notify the Alternate site that a contingency event has been declared and to prepare the facility for the Organization's arrival.

- The Contingency Planning Coordinator is to notify remaining personnel (via notification procedures) on the general status of the incident.

1.8.1.3. Assess the Disaster and its Impact

The first step in recovering from a disaster is to gain *situational awareness*. GTRI will review monitoring for applications and consult with AWS technical support to understand the scale of the disaster and the services that are affected. GTRI will classify disasters on the following scale:

Disaster Recovery Classifications

- **Catastrophic Disaster** - Catastrophic disasters necessitate the recovery of a defined minimal set of services at an alternate location followed by an eventual restoration of all services. An example of a catastrophic disaster would be the failure of an entire AWS region due to a large-scale natural disaster.
- **Critical Disaster** - Critical Disasters necessitate the recovery of a potentially smaller set of services at an alternate site. An example of a critical disaster would be the failure of a core AWS platform service, such as S3.
- **Limited Disaster** - Limited disasters necessitate the restoration of all services at the primary site. An example of a limited disaster would be the failure of an AWS availability zone due to a network issue or a natural disaster.

Impact Analysis

The purpose of an Impact Analysis is to correlate specific system resources with the critical services that they provide and, based on that information, to characterize the consequences of a disruption to each Resource.

To complete this section of the Contingency Plan Template, GTRI will work with the customer to classify applications by their impact to the organization. This will help GTRI prioritize the process of recovering applications.

Assessment Outcomes

Based on the information discovered about the scale of the disaster, the services affected and their impact to the customer, GTRI and the customer will agree on recovery priorities for the system. High priorities are based on the need to restore critical resources within their allowable outage times; moderate and low priorities reflect the requirement to restore full operational capabilities over a longer recovery period. All of this information will come together to form a tactical service recovery plan.

1.8.1.4. Recovery Operations

Starting with the highest priority application, GTRI will work with AWS and the customer to restore the applications to service. The exact methods used to restore an application to service will depend on the disaster recovery strategy selected for the application; details on these strategies can be found in Section 1.8 above. Procedures and runbooks to restore service to each application will be documented during the design and implementation phase of the project, and we will test the contingency plan on a regular basis.

Recovery Steps

While every application will differ in some details, we will take steps such as the following to restore service to an application:

- Pull backup data from S3 locations
- Provision a Virtual Private Cloud in contingency zone.
- Provision subnet and gateway in the VPC.

- Provision Elastic Load Balancers in VPC.
- Provision a Relational Database (RDS) in the specified region
- Restore backups from S3 to the RDS
- Provision EC2 instances.
- Deploy code to newly provisioned servers and build applications.
- Test newly deployed operations
- Coordinate with government DNS personnel on name change requests.
- Validate new DNS entries
- Notify team members and stakeholders that contingency environment is operational.

1.8.1.5. Return to Normal Service

When service in the original region or availability zone has been restored, system operations at the alternate site must be transitioned back. The goal is to provide a seamless transition of operations back to the primary operating environment.

First, after receiving notice from AWS that the issues have been resolved, GTRI will perform operational tests in the environment to validate proper functionality. Once GTRI and AWS have confirmed that the original environment is fit for service, GTRI will notify the customer that the application will be transferred to the original facility.

The exact methods used to restore an application to service will depend on the disaster recovery strategy selected for the application; details on these strategies can be found in Section 1.8. Procedures and runbooks to restore service to each application will be documented during the design and implementation phase of the project, and we will test the contingency plan on a regular basis.

While every application will differ in some details, we will take steps such as the following to restore service to an application:

- Take snapshots and backups of the recovery environment, as appropriate.
- Restore backups to hosts and databases in the original environment.
- Test system and application functionality on newly restored site.
- Create DNS requests for the newly restored systems.
- Once the restored system functionality is restored and DNS requests resolve to the new site, shut down (turn off) the contingency servers after 48 hours of operation.

1.8.2. Methodologies for Backup and Restore Services (8.8.2)

Maintaining backups of an application environment is incredibly important. GTRI believes in automating the backup and restore process as early as possible during the implementation phase of a project. Moreover, GTRI understands that it is important to test the backup and restore processes regularly to ensure that all data is backed up properly. As part of our operational processes, we perform regular contingency plan tests with our customers, where we will restore applications from backups and verify that they are functioning properly.

The exact method of configuring automation to back up and restore applications will differ on an application by application basis, but here are some of the tasks that we will automate and perform:

- Periodic Elastic Block Store snapshots

- RDS hourly/daily/weekly snapshots.
- We will work with the customer to define the desired retention policy for data backups, and we will build an automated snapshot rotation and lifecycle.
- For applications using Amazon's Elastic File System (EFS) we will Setup EFS-to-EFS backup.
- For NoSQL databases, we will set up DynamoDB backups and point-in-time replication
- GTRI will make sure all backups are secured comply with the customer's retention and data lifecycle policies.

All backups are stored in multiple AWS data centers. 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.

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, AWS currently has 18 regions, 54 Availability Zones, and 1 Local Region throughout the world: US East (Northern Virginia), US East (Ohio), US West (Oregon), US West (Northern California), AWS GovCloud (US-West), Canada (Central), EU (Ireland), EU (Frankfurt), EU (London), EU (Paris), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Osaka-Local), South America (Sao Paulo), China (Beijing), and China (Ningxia). Information on each region can be found at the AWS Global Infrastructure webpage.

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.

Amazon S3 runs on the world's largest global cloud infrastructure and is designed from the ground up to deliver 99.99999999% of durability. Data in Amazon S3 Standard, S3 Standard-IA, and Amazon Glacier storage classes is automatically distributed across a minimum of three physical Availability Zones (AZs) that are typically miles apart within an AWS Region. The Amazon S3 One Zone-IA storage class stores data in a single AZ and is ideal for customers who want a lower cost option for infrequently accessed data and do not require the availability and resilience of S3 Standard storage. Amazon S3 can also automatically replicate data to any other AWS Region.

Amazon S3 offers a highly durable, scalable, and secure destination for backing up and archiving your critical data. You can use S3's versioning capability to provide even further protection for your stored data. You can also define lifecycle rules to automatically migrate less frequently accessed data to S3 Standard - Infrequent Access and archive sets of objects to Amazon Glacier.

Amazon S3 and Amazon Glacier provide a range of storage classes to meet the needs of compliance archives for regulated industries or active archives for organizations who need fast, infrequent access to archive data. Amazon Glacier Vault Lock provides write-once-read-many (WORM) storage to meet compliance requirements for records retention. Lifecycle policies make transitioning data from Amazon S3 to Amazon Glacier simple, helping automate the transition based on

Amazon S3's highly durable, secure, global infrastructure offers a robust disaster recovery solution designed to provide superior data protection. Cross-Region Replication (CRR) automatically replicates every S3 object to a destination bucket located in a different AWS Region.

1.9. Data Protection (E) (8.9)

1.9.1. Standard Encryption Technologies and Options to Protect Sensitive Data (8.9.1)

Securing Data at Rest

There are several options for encrypting data at rest, ranging from completely automated AWS encryption solutions (such as AWS Key Management Service [KMS]) to manual, client-side options (such as AWS CloudHSM). 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.

Additionally, the *Encrypting Data at Rest* whitepaper 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 Sockets Layer/Transport Layer Security (SSL/TLS) for protection of data in transit. IPSec is a protocol that extends the IP protocol stack, often in network infrastructure, and allows applications on upper layers to communicate securely without modification. SSL/TLS, on the other hand, operates at the session layer, and while there are third-party SSL/TLS wrappers, it often requires support at the application layer as well.

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

1.9.2. Willingness to Sign Relevant and Applicable Business Associate Agreements (8.9.2)

GTRI has a standard Business Associate Addendum (BAA) we present to customers for signature. It takes into account the unique services AWS provides and accommodates the AWS Shared Responsibility Model.

1.9.3. Only Use Data for Purpose Defined in Master Agreement, Addendum or Related Service Level Agreement (8.9.3)

GTRI will not use customer data for any purpose other than providing services selected by each customer or legally required. AWS does not access or use customer content for any purpose other than as legally required and to provide the AWS services selected by each customer, to that customer and its end users. AWS never uses customer content or derives information from it for other purposes such as marketing or advertising.

1.10. Service Level Agreements (E) (8.10)

1.10.1. Whether Sample Service Level Agreement is Negotiable (8.10.1)

GTRI understands that the level of criticality for every application is different. Some applications can be offline for several hours with minimal impact. Others will result in serious financial impact or ability to function if they stop working, even for a moment. While the SLAs for the underlying AWS services are not negotiable, it is possible to utilize multiple AWS regions to design and build applications with higher

levels of availability and durability than those available within a single region. Companies such as Netflix have built applications with very high availability on AWS using highly-available multi-region architectures. It is important to understand that designing for redundancy across regions does increase the operational cost of an application.

For our managed services, we provide a few distinct SLA categories, and we have designed operational procedures to support them. We have found this to work well in practice and it has allowed us to provide high levels of service at very competitive price points because we can standardize and automate our workflows and procedures.

While it is certainly possible to negotiate a custom SLA for managed services, we would recommend against it because it will certainly cost the customer more because it would increase the level of customization in the environment, and possibly the staffing level necessary to provide the necessary monitoring and operational support.

For AWS, AWS has millions of active Customers and AWS offers the same portfolio of self-service, highly automated web services to its Customers on a one-to-many basis. Because of this, AWS cannot commit to keep the Services or SLAs the same for certain customers but improve or change them for others.

1.10.2. Sample Service Level Agreement (8.10.2)

Amazon provides SLAs for their foundational services. AWS frequently adds new SLA commitments or strengthens existing ones as their services mature, so it is worth checking the AWS Web site for the most recent details. As of this writing, here are a summary of the SLAs that AWS offers along with the service credits if the expected service levels are not met.

AWS Service	Service Type	Service Commitment	SLA Credit	Percentage
RDS	Relational Database	99.95%	Less than 99.95% but equal to or greater than 99.0%	10 %
RDS	Relational Database	99.95%	Less than 99.0%	25 %
S3	Storage	99.9%	Equal to or greater than 99.0% but less than 99.9%	10 %
S3	Storage	99.9%	Less than 99.0%	25 %
EC2	Compute	99.95%	Less than 99.95% but equal to or greater than 99.0%	10 %
EC2	Compute	99.95%	Less than 99.0%	30 %
DynamoDB (Global)	NoSQL Database	99.999%	Less than 99.999% but equal to or greater than 99%	10%
DynamoDB (Global)	NoSQL Database	99.999%	Less than 99.0%	30%

AWS Service	Service Type	Service Commitment	SLA Credit	Percentage
DynamoDB (Local)	NoSQL Database	99.99%	Less than 99.99% but equal to or greater than 99%	10%
DynamoDB (Local)	NoSQL Database	99.99%	Less than 99%	10%
Route 53	DNS + Global Load Balancing	100 %	5 – 30 minutes in a Billing Cycle	1-day Service Credit
Route 53	DNS + Global Load Balancing	100 %	31 minutes – 4 hours in a Billing Cycle	7 days Service Credit
Route 53	DNS + Global Load Balancing	100 %	More than 4 hours in a Billing Cycle	30 days Service Credit
CloudFront	Content Distribution Network	99.9%	Equal to or greater than 99% but less than 99.9%	10%
CloudFront	Content Distribution Network	99.9%	Less than 99%	30%
AWS Shield Advanced	DDos Mitigation	100%	Unavailable during a 24-hour period	1-day service credit

Table 7 –Foundational Services Sample SLAs

1.11. Data Disposal (E) (8.11)

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

GTRI often recommends that our AWS cloud customers use AWS Key Management Service (KMS). AWS Key Management Service (KMS) is a managed symmetric key service that allows customers to retain control of your regional Customer Master Key (CMK). The AWS S3 SSE-KMS encryption of objects leverages KMS and your CMK. Every customer gets a master key in the KMS and they can use this master key to create sub-keys on the keychain (hierarchy). It is then easy to perform a one-click encryption of server and database storage (RDS). Customers have centralized key management (create, delete, view, set policies) to enforced, automatic key rotation. AWS KMS has visibility into any changes via CloudTrail.

When it comes to destroying the encrypted data, you simply delete the encryption key that was used to encrypt it and the data is essentially “crypto-shredded”.

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.

1.12. Performance Measures and Reporting (E) (8.12)

1.12.1. Ability to Guarantee Reliability and Uptime Greater than 99.5% (8.12.1)

GTRI can guarantee high application availability through its use of AWS services and through proper design and implementation of application environments. The foundational AWS storage, compute and database services are very well designed, and have seen years of battle testing. These form the building blocks of highly available applications. Also, every US-based AWS region consists of multiple geographically separate datacenters, so it is remarkably easy to deploy applications in multiple physical datacenters with AWS.

GTRI can also leverage the immense scale by designing applications so that they can operate in distinct geographical regions, within the US or worldwide.

The following table illustrates the service levels that GTRI can realistically provide leveraging high-availability architecture and AWS services for an application with servers, object storage and a relational database:

AWS Service	Single Region SLA	Multi-Region SLA
EC2	99.95%	99.99998%
S3	99.90%	99.99990%
RDS	99.95%	99.99998%
Combined SLA	99.80%	99.99960%

Table 8 –High Availability SLA’s

The combined SLA is lower than the individual services because the failure of any one of the services can bring down the entire application. The multi-region SLA is much higher because multiple regions would need to fail in order for the application to become out of service.

The AWS Design Philosophy

AWS takes extensive precautions to help ensure that we will remain fully operational, with no loss of service for our hosted applications. AWS replicates critical system components across multiple Availability Zones to ensure high availability both under normal circumstances and during disasters such as fires, tornadoes, or floods. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity and housed in separate facilities. Each AWS Availability Zone runs on its own independent infrastructure, engineered to be highly reliable so that even extreme disasters or weather events should only affect a single Availability Zone. The data centers’ electrical power systems are designed to be fully redundant and maintainable without impact to operations. Common points of failure, such as generators, UPS units, and air conditioning, are not shared across Availability Zones.

At AWS, we plan for failure by maintaining contingency plans and regularly rehearsing our responses. In the words of Werner Vogels, Amazon's CTO: "Everything fails, all the time." We regularly perform preventative maintenance on our generators and UPS units to ensure that equipment is ready when needed. We also maintain a series of incident response plans covering both common and uncommon events and update them regularly to incorporate lessons learned and prepare for emerging threats.

How reliable is an application hosted by AWS?

In 2014, Nucleus Research surveyed 198 AWS customers that reported moving existing workloads from on-premises to AWS and found that they were able to reduce unplanned downtime by 32%.

While AWS goes to great lengths to provide availability of the cloud, our customers share responsibility for ensuring availability within the cloud. These customers and others like them have succeeded because they designed for failure and have adopted best practices for high availability, such as taking advantage of multiple Availability Zones and configuring Auto Scaling groups to replace unhealthy instances. The *Building Fault-Tolerant Applications on AWS* whitepaper is a great introduction to achieving high availability in the cloud. In addition, the AWS Well-Architected Framework codifies the experiences of thousands of customers, helping customers assess and improve their cloud-based architectures and mitigate disruptions.

In addition, the AWS Architecture Center is designed to provide customers with the necessary guidance and application architecture best practices to build highly scalable and reliable applications in the AWS Cloud. These resources will help you understand the AWS Cloud, its services and features, and will provide architectural guidance for design and implementation of systems that run on the AWS infrastructure.

1.12.2. Standard Uptime Service and Related SLA Criteria (8.12.2)

AWS far exceeds the Uptime Institute Tiering certification. Tiering aspects do not take into consideration the nature of the services of the cloud environment, and although the uptime institution tiering can be a great guide, it ultimately does not accurately map to a cloud service provider organization. AWS does not have a Certified Uptime Tiering level; however, we operate a data center environment using N+1 architecture. AWS offers SLAs for services such as Amazon EC2 and Amazon EBS at 99.95%, and Amazon S3 with an SLA of 99.9%. Our generator backup capabilities are detailed in our System and Organization Control (SOC) Reports (as is our discussion regarding business continuity planning and N+1 architecture).

1.12.3. Process Used for Customer to Call/Contact for Support (8.12.3)

GTRI detailed our process for a customer to contact for support in Section 1.4.1 above.

1.12.4. Consequences/SLA Remedies if Fail to Meet Incident Response and Incident Fix time (8.12.4)

GTRI stands behind its service commitments and those of its partners. AWS provides service when it does not meet its service commitments. Those credits are defined in Table 7 –Foundational Services Sample SLAs above.

For managed services, our service commitments are defined in Section 1.4.1 above. In the event that GTRI does not meet its service commitment, GTRI will offer the customer a credit of 2% of the total monthly service fees per event in the event billing cycle.

1.12.5. Procedures and Schedules for Planned Downtime (8.12.5)

In general, GTRI reserves a nightly window at 0900 UTC to apply system updates and reboot systems in the event of a security vulnerability. For systems that must be highly-available, redundant servers will be rebooted in sequence to ensure uptime. For other updates and maintenance, especially ones that may

require assistance from customer application developers, we will work with the customer to schedule a maintenance window.

AWS does not require systems to be brought offline to perform regular maintenance and system patching, and AWS's own maintenance and system patching generally do not impact customers. There may be occasions when AWS might schedule a customer instance for a reboot for necessary maintenance, such as to apply updates that require a reboot. No action is required on the customer's part; we recommend that customers wait for the reboot to occur within its scheduled window. These scheduled events are not frequent and if a customer instance will be affected by a scheduled event, they will receive an email prior to the scheduled event with details about the event, as well as a start and end date. Customers can also view scheduled events for their instance(s) by using the Amazon EC2 Console, API, or CLI. AWS will communicate with customers, either via email, or through the AWS Service Health Dashboard if service use is likely to be adversely affected.

1.12.6. Consequences/SLA Remedies if Disaster Recovery Metrics are Not Met (8.12.6)

If GTRI does not meet the Recovery Time Objective, meaning that the elapsed time between outage notification and service restoration was greater than the RTO defined for the application, GTRI will credit 10% of the monthly managed service fees for the month that the incident started.

If GTRI does not meet the Recovery Point Objective, meaning that data loss was greater than the RPO defined for the application, GTRI will credit 100% of the monthly managed service fees for the month that the incident started.

GTRI has provided our RTO and RPO objectives in Section 1.8 above.

1.12.7. Sample Performance Reports (8.12.07)

AWS customers can leverage AWS Cloud monitoring tools such as Amazon CloudWatch, AWS Trusted Advisor, AWS Health Checks, and third-party monitoring tools to extract metrics and system analytics.

The AWS Service Health Dashboard provides current and historical data across regions for each service offered. The status can be monitored in real time or subscribed to as an RSS feed by service.

In addition, GTRI recommends to customers to use CloudCheckr for more granular reports.

CloudCheckr's Utilization reporting helps you keep an eye on your AWS deployment, showing how much your resources are being utilized to help ensure your infrastructure is running smoothly. In the Utilization reports you'll find several different report types across multiple AWS services. The utilization statistics, which are taken from CloudWatch, are saved for the lifetime of your CloudCheckr account. This allows you to view historic utilization trends and makes future planning much easier.

CloudCheckr's Utilization reporting offers the following report types:

Summary: reports take a look across all resources within the selected service and provide important insight into how much those resources are used. The summary reports, which can be broken into different

utilization categories (such as CPU and Network), provide key statistics such as the average usage over the past 30 days, usage by day of the week, resource type, and hour of day.




Figure 6 - CloudCheckr Summary Performance Report

Details: For a full utilization breakdown for each instance CloudCheckr offers Utilization Details reporting. These reports will give you a full list of each of your resources running within the service and offers comprehensive utilization statistics and graphs for reach. Amongst the information these reports provide is an average usage over the past 7, 30, and 90 days, peak and lowest utilization, top 10 hours of highest utilization, and usage by hour of the day, and day of the week.




Figure 7 - CloudCheckr Detailed Utilization Report

Heat Maps: provide you with a visual breakdown how heavily your resources were utilized each hour over a 7-day period. You can refresh the report using absolute values (0-100) or relative. There are also several filtering options to look at the data by statistic type (bytes in versus out, for example), resource, or resource tag.




Figure 8 - CloudCheckr Heat Map Report

Historical Export: The CloudWatch Historical Export allows you to view and export CloudWatch data to CSV for any user-defined time frame. With this report you can export a complete, hourly list of your CloudWatch metrics into a CSV file. This historic view allows you to spot trends and predict future usage patterns. You can pull the metrics for all resources that ran during the entered time period, or you can filter the data to a specific resource.




Figure 9 - Export of historical performance data

1.12.8. Ability to Print Historical, Statistical, and Usage Reports Locally (8.12.8)

By using the CloudCheckr platform included with the AWS IaaS and PaaS offerings as described directly above, GTRI can provide historical, statistical and usage reports. Each Purchasing Agency can obtain their own access to their systems and view the reports locally without GTRI support, if needed

1.12.9. On-Demand Deployment Services Supported 24x365 (8.12.9)

In AWS, on-demand deployments can happen on a self-service basis either through a web-based control panel, via a command-line interface, or via an API, without any human intervention. Deploying infrastructure on-demand is available 24x7x365.

1.12.10. Scale-up and Scale-down Availability 24x365 (8.12.10)

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, customers 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%; and similarly, customers 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. Customers can configure a cool down 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 customers 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. Customers 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. 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 database instances, as well as custom metrics generated by applications and services and any log files your 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 and keep application 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 10.




Figure 10 – Auto Scaling with Elastic Load Balancing and Amazon CloudWatch alarms

1.13. Cloud Security Alliance (E) (8.13)

GTRI employs cloud consultants who possess the Cloud Security Alliance (CSA) Certificate of Cloud Security Knowledge (CCSK) and the (ISC)2 Certified Cloud Security Professional (CCSP) certifications. As you know, the CSA Security, Trust & Assurance Registry (STAR) program is based on the Cloud Controls Matrix (CCM) that the Cloud Security Alliance (CSA) publishes. This Cloud Controls Matrix (CCM) is a focus point of the CCSK and the CCSP certification. Therefore, GTRI's cloud security consultants have in-depth knowledge of these security controls.

CSA STAR Level 1: CSA STAR Self-Assessment

AWS has completed the CSA STAR Self-Assessment and published the results to the AWS website. Please refer to the CSA Consensus Assessments Initiative Questionnaire. This is the latest CAIQ (v3) released by the CSA.

Please see Attachment B with the completed CSA Self- Assessment.

CSA STAR Level 2: CSA STAR Attestation and Certification

Per the CSA definitions, AWS aligns with the CSA STAR Attestation and Certification via the determinations in our third-party audits for SOC and ISO:

CSA STAR Level 2 Attestation is based on SOC2, which can be requested with AWS Artifact - 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.

CSA STAR Level 3: Continuous Monitoring

As noted on the CSA website, CSA is still defining the Level 3 Continuous Monitoring requirements. Although, for this reason, AWS cannot determine alignment, AWS does provide customers with the tools they need to meet continuous monitoring requirements. Customers can leverage the AWS Security by Design (SbD) program by providing control responsibilities outlines, the automation of security baselines, the configuration of security and the customer audit of controls for AWS customer infrastructure, operating systems, services and applications running in AWS. This standardized, automated, prescriptive and repeatable design can be deployed for common use cases, security standards and audit requirements across multiple industries and workloads. For more information visit the Security by Design page.

CSA STAR Level 2 Certification is based on [ISO 27001:2005](#). Following is the link to the AWS certification examination in conformity with defined requirements in [ISO/IEC 17021-1:2015](#) and [ISO/IEC 27006:2015](#), the Information Security Management System as defined and implemented by AWS.

1.14. Service Provisioning (E)

1.14.1. Processing Emergency or Rush Service Implementations (8.14.1)

We work with organization that perform search and rescue and disaster response, so we understand the needs of customers who have urgent needs. If we have a contract with the requesting Purchasing Entity, as soon as we receive a purchase order, we will initiate an emergency implementation. If necessary, we will pull resources from less urgent projects. We expect that the customer will provide a dedicated point of contact will be available throughout the rush implementation. The customer must also understand that often "rush jobs" require remediation afterwards, and we expect that the customer will also provide resources and funding to support remediation after the emergency ends.

The customer does not need to worry about adequate compute and storage being available since GTRI works with "hyper-scale" cloud providers such as AWS, who provide a virtually limitless amount of compute and storage resources.

1.14.2. Standard Lead-time for Provisioning Solutions (8.14.2)

The time to provision a solution depends on the type of solution being provisioned and the SLA level required from the customer. It takes a longer time to provision a redundant multi-region highly available environment than it does to create a developer sandbox environment. Our cloud service providers operate at sufficient scale that they can provide raw resources within seconds or minutes; however, it takes more time to construct the environment according to the customers specifications.

The table below describes our standard process that we follow to provision a new environment along with customer responsibilities and our provisioning SLAs:

Onboarding Services	Description	Customer Responsibility	SLA
Billing & Invoice Setup	Set up appropriate billing codes and invoicing for the customer	Provide billing point of contact	24 hours after authorization
Configure Help Desk Portal	Ensure that customer can submit helpdesk tickets	n/a	24 hours after authorization
AWS Account Provisioning ¹	- Create or migrate AWS account. - Provide access to billing and utilization dashboard	- Provide administrative point of contact. - Provide alias that customer wishes to use for the account	2 days after authorization
AWS Services Configuration	- Configure audit logging, - Define password policies - Create default user and group policies	Provide any organization-specific roles, password policies, etc.	3 days after authorization
User Configuration	Create IAM users for customer AWS account access.	Provide a list of users and desired level of access.	3 days after authorization
Network Configuration	- Create Software Defined Networking (SDN) configuration via AWS Virtual Private Cloud (VPC) networks. - Define Subnets - Create routing tables - Create network Access Control Lists (NACLs)	CIDR block(s) to use for VPC networks (Only if customer wishes to connect internal network)	5 Business days after authorization

¹ **Note on Existing AWS Accounts:** If GTRI is taking over an existing AWS account, account ownership needs to be transferred to GTRI, a process that requires the government, the existing account owner and Amazon to execute agreements to transfer the account. The amount of time this process takes depends on how quickly the customer and the existing account owner execute the required agreements. GTRI can only offer a *best-efforts* SLA in this case.

Onboarding Services	Description	Customer Responsibility	SLA
Storage Configuration	- Create object storage buckets - Ensure that authorized users can read / write to the buckets.	Users authorized to read and/or write data to S3 buckets.	5 Business days after authorization
RDS Database Provisioning	- Create RDS database instance(s) - Provision database users - Configure networking and backups	- Applications and users who need access to the database itself. - For SQL Server or Oracle database: provide licensing information.	5 Business days after authorization
Custom Database Provisioning	Create database instance(s)	- Database software installation - Configuration and tuning of the database instance.	5 Business days after authorization
Backups	Setup and Configure Backups of environment stack and retain per customer policy.	Provide custom retention schedule if required more often or stored longer than customer policy.	5 Business days after authorization
EC2 Provisioning	- Launch EC2 virtual server instances. - Configure security group (firewall) settings. - Install available patches and updates. - Harden operating system - Set up users and groups	List of users who require administrative, power user, or regular user access to servers.	5 Business days after authorization

Table 9 –Standard Setup SLA's

1.15. Back Up and Disaster Plan (E)

1.15.1. Ability to Apply Legal Retention Periods and Disposition By Agency (8.15.1)

GTRI process: GTRI follows agencies retention policies for all data hosted in Infrastructure we provide. If a customer requires a legal retention longer than normal retention policies, we can change the storage times in our systems. Additionally, if needed we can archive data in Glacier for long term storage.

AWS customers control the entire life-cycle of their content on AWS and manage their content in accordance with their own specific needs, including content classification, access control, retention and deletion.

1.15.2. Known Inherent Disaster Recovery Risks and Potential Mitigation Strategies (8.15.2)

Cloud service providers provide a high level of redundancy for their infrastructure and a rich set of tools to implement disaster recovery strategies. The biggest risk to a speedy recovery is that an application's architecture does not support it. If an application is not 'cloud native', requires manual processes to recover, stores large amounts of state locally, or is otherwise difficult to configure or deploy in a highly-available manner, it will be difficult to restore in any environment.

Regardless of the application architecture, the cloud provides a variety of strategies to mitigate disaster recovery risks. When we bring a customer onboard, we will discuss disaster recovery strategies as part of our initial planning process.

Our infrastructure has a high level of availability and we provide you with the features you need to deploy a resilient IT architecture. Our systems are designed to tolerate system or hardware failures with minimal customer impact. The AWS Cloud supports many popular disaster recovery architectures, ranging from "pilot light" environments that are ready to scale up at a moment's notice to "hot standby" environments that enable rapid failover. All data centers are online and serving customers; no data center is "cold." In the case of a failure, automated processes move your data traffic away from the affected area. By distributing applications across multiple AWS Availability Zones, you can remain resilient in the face of most failure modes, including natural disasters or system failures.

You can build highly resilient systems in the cloud by employing multiple instances in multiple AWS Availability Zones and using data replication to achieve extremely high recovery time and recovery point objectives. We can help managing and testing the backup and recovery of your information system that is built on the AWS infrastructure. You can use the AWS infrastructure to enable faster disaster recovery of your critical IT systems without incurring the infrastructure expense of a second physical site.

1.15.3. Infrastructure Supports Multiple Data Centers in US for Failover Capability (8.15.3)

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 customers the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible with a single data center.

AWS currently has 18 regions, 55 Availability Zones, and 1 Local Region throughout the world: US East (Northern Virginia), US East (Ohio), US West (Oregon), US West (Northern California), AWS GovCloud (US-West), Canada (Central), EU (Ireland), EU (Frankfurt), EU (London), EU (Paris), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Osaka-Local), South America (Sao Paulo), China (Beijing), and China (Ningxia). Information on each region can be found at the AWS Global Infrastructure webpage. Figure 11 depicts the current AWS Regions and Availability Zones, along with the four new regions that AWS has announced plans for.

The AWS products and services that are available in each region are listed at the Region Table webpage.



Figure 11 – Global Map of AWS Regions and Availability Zones

Figure 12 illustrates the relationship between regions and Availability Zones.




Figure 12 – Regions and Availability Zones

1.16. Hosting and Provisioning (E) (8.16)

1.16.1. Documented Cloud Hosting Provisioning Process and Defined/Standard Cloud Provisioning Stack (8.16.1)

Our provisioning process along with SLAs and customer responsibilities is documented in Table 9 – Standard Setup SLA's above.

When we create AWS environments for customers, our stack is as follows:

AWS CloudFormation: CloudFormation automates the creation of multiple AWS resources in a repeatable and self-documenting manner using templates that describe the resources that should exist and their configuration

AWS Quick Starts for NIST-Based assurance frameworks: AWS provides a set of CloudFormation templates that are designed to comply with FedRAMP and the NIST 800-53 controls. They have been assessed by a third-party assessment organization (3PAO) and include a controls matrix documenting how the NIST 800-53 controls have been implemented and the customer's responsibilities.

Ansible: GTRI utilizes Ansible to provide a repeatable mechanism to configure and launch the NIST CloudFormation templates. We also use Ansible to launch, harden and configure individual hosts, storage buckets and other cloud resources. We also use Ansible to set up users, groups and permissions. In general, we use Ansible when possible to perform system administrative activities because we can run the playbooks on a regular basis to automatically remediate any configuration drift.

When necessary, we will utilize the AWS management console and command-line interface to perform tasks.

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. Customers 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 customers provision resources across multiple regions.

Command Line Interface

The AWS Command Line Interface (CLI) is a unified tool used to manage AWS Cloud services. With just one tool to download and configure, customers can control multiple AWS resources from the command line and automate them through scripts. The AWS CLI introduces a new set of simple file commands for efficient file transfers to and from Amazon Simple Storage Service (Amazon S3).

1.16.2. Tools Sets Used (8.16.2)

AWS offerings are provided with a range of supporting components like management tools, networking services, and application augmentation services, with multiple interfaces to AWS Application Programming Interface (API)-based services, including Software Development Kits (SDKs), Integrated Development Environment (IDE) toolkits, and Command Line Tools. Deploying new servers (determining configuration for both stand alone or part of an existing server farm, etc.)

1.16.2.1. DevOps Automation

Here are some of the services used as part of GTRI's DevOps automation platform:

DevOps Tools	Description
Github	Github is the leading provider of source code management tools, and GTRI uses Github to manage its source code.
Jenkins	Jenkins is an open source build server. We use Jenkins as part of a process to build and test changes to our source code.
Docker	GTRI likes to use container technology to ease application deployments by providing a standardized deployment environment and means of applying host configuration
RunDeck	RunDeck is an open-source service catalog that we use to deploy and provision services consistently across multiple accounts

Table 10 –DevOps Automation Tools

1.16.2.2. Creating and Storing Server Images for Future Multiple Deployments

Creating and storing server images for future multiple deployments is a very important use case for GTRI, so we use Ansible and Packer together to create a cloud-agnostic platform to build machine images:

Ansible – Ansible, an open-source configuration management platform provided by RedHat, is used to configure AWS infrastructure and individual machines. We use Ansible to install software and apply and reinforce hardening baselines.

Packer – Packer is an open-source command line tool designed to build machine images. We use Packer to ensure that customers always have fully-patched, hardened, up to date OS images to use for application deployments.

1.16.2.3. Securing Additional Storage Space

In the cloud, purchasing entities are charged for the storage that they consume. In AWS, storage is a quick to provision, elastic service. It is best to not provision a large amount of extra storage because it is always easy to get more storage on demand. The following AWS services are the basis of an elastic and cost-effective storage platform:

Elastic Block Store – AWS Elastic Block Store is used to dynamically create block storage volumes. There is a snapshot capability that makes it easy to perform consistent point-in-time snapshots. It is also possible to increase the size of storage volumes dynamically, without shutting down instances, so volumes can be expanded without any downtime.

Amazon S3 – S3 is an object store allowing users to store a virtually unlimited amount of data in storage buckets. S3 is ideal for large amounts of data that does not change frequently. Users are only charged for the amount of data they actually use.

AWS provides a wealth of items in their service catalog that can be used to deploy new services either standalone or as part of a server farm. GTRI can consult with a purchasing entity to help them decide what services will work the best in their environment.

AWS CodeCommit-AWS CodeCommit is a fully-managed source control service that makes it easy for companies to host secure and highly scalable private Git repositories. AWS CodeCommit eliminates the need to operate a source control system or worry about scaling its infrastructure. Customers can use AWS CodeCommit to securely store anything from source code to binaries, and it works seamlessly with existing Git tools.

AWS CodePipeline- AWS CodePipeline is a continuous integration and continuous delivery service for fast and reliable application and infrastructure updates. AWS CodePipeline builds, tests, and deploys a customer's code every time there is a code change, based on the release process models they define. This enables customers to rapidly and reliably deliver features and updates. Customers can easily build out an end-to-end solution by using our pre-built plugins for popular third-party services like GitHub or integrating custom plugins into any stage of the release process. With AWS CodePipeline, customers only pay for what they use. There are no upfront fees or long-term commitments.

AWS CodeBuild-AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. With AWS CodeBuild, customers don't need to provision, manage, and scale their own build servers. AWS CodeBuild scales continuously and processes multiple builds concurrently, so builds are not left waiting in a queue. Customers can get started quickly by using prepackaged build environments, or they can create custom build environments that use their own build tools. With AWS CodeBuild, customers are charged by the minute for the compute resources they use.

AWS CodeDeploy- AWS CodeDeploy is a service that automates code deployments to any instance, including Amazon EC2 instances and instances running on-premises. AWS CodeDeploy makes it easier

for customers to rapidly release new features, helps them avoid downtime during application deployment, and handles the complexity of updating their applications. Customers can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations. The service also scales with infrastructure, so customers can easily deploy to one instance or thousands.

Management Tools- AWS provides a broad set of services that help IT administrators, systems administrators, and developers more easily manage and monitor their resources. Using these fully managed services, customers can automatically provision, configure, and manage their AWS or on-premises resources at scale. Customers can also monitor infrastructure logs and metrics using real-time dashboards and alarms. AWS also helps customers monitor, track, and enforce compliance and security.

Resource Provisioning AWS CloudFormation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion. Customers can use AWS CloudFormation's sample templates or create their own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run their application.

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. AWS Marketplace features many software categories including databases, application servers, testing tools, monitoring tools, content management, and business intelligence. Visitors to AWS Marketplace can use 1-Click deployment to quickly launch pre-configured software and pay only for what they use, by the hour or month. AWS handles billing and payments, and software charges appear on customers' AWS bill.

AWS Service Catalog allows organizations to create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows customers to centrally manage commonly deployed IT services, helps them achieve consistent governance, and helps them meet their compliance requirements, all while enabling users to quickly deploy only the approved IT services they need.

Configuration Management- AWS OpsWorks is a configuration management service that helps customers configure and operate applications of all shapes and sizes using Chef. Customers can define the application's architecture and the specification of each component including package installation, software configuration, and resources such as storage. Customers can start from templates for common technologies like application servers and databases or build their own to perform any task that can be scripted. AWS OpsWorks includes automation to scale applications based on time or load and dynamic configuration to orchestrate changes as an environment scales.

AWS Systems Manager allows customers to centralize operational data from multiple AWS services and automate tasks across AWS resources. Customers can create logical groups of resources such as applications, different layers of an application stack, or production versus development environments. With AWS Systems Manager, customers can select a resource group and view its recent API activity, resource configuration changes, related notifications, operational alerts, software inventory, and patch compliance status. AWS Systems Manager provides a central place to view and manage AWS resources, so customers can have complete visibility and control over their operations.

1.16.2.4. 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).

GTRI has several monitoring tools that are available for Participating Entities. Descriptions of each are listed below.

Monitoring and Performance

Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications that customers run on AWS. Customers can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in their AWS resources. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon Relational Database Service (Amazon RDS) DB instances, as well as custom metrics generated by the 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 and then use those insights to react and keep their application running smoothly.

Governance and Compliance

AWS Config is a fully managed service that provides customers with an AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. AWS Config Rules enables customers to create rules that automatically check the configuration of AWS resources recorded by AWS Config. With AWS Config, customers can discover existing and deleted AWS resources, determine their overall compliance against rules, and dive into configuration details of a resource at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.

AWS CloudTrail is a web service that records AWS API calls for a customer's account and delivers log files to them. 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. With AWS CloudTrail, customers 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 Cloud services (such as AWS CloudFormation). The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

AWS Service Catalog allows organizations to create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows customers to centrally manage commonly deployed IT services, helps them achieve consistent governance, and helps them meet their compliance requirements, all while enabling users to quickly deploy only the approved IT services they need.

Resource Optimization

AWS Trusted Advisor is an online resource to help customers reduce cost, increase performance, and improve security by optimizing their AWS environment. AWS Trusted Advisor provides real-time guidance to help customers provision their resources following AWS best practices.

1.17. Trial and Testing Periods (Pre and Post Purchase) (E) (8.17)

Finding the best mode of operation in a cloud environment can require some experimentation. Many cloud service providers have a lot of breadth and depth to their service catalogs. The best way to determine what tools and services will work the best is by experimentation. This is one of the strengths of a cloud environment – it is possible to create a sandbox and try out different approaches for a minimal charge, sometimes free of charge.

1.17.1. Testing and Training Periods (8.17.1)

One of the ways that GTRI can accelerate a customer's initiatives is through our architecture workshops. We will deliver customized training and meet with a customer's stakeholders to determine requirements

and needs. We will provide the customer with a clean sandbox environment that they can use for whatever purpose they see fit. If the customer has a specific application they wish to deploy in a cloud environment, we can also propose an architecture based on our experience delivering cloud solutions for our customers. We describe this offering in more detail in Section 1.18.2.

The testing and training periods can be as long or as short as the customer wishes, there is no set limitation. If the customer no longer wishes to pay for the resources that they have created in their cloud sandbox, we can shut down the environment.

1.17.2. Providing Test and/or Proof of Concept Environment for Evaluation (8.17.2)

GTRI can provide test and/or proof of concept environments in a very short time-frame. AWS accounts can be created free of charge, and there is no limit on the number of accounts that a customer can have. We can create a test or sandbox account very quickly once we have received authorization from a customer. Here are the steps necessary to create a test or PoC environment:

Onboarding Services	Description	Customer Responsibility	SLA
Billing & Invoice Setup	Set up appropriate billing codes and invoicing for the customer	Provide billing point of contact	24 hours after authorization
Configure Help Desk Portal	Ensure that customer can submit helpdesk tickets	n/a	24 hours after authorization
AWS Account Provisioning	- Create or migrate AWS account. - Provide access to billing and utilization dashboard	Provide administrative point of contact. Provide alias that customer wishes to use for the account	2 days after authorization
AWS Services Configuration	- Configure audit logging, - Define password policies - Create default user and group policies	Provide any organization-specific roles, password policies, etc.	3 days after authorization
User Configuration	Create IAM users for customer AWS account access.	Provide a list of users and desired level of access.	3 days after authorization

Table 11 –Proof of Concept SLA's

At the end of this process, the customer will have a clean, but securely-configured AWS account that they can use for testing or PoC purposes.

If a customer simply wants to try out a service, AWS provides a number of hands-on labs on their AWS training site free of charge, or through QuikLabs, a training partner. Many of these labs and trainings are completely free of charge and do not require an AWS account, or even an interaction or purchasing relationship with GTRI.

1.17.3. Training and support Provided at No Additional Cost (8.17.3)

GTRI Bootcamps

Almost all AWS services have a free tier that allows customers to try services, or even deploy small-scale applications free of charge. A customer can provision servers, databases, networking, users, notification services, message queues and more at a small scale without needing to pay. Information about the AWS free tier of services is available on the AWS Web Site under the Pricing Page for each service. Amazon also provides a Simple Pricing Calculator that will allow a customer to confirm that the environment they are about to create falls into the free tier of services.

If a customer simply wants to try out a service, AWS provides a number of hands-on labs on their AWS training site free of charge, or through QuikLabs, a training partner. Many of these labs and trainings are completely free of charge and do not require an AWS account, or even an interaction or purchasing relationship with GTRI.

GTRI also provides free labs and bootcamps to the general public on a variety of topics, most notably around Splunk's operational and enterprise security features. Anyone can sign up for these bootcamps on the GTRI Web site, without needing to have a purchasing relationship with GTRI.

1.18. Integration and Customization (E) (8.18)

1.18.1. Solutions Can Be Integrated into Complementary Applications (8.18.1)

AWS provides documented HTTP-based REST APIs to support integration of its entire service catalog with complementary applications. To further support integration with complementary applications, AWS also provides Software Development Kits (SDKs) to integrate AWS with all major operating systems and application platforms.

Given AWS's status as the most popular cloud service provider, most complementary applications already support AWS. For example, all of the major DevOps tools, such as Ansible and Jenkins already integrate with AWS. Popular backup and recovery utilities such as Commvault, Veeam and Zerto support AWS out of the box with minimal configuration. IT monitoring tools, such as Splunk, provide free add-ons to integrate with AWS.

If a purchasing entity needs to integrate a complementary application with AWS that does not already have support for the platform, GTRI's professional services team can perform bespoke integrations. Please see Section 3.4 for more details.

1.18.2. Ways to Customize and Personalize Solutions for Purchasing Entities (8.18.2)

GTRI specializes in the configuration and operation of AWS cloud infrastructure and can provide a wide depth and breadth of services to meet any specific needs of a purchasing entity and supporting all phases of the application life cycle. We provide the following services to tailor our cloud offerings to the purchasing entity's needs:

- **Architecture Workshops** - If a purchasing entity is seeking guidance on how to get started in the cloud, we provide tailored architecture workshops that can cover the service catalog, cloud economics and cost optimization, and application architecture in the cloud.
- **Landing Zone Implementation** – Starting out with solid design principles, an understanding of how the environment will grow over time, and how to build in security from day one is vital to operating successfully in the cloud. GTRI can help build your cloud “landing zone”.

- **Professional Services** – GTRI has a deep bench of AWS certified engineers who can build out environments on behalf of purchasing entities or augment their existing development, operations and security teams.
- **Managed Services** – GTRI can help purchasing entities focus on their missions by handling the day-to-day business of keeping a cloud environment secure, performant and operational.

These services are described in more detail in section 3 below.

2. Marketing Plan (E) (8.19)

The GTRI Account Team will provide awareness training of our company and capabilities to all States who currently participate or are considering participating in this NASPO contract.

GTRI will update our web site with the respective NASPO contract information indicating our ability to provide Cloud Solutions through this purchasing vehicle as well as the depth and breadth of the offerings we provide

GTRI will provide contract awareness during the various conferences we attend in each state

GTRI has a dedicated Marketing Team who will assist with awareness through a combination of direct and indirect marketing to potential customers (states).

GTRI will utilize the respective marketing functions of our Partners such as AWS, Splunk, Facion and others to market to their existing and future customer base regarding this NASPO contract.

3. Related Value-Added Services to Cloud Solutions (E) (8.20)

3.1. Architecture Workshops

Many purchasing entities are new to operating applications in the cloud. It is possible to cut costs and operate efficiently in the cloud, but it requires a change in mindset, practices and culture from the traditional on-premises mode of operation. A successful transition to the cloud requires careful assessment and planning, and GTRI offers architecture workshops for purchasing entities at all stages of maturity.

GTRI can tailor workshops to provide the maximum benefit to the purchasing entity, including the following topics:

- **The AWS Service Catalog** – AWS has thousands of features and hundreds of services in its service catalog. We can go over the major services and discuss what they are, how they are priced and when to use them.
- **Cloud Economics** - It is important to know how AWS services are priced in order to operate efficiently in the cloud. We can talk about strategies to measure utilization and cut costs.
- **Application Architecture** – In AWS, it is easier than ever to build applications that span multiple datacenters, that are highly available, and can scale based on demand. We can discuss architecture best practices in the cloud, and strategies for migrating existing applications to the cloud.

At the completion of the architecture workshops, GTRI can provide the customer with architecture documentation, up-front and operational cost estimates for a workload.

3.2. Landing Zone Implementation

One of the most difficult parts of building and AWS environment is the initial environment configuration. GTRI incorporates lessons learned and best practices from setting up numerous AWS environments for customers. By starting with a clear strategy of to govern access into the cloud environment, how applications will be deployed, how networks will be properly segmented, traffic inspected, and access controlled in order to meet compliance and security needs, customers can ensure that they get started in the best way possible and avoid costly mistakes by designing their environment correctly.

GTRI provides consulting and implementation services to build out landing zones for customers.

3.3. Professional Services

GTRI knows that each enterprise market presents its own challenges. Specific industry needs differ, and so do requirements for data security, regulatory compliance, availability and performance. We have the flexibility—and expertise—to scale our solutions to your requirements, from physical infrastructure, to integrated solutions, to network monitoring and management.

GTRI's Professional Services consulting practice adds value beyond connectivity by combining our multidisciplinary expertise to deliver business applications encompassing advanced network engineering, network security, collaborative solutions and virtualization. The Professional Services team leverages its business and specialized technical expertise and employs a collaborative, structured, repeatable approach to deliver best of breed technology solutions. Our industry expert consultants listen and work closely with our clients to deliver solutions that effectively address your business objectives. By combining multiple practices under one business unit, our experts can collaborate across each layer and determine the best solution. Instead of isolated technology, our clients receive integrated solutions that meet business needs and provide expected ROI.



Figure 13 – GTRI's Service Offerings

3.3.1. Hybrid IT

Networking

Every data, voice and video applications are expected to run seamlessly and integrate with ease. You must serve users from anywhere with any device. Networks must also be flexible, scalable, agile and elastic so they can adapt to shifting business requirements. GTRI can provide a comprehensive network solution for your most important business and information technology (IT) challenges.

Network Infrastructure

- Data center and virtualized networking
- Load balancing and traffic optimization
- Local-area network (LAN), wide area network (WAN), and Internet engineering
- Wireless LAN and RF engineering
- Network management systems

Advanced Network Technologies

- Cloud networking
- Data center interconnect protocols
- Software Defined Networking (SDN)
- Internet Protocol Version 6 (IPv6) deployment
- Virtualized network topologies

Mobility

Now applications and systems are accessible from anywhere with the onset of mobile devices, tablets or desktop computers. Employees can maximize productivity and collaborate virtually. GTRI provides cost-effective and secure segmented wireless Internet access, while reducing wireless local-area network (WLAN) operating costs. Scalability, reliability and network management are all valuable and provided by GTRI.

Data Center

Today's data center architecture is dramatically changing the way information technology (IT) provides business agility. The modern data center includes technology to virtualize and pool resources. Plus, data centers now allow self-service capabilities, as well as orchestration and management tools. Data centers provide unmatched levels of uptime, responsiveness and flexibility. Improvements in capacity planning, provisioning, and policy control, visibility, application development and application maintenance help lower your costs and reduce your business risk. And now multi-tenant, shared infrastructure platforms support a variety of workloads, while still meeting unique security, storage and performance requirements from any user group, worldwide.

Maintaining older systems is becoming more and more expensive from both a knowledge and time perspective. Application inter-dependency makes it more challenging to consolidate and standardize platforms—upgrading one component breaks two others. Backup, recovery and disaster recovery procedures need to be rebuilt for virtual and hybrid environments and expanded to accommodate the Cloud. With virtualization, new challenges arise, leaving you trying to figure out who is using what. General lack of knowledge creates trust, security and compliance concerns around moving applications to the Cloud.

- Data Center Virtualization
- o Virtualization Architecture
- o Network Functions Virtualization
- Converged Infrastructure
- High Performance Computing

Cloud

The promise of the Cloud is real. Now there are methods for provisioning, measuring, and managing information technology (IT) infrastructure in a virtual environment. Not all Cloud applications are created equal however. GTRI provides careful planning and operational considerations to ensure sound business decisions are still at the core of your cloud strategy. We know the cloud is a business and operational model, not a just a technology or stand-alone strategy.

GTRI can assist you evaluate a variety of concerns when selecting cloud services and partners including, governance, controls, security, data ownership, privacy and regulatory requirements. We assist with

validating and monitoring reliability, ease-of-management, integration with other systems and tools, migrations and performance issues to name a few concerns. GTRI can help you with financial analysis and vendor choice—evaluating services, support, direction and financial strength for both hard and soft costs.

3.3.2. Collaboration

Smart collaboration tools help speed up and increase the quality of workforce communications. They foster innovation, strengthen personal connections and enable anytime, anywhere access to vital business information. The drive to offer these tools through personal devices and social media tools requires a unique understanding of the underlying communication platforms that frame the way we talk, meet, interact and analyze and share data.

GTRI understands there are critical collaboration and communication issues today. The need for real-time, anywhere access to co-workers and critical data is crucial. The typical method of emailing and phoning in, depersonalizes communication and creates chaos. What's more, those methods require constant maintenance and support. GTRI provides the ability to work virtually with employees and customers, as well as suppliers.

3.3.3. Security

In this era of hyper-connectivity, security threats are more prevalent than ever. The emergence and rapid evolution of Cloud consumption models and the explosion of mobile devices create both enormous opportunity and significant security and privacy threats. With everything from insider threats and the emergence of "hacktivism," to the proliferation of state-sponsored attacks, information security has risen to the top of information technology (IT) strategic planning. Security models today require comprehensive architectures, enterprise visibility, situational awareness, governance, automation and a skilled team to face modern security issues.

Beyond the typical security concerns, GTRI explains what technologies are available to help protect the highly virtualized data center. Furthermore, we explain the risks, security, policies and controls of moving to the Cloud with all complexities. Also, GTRI can assist with compliance issues, like Health Insurance Portability and Accountability Act (HIPAA) and Payment Card Industry Data Security Standard (PCI DSS).

3.3.4. Big Data Analytics

To enhance operational and security awareness, GTRI partners with and leverages Splunk. Splunk software is a robust security platform that collects, analyzes and reports on all types of machine generated data and turns it into valuable insights. It also provides visibility across disparate systems, enabling its users to make informed decisions across silos.

Splunk has numerous use cases including information technology (IT) operational awareness, security, application management, web analytics and business intelligence to name a few. It can index any machine generated data whether it's structured, semi-structured, or completely unstructured. Data is stored in its raw form with full fidelity without the need for a relational database. It applies schema on the fly at search to make it available for analysis and insights.

GTRI is the first Splunk Elite Partner in the U.S. and is one of just two authorized training centers in North America. We are the most credentialed Splunk partner, holding nearly 50 certifications across our sales, engineering, administration and architecture teams. As such, GTRI can provide end-to-end support for Splunk from licensing and pre-sales support to certified training, professional services and Splunk Managed Services.

Splunk Use Cases

- Application management

- IT operations
- Security
- Compliance
- Fraud
- Business intelligence

3.4. Managed Services

Many customers do not have dedicated system administration staff, database administrators, information security staff and many other IT Operational roles, yet they need to utilize applications in a cloud-hosted environment. GTRI can bridge the gap by providing a secure, performant, reliable AWS-based managed cloud platform. GTRI can handle all of the IT operations so the customers can focus on their mission.

The diagram below illustrates how GTRI and AWS can work together to handle most of the compliance, security and operational details:





Figure 14 – Roles in Security and Operations

Inside of a customer's cloud infrastructure, GTRI can handle many of the routine tasks that are necessary to ensure that systems are secure and operational:



Note: Patching includes Minor upgrades. Major version upgrades will require T&M Services

Figure 15 – GTRI's Responsibilities for Routine Tasks

3.5. Application Assessments

If a customer is just starting their journey towards cloud solutions and is looking to better understand their application portfolio, GTRI performs application assessments. This is a light-weight engagement designed to help customers better understand the applications they are running, how they communicate with each other, how easy it will be to move the application to the cloud, and how much it would cost to operate those applications in a cloud environment.

GTRI will use automated processes to discover the applications on a customer's network, the resources they utilize and how they communicate with each other. As a deliverable, GTRI will provide the customer a report detailing the applications that they are running, their dependencies and the resources they are using. We will also provide some recommendations on what applications would be the best fit for a cloud environment.

If a customer desires more information, or a full application portfolio analysis or cloud migration roadmap, GTRI's professional services organization can perform those services as a custom engagement.

3.6. GIS Application Services

GTRI is contracted to host GIS applications on behalf of the Federal Geographic Data Community as part of a task order awarded by the Department of the Interior. As such we are subject matter experts in hosting and operating GIS applications. Our capabilities include:

- Hosted GIS Server environments
- Extract-Transform-Load (ETL) of GIS data
- Virtual Desktop Environments for GIS analysts
- On-demand environments for disaster response

4. Supporting Infrastructure (E) (8.22)

GTRI's IaaS Solution and IaaS/PaaS solution with AWS do not require any additional infrastructure to deploy our solutions.

CSA Consensus Assessments Initiative Questionnaire

May 2017



© 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Introduction	1
CSA Consensus Assessments Initiative Questionnaire	1
Further Reading	56
Document Revisions	56

Abstract

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.

Introduction

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.” For more information, see <https://cloudsecurityalliance.org/about/>.

A wide range of industry security practitioners, corporations, and associations participate in this organization to achieve its mission.

CSA Consensus Assessments Initiative Questionnaire

Control Group	CID	Consensus Assessment Questions	AWS Response
Application & Interface Security <i>Application Security</i>	AIS-01.1	Do you use industry standards (Build Security in Maturity Model [BSIMM] benchmarks, Open Group ACS Trusted Technology Provider Framework, NIST, etc.) to build in security for your Systems/Software Development Lifecycle (SDLC)?	The AWS system development lifecycle incorporates industry best practices which include formal design reviews by the AWS Security Team, threat modeling and completion of a risk assessment. Refer to the AWS Overview of Security Processes for further details. AWS has procedures in place 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?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	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?	
Application & Interface Security <i>Customer Access Requirements</i>	AIS-02.1	Are all identified security, contractual and regulatory requirements for customer access contractually addressed and remediated prior to granting customers access to data, assets and information systems?	AWS Customers retain responsibility to ensure their usage of AWS is in compliance with applicable laws and regulations. AWS communicates its security and control environment to customers through industry certifications and third-party attestations, white papers (available at http://aws.amazon.com/compliance) and providing certifications, reports and other relevant documentation directly to AWS Customers.
	AIS- 02.2	Are all requirements and trust levels for customers' access defined and documented?	
Application & Interface Security <i>Data Integrity</i>	AIS-03.1	Are data input and output integrity routines (i.e., reconciliation and edit checks) implemented for application interfaces and databases to prevent manual or systematic processing errors or corruption of data?	AWS data integrity controls as described in AWS SOC reports illustrates the data integrity controls maintained through all phases including transmission, storage and processing. In addition, refer to ISO 27001 standard, Annex A, domain 14 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Application & Interface Security <i>Data Security / Integrity</i>	AIS-04.1	Is your Data Security Architecture designed using an industry standard (e.g., CDSA, MULITSAFE, CSA Trusted Cloud Architectural Standard, FedRAMP, CAESARS)?	AWS Data Security Architecture was designed to incorporate industry leading practices. Refer to AWS Certifications, reports and whitepapers for additional details on the various leading practices that AWS adheres to (available at http://aws.amazon.com/compliance).
Audit Assurance & Compliance <i>Audit Planning</i>	AAC-01.1	Do you produce audit assertions using a structured, industry accepted format (e.g., CloudAudit/A6 URI Ontology, CloudTrust, SCAP/CYBEX, GRC XML, ISACA's Cloud Computing Management Audit/Assurance Program, etc.)?	AWS obtains certain industry certifications and independent third-party attestations and provides certain certifications, reports and other relevant documentation directly to AWS Customers.
Audit Assurance & Compliance <i>Independent Audits</i>	AAC-02.1	Do you allow tenants to view your SOC2/ISO 27001 or similar third-party audit or certification reports?	AWS provides third-party attestations, certifications, Service Organization Controls (SOC) reports and other relevant compliance reports directly to our customers under NDA. The AWS ISO 27001 certification can be downloaded here . The AWS SOC 3 report can be downloaded here . 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 audits and risk assessments. AWS engages with external certifying bodies and independent
	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?	
	AAC-02.4	Do you conduct internal audits regularly as prescribed by industry best	

Control Group	CID	Consensus Assessment Questions	AWS Response
		practices and guidance?	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 <i>Information System Regulatory Mapping</i>	AAC-03.1	Do you have the ability to logically segment or encrypt customer data such that data may be produced for a single tenant only, without inadvertently accessing another tenant's data?	All data stored by AWS on behalf of customers has strong tenant isolation security and control capabilities. Customers retain control and ownership of their data, thus it is their responsibility to choose to encrypt the data. AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, SimpleDB and EC2. IPSec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security 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
	AAC-03.2	Do you have capability to recover data for a specific customer in the case of a failure or data loss?	

Control Group	CID	Consensus Assessment Questions	AWS Response
			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-03.3	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. For a complete list of available regions, see the AWS Global Infrastructure page.
	AAC-03.4	Do you have a program in place that includes the ability to monitor changes to the regulatory requirements in relevant jurisdictions, adjust your security program for changes to legal requirements, and ensure compliance with relevant regulatory requirements?	AWS monitors relevant legal and regulatory requirements. Refer to ISO 27001 standard Annex 18 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Business Continuity Management & Operational Resilience <i>Business Continuity Planning</i>	BCR-01.1	Do you provide tenants with geographically resilient hosting options?	Data centers are built in clusters in various global regions. AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Customers should architect their AWS usage to take advantage of multiple Regions and Availability Zones. Refer to AWS Overview of Cloud Security whitepaper for additional details - available at http://aws.amazon.com/security .
	BCR-01.2	Do you provide tenants with infrastructure service failover capability to other providers?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Business Continuity Testing</i>	BCR-02.1	Are business continuity plans subject to test at planned intervals or upon significant organizational or environmental changes to ensure continuing effectiveness?	AWS Business Continuity Policies and Plans have been developed and tested in alignment with ISO 27001 standards. Refer to ISO 27001 standard, annex A domain 17 for further details on AWS and business continuity.
Business Continuity Management & Operational Resilience <i>Power / Telecommunications</i>	BCR-03.1	Do you provide tenants with documentation showing the transport route of their data between your systems?	AWS Customers designate in which physical region their data and servers will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. AWS SOC reports provides additional details. Customers can also choose their network path to AWS facilities, including over dedicated, private networks where the customer controls the traffic routing.
	BCR-03.2	Can tenants define how their data is transported and through which legal jurisdictions?	
Business Continuity Management & Operational Resilience Documentation	BCR-04.1	Are information system documents (e.g., administrator and user guides, architecture diagrams, etc.) made available to authorized personnel to ensure configuration, installation and operation of the information system?	Information System Documentation is made available internally to AWS personnel through the use of Amazon's Intranet site. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security/ . Refer to ISO 27001 Appendix A Domain 12.
Business Continuity Management & Operational Resilience <i>Environmental Risks</i>	BCR-05.1	Is physical protection against damage (e.g., natural causes, natural disasters, deliberate attacks) anticipated and designed with countermeasures applied?	AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 11.

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Equipment Location</i>	BCR-06.1	Are any of your data centers located in places that have a high probability/occurrence of high-impact environmental risks (floods, tornadoes, earthquakes, hurricanes, etc.)?	AWS data centers incorporate physical protection against environmental risks. AWS' physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 11.
Business Continuity Management & Operational Resilience <i>Equipment Maintenance</i>	BCR-07.1	If using virtual infrastructure, does your cloud solution include independent hardware restore and recovery capabilities?	EBS Snapshot functionality allows customers to capture and restore virtual machine images at any time. Customers can export their AMIs and use them on premise or at another provider (subject to software licensing restrictions). Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	BCR-07.2	If using virtual infrastructure, do you provide tenants with a capability to restore a Virtual Machine to a previous state in time?	
	BCR-07.3	If using virtual infrastructure, do you allow virtual machine images to be downloaded and ported to a new cloud provider?	
	BCR-07.4	If using virtual infrastructure, are machine images made available to the customer in a way that would allow the customer to replicate those images in their own off-site storage location?	
	BCR-07.5	Does your cloud solution include software/provider independent restore and recovery capabilities?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Business Continuity Management & Operational Resilience <i>Equipment Power Failures</i>	BCR-08.1	Are security mechanisms and redundancies implemented to protect equipment from utility service outages (e.g., power failures, network disruptions, etc.)?	AWS equipment is protected from utility service outages in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS SOC reports provides additional details on controls in place to minimize the effect of a malfunction or physical disaster to the computer and data center facilities. In addition, refer to the AWS Cloud Security Whitepaper - available at http://aws.amazon.com/security .
Business Continuity Management & Operational Resilience <i>Impact Analysis</i>	BCR-09.1	Do you provide tenants with ongoing visibility and reporting of your operational Service Level Agreement (SLA) performance?	AWS CloudWatch provides monitoring for AWS cloud resources and the applications customers run on AWS. Refer to aws.amazon.com/cloudwatch for additional details. AWS also publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Refer to status.aws.amazon.com .
	BCR-09.2	Do you make standards-based information security metrics (CSA, CAMM, etc.) available to your tenants?	
	BCR-09.3	Do you provide customers with ongoing visibility and reporting of your SLA performance?	
Business Continuity Management & Operational Resilience <i>Policy</i>	BCR-10.1	Are policies and procedures established and made available for all personnel to adequately support services operations' roles?	Policies and Procedures have been established through AWS Security framework based upon NIST 800-53, ISO 27001, ISO 27017, ISO 27018, ISO 9001 standard and the PCI DSS requirements. Refer to AWS Risk and Compliance Whitepaper for additional details - available at http://aws.amazon.com/compliance .
Business Continuity Management	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

Control Group	CID	Consensus Assessment Questions	AWS Response
& Operational Resilience <i>Retention Policy</i>	BCR-11.2	Do you have a documented procedure for responding to requests for tenant data from governments or third parties?	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-11.4	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 <i>New Development / Acquisition</i>	CCC-01.1	Are policies and procedures established for management authorization for development or acquisition of new applications, systems, databases, infrastructure, services, operations and facilities?	Policies and Procedures have been established through AWS Security framework based upon NIST 800-53, ISO 27001, ISO 27017, ISO 27018, ISO 9001 standard and the PCI DSS requirements. Whether a customer is new to AWS or an advanced user, useful information about the services, ranging from introductions to advanced features, can be found on the AWS Documentation section of our website at https://aws.amazon.com/documentation/ .
	CCC-01.2	Is documentation available that describes the installation, configuration and use of products/services/features?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Change Control & Configuration Management <i>Outsourced Development</i>	CCC-02.1	Do you have controls in place to ensure that standards of quality are being met for all software development?	AWS does not generally outsource development of software. AWS incorporates standards of quality as part of the system development lifecycle (SDLC) processes. 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 <i>Quality Testing</i>	CCC-03.1	Do you provide your tenants with documentation that describes your quality assurance process?	AWS maintains an ISO 9001 certification. This is an independent validation of AWS quality system and determined that AWS activities comply with ISO 9001 requirements. AWS Security Bulletins notify customers of security and privacy events. Customers can subscribe to the AWS Security Bulletin RSS feed on our website. Refer to aws.amazon.com/security/security-bulletins/ . 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?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Change Control & Configuration Management <i>Unauthorized Software Installations</i>	CCC-04.1	Do you have controls in place to restrict and monitor the installation of unauthorized software onto your systems?	AWS' program, processes and procedures for managing malicious software is in alignment with ISO 27001 standards. Refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Change Control & Configuration Management <i>Production Changes</i>	CCC-05.1	Do you provide tenants with documentation that describes your production change management procedures and their roles / rights / responsibilities within it?	AWS SOC reports provides an overview of the controls in place to manage change management in the AWS environment. In addition, refer to ISO 27001 standard, Annex A, domain 12 for further details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Data Security & Information Lifecycle Management <i>Classification</i>	DSI-01.1	Do you provide a capability to identify virtual machines via policy tags/metadata (e.g., tags can be used to limit guest operating systems from booting/instantiating/tr ansporting 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

Control Group	CID	Consensus Assessment Questions	AWS Response
	DSI-01.5	Can you provide the physical location/geography of storage of a tenant's data in advance?	move customers' content from the selected Regions without notifying the customer, unless required to comply with the law or requests of governmental entities. For a complete list of available regions, see the AWS Global Infrastructure page.
	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. For a complete list of available regions, see the AWS Global Infrastructure page.
Data Security & Information Lifecycle Management <i>Data Inventory / Flows</i>	DSI-02.1	Do you inventory, document, and maintain data flows for data that is resident (permanent or temporary) within the services' applications and infrastructure network and systems?	AWS Customers designate in which physical region their content will be located. AWS will not move customers' content from the selected regions without notifying the customer, unless required to comply with the law or requests of governmental entities. For a complete list of available regions, see the AWS Global Infrastructure page.
	DSI-02.2	Can you ensure that data does not migrate beyond a defined geographical residency?	
Data Security & Information Lifecycle Management <i>eCommerce Transactions</i>	DSI-03.1	Do you provide open encryption methodologies (3.4ES, AES, etc.) to tenants in order for them to protect their data if it is required to move through public	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 Group	CID	Consensus Assessment Questions	AWS Response
		networks (e.g., the Internet)?	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 <i>Handling / Labeling / Security Policy</i>	DSI-04.1	Are policies and procedures established for labeling, handling and the security of data and objects that contain data?	AWS Customers retain control and ownership of their data and may implement a labeling and handling policy and procedures to meet their requirements.
	DSI-04.2	Are mechanisms for label inheritance implemented for objects that act as aggregate containers for data?	
Data Security & Information Lifecycle Management <i>Nonproduction Data</i>	DSI-05.1	Do you have procedures in place to ensure production data shall not be replicated or used in non-production environments?	AWS Customers retain control and ownership of their own data. AWS provides customers the ability to maintain and develop production and non-production environments. It is the responsibility of the customer to ensure that their production data is not replicated to non-production environments.
Data Security & Information Lifecycle Management <i>Ownership / Stewardship</i>	DSI-06.1	Are the responsibilities regarding data stewardship defined, assigned, documented and communicated?	AWS Customers retain control and ownership of their own data. Refer to the AWS Customer Agreement for additional information.

Control Group	CID	Consensus Assessment Questions	AWS Response
Data Security & Information Lifecycle Management <i>Secure Disposal</i>	DSI-07.1	Do you support secure deletion (e.g., degaussing/cryptographic wiping) of archived and backed-up data as determined by the tenant?	When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in NIST 800-88 ("Guidelines for Media Sanitization") as part of the decommissioning process. 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 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).
	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?	
Datacenter Security <i>Asset Management</i>	DCS-01.1	Do you maintain a complete inventory of all of your critical assets that includes ownership of the asset?	In alignment with ISO 27001 standards, AWS Hardware assets are assigned an owner, tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 8 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	DCS-01.2	Do you maintain a complete inventory of all of your critical supplier relationships?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Datacenter Security <i>Controlled Access Points</i>	DCS-02.1	Are physical security perimeters (e.g., fences, walls, barriers, guards, gates, electronic surveillance, physical authentication mechanisms, reception desks and security patrols) implemented?	Physical security controls include but are not limited to perimeter controls such as fencing, walls, security staff, video surveillance, intrusion detection systems and other electronic means. The AWS SOC reports provides additional details on the specific control activities executed by AWS. Refer to ISO 27001 standards; Annex A, domain 11 for further information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Datacenter Security <i>Equipment Identification</i>	DCS-03.1	Is automated equipment identification used as a method to validate connection authentication integrity based on known equipment location?	AWS manages equipment identification in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Datacenter Security <i>Offsite Authorization</i>	DCS-04.1	Do you provide tenants with documentation that describes scenarios in which data may be moved from one physical location to another? (e.g., offsite backups, business continuity failovers, replication)	AWS Customers can designate which physical region their data will be located. AWS will not move customers' content from the selected Regions without notifying the customer unless required to comply with the law or requests of governmental entities. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Datacenter Security <i>Offsite equipment</i>	DCS-05.1	Can you provide tenants with evidence documenting your policies and procedures governing asset management and repurposing of equipment?	In alignment with ISO 27001 standards, when a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in NIST 800-88 ("Guidelines for Media Sanitization") as part of the decommissioning process. 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.

Control Group	CID	Consensus Assessment Questions	AWS Response
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. For a complete list of available regions, see the AWS Global Infrastructure page.
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

Control Group	CID	Consensus Assessment Questions	AWS Response
Datacenter Security <i>User Access</i>	DCS-09.1	Do you restrict physical access to information assets and functions by users and support personnel?	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.
Encryption & Key Management <i>Entitlement</i>	EKM-01.1	Do you have key management policies binding keys to identifiable owners?	AWS provides customers the ability to use their own encryption mechanism for nearly all services including S3, EBS and EC2. VPC sessions are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Internally, AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications. AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
Encryption & Key Management <i>Key Generation</i>	EKM-02.1	Do you have a capability to allow creation of unique encryption keys per tenant?	AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS and EC2. IPsec tunnels to VPC are also encrypted. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). 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
	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?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	EKM-02.5	Do you utilize any third party/open source/proprietary frameworks to manage encryption keys?	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.
Encryption & Key Management <i>Encryption</i>	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 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-03.2	Do you leverage encryption to protect data and virtual machine images during transport across and between networks and hypervisor instances?	
	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	EKM-04.1	Do you have platform and data appropriate encryption that uses open/validated formats	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

Control Group	CID	Consensus Assessment Questions	AWS Response
Storage and Access		and standard algorithms?	Management Systems (KMS) to create and control encryption keys (refer to https://aws.amazon.com/kms/). Refer to AWS SOC reports for more details on KMS.
	EKM-04.2	Are your encryption keys maintained by the cloud consumer or a trusted key management provider?	AWS establishes and manages cryptographic keys for required cryptography employed within the AWS infrastructure. AWS produces, controls and distributes symmetric cryptographic keys using NIST approved key management technology and processes in the AWS information system. An AWS developed secure key and credential manager is used to create, protect and distribute symmetric keys and is used to secure and distribute: AWS credentials needed on hosts, RSA public/private keys and X.509 Certifications. AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.
	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 <i>Baseline Requirements</i>	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 <i>Risk Assessments</i>	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-faq/ .
Governance and Risk Management <i>Management Oversight</i>	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 <i>Management Program</i>	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

Control Group	CID	Consensus Assessment Questions	AWS Response
	GRM-04.2	Do you review your Information Security Management Program (ISMP) least once a year?	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/ .
Governance and Risk Management <i>Management Support / Involvement</i>	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).
Governance and Risk Management <i>Policy</i>	GRM-06.1	Do your information security and privacy policies align with industry standards (ISO-27001, ISO-22307, CoBIT, etc.)?	AWS manages third-party relationships in alignment with ISO 27001 standards.
	GRM-06.2	Do you have agreements to ensure your providers adhere to your information security and privacy policies?	AWS Third Party requirements are reviewed by independent external auditors during audits for our PCI DSS, ISO 27001 and FedRAMP compliance.
	GRM-06.3	Can you provide evidence of due diligence mapping of your controls, architecture and processes to regulations and/or standards?	Information about the AWS Compliance programs is published publicly on our website at http://aws.amazon.com/compliance/ .
	GRM-06.4	Do you disclose which controls, standards, certifications and/or regulations you comply with?	
Governance and Risk Management <i>Policy Enforcement</i>	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

Control Group	CID	Consensus Assessment Questions	AWS Response
	GRM-07.2	Are employees made aware of what actions could be taken in the event of a violation via their policies and procedures?	(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.
Governance and Risk Management <i>Business / Policy Change Impacts</i>	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 <i>Policy Reviews</i>	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 <i>Assessments</i>	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	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.

Control Group	CID	Consensus Assessment Questions	AWS Response
		subject to background verification?	
Human Resources <i>Employment Agreements</i>	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?	
	HRS-03.4	Is successful and timed completion of the training program considered a prerequisite for acquiring and maintaining access to sensitive systems?	
	HRS-03.5	Are personnel trained and provided with awareness programs at least once a year?	
Human Resources <i>Employment Termination</i>	HRS-04.1	Are documented policies, procedures and guidelines in place to govern change in employment and/or termination?	AWS Human Resources team defines internal management responsibilities to be followed for termination and role change of employees and vendors. AWS SOC reports provide additional details.

Control Group	CID	Consensus Assessment Questions	AWS Response
	HRS-04.2	Do the above procedures and guidelines account for timely revocation of access and return of assets?	Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked. AWS SOC reports provide further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 7 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Human Resources <i>Portable / Mobile Devices</i>	HRS-05.1	Are policies and procedures established and measures implemented to strictly limit access to your sensitive data and tenant data from portable and mobile devices (e.g. laptops, cell phones and personal digital assistants (PDAs)), which are generally higher-risk than non-portable devices (e.g., desktop computers at the provider organization's facilities)?	Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
Human Resources <i>Nondisclosure Agreements</i>	HRS-06.1	Are requirements for non-disclosure or confidentiality agreements reflecting the organization's needs for the protection of data and operational details identified, documented and reviewed at planned intervals?	Amazon Legal Counsel manages and periodically revises the Amazon NDA to reflect AWS business needs.

Control Group	CID	Consensus Assessment Questions	AWS Response
Human Resources <i>Roles / Responsibilities</i>	HRS-07.1	Do you provide tenants with a role definition document clarifying your administrative responsibilities versus those of the tenant?	The AWS Cloud Security Whitepaper and the AWS Risk and Compliance Whitepaper provide details on the roles and responsibilities of AWS and those of our Customers. The whitepapers are available at: http://aws.amazon.com/security and http://aws.amazon.com/compliance .
Human Resources <i>Acceptable Use</i>	HRS-08.1	Do you provide documentation regarding how you may or access tenant data and metadata?	AWS has a formal access control policy that is reviewed and updated on an annual basis (or when any major change to the system occurs that impacts the policy). The policy addresses purpose, scope, roles, responsibilities and management commitment. AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content. Refer to the ISO 27001 standard and 27018 code of practice for additional information. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 and ISO 27018.
	HRS-08.2	Do you collect or create metadata about tenant data usage through inspection technologies (search engines, etc.)?	
	HRS-08.3	Do you allow tenants to opt out of having their data/metadata accessed via inspection technologies?	
Human Resources <i>Training / Awareness</i>	HRS-09.1	Do you provide a formal, role-based, security awareness training program for cloud-related access and data management issues (e.g., multi-tenancy, nationality, cloud delivery model segregation of duties implications and conflicts of interest) for all persons with access to tenant data?	In alignment with ISO 27001 standard, all AWS employees complete periodic Information Security training which requires an acknowledgement to complete. Compliance audits are periodically performed to validate that employees understand and follow the established policies. AWS roles and responsibilities are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
	HRS-09.2	Are administrators and data stewards properly educated on their legal responsibilities with	

Control Group	CID	Consensus Assessment Questions	AWS Response
		regard to security and data integrity?	
Human Resources <i>User Responsibility</i>	HRS-10.1	Are users made aware of their responsibilities for maintaining awareness and compliance with published security policies, procedures, standards and applicable regulatory requirements?	AWS has implemented various methods of internal communication at a global level to help employees understand their individual roles and responsibilities and to communicate significant events in a timely manner. These methods include orientation and training programs for newly hired employee as well as electronic mail messages and the posting of information via the Amazon intranet. Refer to ISO 27001 standard, Annex A, domain 7 and 8. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. In addition the AWS Cloud Security Whitepaper provides further details - available at http://aws.amazon.com/security .
	HRS-10.2	Are users made aware of their responsibilities for maintaining a safe and secure working environment?	
	HRS-10.3	Are users made aware of their responsibilities for leaving unattended equipment in a secure manner?	
Human Resources <i>Workspace</i>	HRS-11.1	Do your data management policies and procedures address tenant and service level conflicts of interests?	AWS data management policies are in alignment with ISO 27001 standard. Refer to ISO 27001 standard, Annex A, domain 8 and 9. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. AWS SOC reports 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-
	HRS-11.2	Do your data management policies and procedures include a tamper audit or software integrity function for unauthorized access to tenant data?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	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?	related events in accordance with requirements. Audit records contain a set of data elements in order to support necessary analysis requirements. In addition, audit records are available for AWS Security team or other appropriate teams to perform inspection or analysis on demand, and in response to security-related or business-impacting events.
Identity & Access Management <i>Audit Tools Access</i>	IAM-01.1	Do you restrict, log and monitor access to your information security management systems? (E.g., hypervisors, firewalls, vulnerability scanners, network sniffers, APIs, etc.)	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	IAM-01.2	Do you monitor and log privileged access (administrator level) to information security management systems?	AWS has identified auditable event categories across systems and devices within the AWS system. Service teams configure the auditing features to record continuously the security-related events in accordance with requirements. The log storage system is designed to provide a highly scalable, highly available service that automatically increases capacity as the ensuing need for log storage grows. Audit records contain a set of data elements in order to support necessary analysis requirements. In addition, audit records are available for AWS Security team or other appropriate teams to perform inspection or analysis on demand, and in response to security-related or business-impacting events. Designated personnel on AWS teams receive automated alerts in the event of an audit processing failure. Audit processing failures include, for example, software/hardware errors. When alerted, on-call personnel issue a trouble ticket and track the event until it is resolved. AWS logging and monitoring processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP compliance.

Control Group	CID	Consensus Assessment Questions	AWS Response
Identity & Access Management <i>User Access Policy</i>	IAM-02.1	Do you have controls in place ensuring timely removal of systems access that is no longer required for business purposes?	AWS SOC reports provides further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 9 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	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?	
Identity & Access Management <i>Diagnostic / Configuration Ports Access</i>	IAM-03.1	Do you use dedicated secure networks to provide management access to your cloud service infrastructure?	Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored per the AWS access policy. In addition, customer data is and server instances are logically isolated from other customers by default. Privileged user access controls are reviewed by an independent auditor during the AWS SOC, ISO 27001, PCI, ITAR, and FedRAMP audits.
Identity & Access Management <i>Policies and Procedures</i>	IAM-04.1	Do you manage and store the identity of all personnel who have access to the IT infrastructure, including their level of access?	
	IAM-04.2	Do you manage and store the user identity of all personnel who have network access, including their level of access?	
Identity & Access Management <i>Segregation of Duties</i>	IAM-05.1	Do you provide tenants with documentation on how you maintain segregation of duties within your cloud service offering?	Customers retain the ability to manage segregations of duties of their AWS resources. Internally, AWS aligns with ISO 27001 standards for managing segregation of duties. Refer to ISO 27001 standard, Annex A, domain 6 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Identity & Access Management <i>Source Code Access Restriction</i>	IAM-06.1	Are controls in place to prevent unauthorized access to your application, program or object source code, and assure it is restricted to authorized personnel only?	In alignment with ISO 27001 standards, AWS has established formal policies, procedures to delineate the minimum standards for logical access to AWS resources. AWS SOC reports outlines the controls in place to manage access provisioning to AWS resources. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IAM-06.2	Are controls in place to prevent unauthorized access to tenant application, program or object source code, and assure it is restricted to authorized personnel only?	
Identity & Access Management <i>Third Party Access</i>	IAM-07.1	Do you provide multi-failure disaster recovery capability?	AWS provides customers the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area. AWS SOC reports provides further details. ISO 27001 standard Annex A, domain 15 provides additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification.
	IAM-07.2	Do you monitor service continuity with upstream providers in the event of provider failure?	
	IAM-07.3	Do you have more than one provider for each service you depend on?	
	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?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Identity & Access Management <i>User Access Restriction / Authorization</i>	IAM-08.1	Do you document how you grant and approve access to tenant data?	AWS Customers retain control and ownership of their data. Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored. In addition, customer data and server instances are logically isolated from other customers by default. Privileged user access controls are reviewed by an independent auditor during the AWS SOC, ISO 27001, PCI, ITAR, and FedRAMP audits.
	IAM-08.2	Do you have a method of aligning provider and tenant data classification methodologies for access control purposes?	
Identity & Access Management <i>User Access Authorization</i>	IAM-09.1	Does your management provision the authorization and restrictions for user access (e.g. employees, contractors, customers (tenants), business partners and/or suppliers) prior to their access to data and any owned or managed (physical and virtual) applications, infrastructure systems and network components?	Unique user identifiers are created as part of the onboarding workflow process in the AWS human resources management system. The device provisioning process helps ensure unique identifiers for devices. Both processes include manager approval to establish the user account or device. Initial authenticators are delivered to user's in-person and to devices as part of the provisioning process. Internal users can associate SSH public keys with their account. System account authenticators are provided to the requestor as part of the account creation process after the identity of the requestor is verified.
	IAM-09.2	Do you provide upon request user access (e.g. employees, contractors, customers (tenants), business partners and/or suppliers) to data and any owned or managed (physical and virtual) applications, infrastructure systems and network components?	AWS has established controls to address the threat of inappropriate insider access. All certifications and third-party attestations evaluate logical access preventative and detective controls. In addition, periodic risk assessments focus on how insider access is controlled and monitored.

Control Group	CID	Consensus Assessment Questions	AWS Response
Identity & Access Management <i>User Access Reviews</i>	IAM-10.1	Do you require at least annual certification of entitlements for all system users and administrators (exclusive of users maintained by your tenants)?	In alignment with ISO 27001 standard, all access grants are reviewed on a periodic basis; explicit re-approval is required or access to the resource is automatically revoked. Controls specific to User Access reviews are outlined in the SOC reports. Exceptions in the User entitlement controls are documented in the SOC reports. Refer to ISO 27001 standards, Annex A, domain 9 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IAM-10.2	If users are found to have inappropriate entitlements, are all remediation and certification actions recorded?	
	IAM-10.3	Will you share user entitlement remediation and certification reports with your tenants, if inappropriate access may have been allowed to tenant data?	
Identity & Access Management <i>User Access Revocation</i>	IAM-11.1	Is timely deprovisioning, revocation or modification of user access to the organizations systems, information assets and data implemented upon any change in status of employees, contractors, customers, business partners or involved third parties?	Access is automatically revoked when an employee's record is terminated in Amazon's Human Resources system. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked. AWS SOC reports provides further details on User access revocation. In addition the AWS Security White paper, section "Employee Lifecycle" provides additional information. Refer to ISO 27001 Annex A, domain 9 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IAM-11.2	Is any change in user access status intended to include termination of employment, contract or agreement, change of employment or transfer within the organization?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Identity & Access Management <i>User ID Credentials</i>	IAM-12.1	Do you support use of, or integration with, existing customer-based Single Sign On (SSO) solutions to your service?	The AWS Identity and Access Management (IAM) service provides identity federation to the AWS Management Console. Multi-factor authentication is an optional feature that a customer can utilize. Refer to the AWS website for additional details - http://aws.amazon.com/mfa. AWS Identity and Access Management (IAM) supports identity federation for delegated access to the AWS Management Console or AWS APIs. With identity federation, external identities (federated users) are granted secure access to resources in your AWS account without having to create IAM users. These external identities can come from your corporate identity provider (such as Microsoft Active Directory or from the AWS Directory Service) or from a web identity provider, such as Amazon Cognito, Login with Amazon, Facebook, Google or any OpenID Connect (OIDC) compatible provider.
	IAM-12.2	Do you use open standards to delegate authentication capabilities to your tenants?	
	IAM-12.3	Do you support identity federation standards (SAML, SPML, WS-Federation, etc.) as a means of authenticating/authorizing users?	
	IAM-12.4	Do you have a Policy Enforcement Point capability (e.g., XACML) to enforce regional legal and policy constraints on user access?	
	IAM-12.5	Do you have an identity management system (enabling classification of data for a tenant) in place to enable both role-based and context-based entitlement to data?	
	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?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	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 <i>Utility Programs Access</i>	IAM-13.1	Are utilities that can significantly manage virtualized partitions (e.g., shutdown, clone, etc.) appropriately restricted and monitored?	In alignment with ISO 27001 standards, system utilities are appropriately restricted and monitored. AWS SOC reports provides details on the specific control activities executed by AWS. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IAM-13.2	Do you have a capability to detect attacks that target the virtual infrastructure directly (e.g., shimming, Blue Pill, Hyper jumping, etc.)?	
	IAM-13.3	Are attacks that target the virtual infrastructure	

Control Group	CID	Consensus Assessment Questions	AWS Response
		prevented with technical controls?	
Infrastructure & Virtualization Security <i>Audit Logging / Intrusion Detection</i>	IVS-01.1	Are file integrity (host) and network intrusion detection (IDS) tools implemented to help facilitate timely detection, investigation by root cause analysis and response to incidents?	AWS Incident response program (detection, investigation and response to incidents) has been developed in alignment with ISO 27001 standards, system utilities are appropriately restricted and monitored. AWS SOC reports provides additional details on controls in place to restrict system access. Refer to AWS Overview of Security Processes for additional details - available at http://aws.amazon.com/security .
	IVS-01.2	Is physical and logical user access to audit logs restricted to authorized personnel?	
	IVS-01.3	Can you provide evidence that due diligence mapping of regulations and standards to your controls/architecture/processes has been done?	
	IVS-01.4	Are audit logs centrally stored and retained?	In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard. 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
	IVS-01.5	Are audit logs reviewed on a regular basis for security events (e.g., with automated tools)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
			are quickly and reliably communicated to operations personnel. Refer to AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
Infrastructure & Virtualization Security <i>Change Detection</i>	IVS-02.1	Do you log and alert any changes made to virtual machine images regardless of their running state (e.g. dormant, off or running)?	Virtual Machines are assigned to customers as a part of the EC2 service. Customers retain control over what resources are being used and where resources reside. Refer to the AWS website for additional details - http://aws.amazon.com .
	IVS-02.2	Are changes made to virtual machines, or moving of an image and subsequent validation of the image's integrity, made immediately available to customers through electronic methods (e.g. portals or alerts)?	
Infrastructure & Virtualization Security <i>Clock Synchronization</i>	IVS-03.1	Do you use a synchronized time-service protocol (e.g., NTP) to ensure all systems have a common time reference?	In alignment with ISO 27001 standards, AWS information systems utilize internal system clocks synchronized via NTP (Network Time Protocol). AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
Infrastructure & Virtualization Security <i>Capacity / Resource Planning</i>	IVS-04.1	Do you provide documentation regarding what levels of system (network, storage, memory, I/O, etc.) oversubscription you maintain and	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 .

Control Group	CID	Consensus Assessment Questions	AWS Response
		under what circumstances/scenarios?	AWS manages capacity and utilization data in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-04.2	Do you restrict use of the memory oversubscription capabilities present in the hypervisor?	
	IVS-04.3	Do your system capacity requirements take into account current, projected and anticipated capacity needs for all systems used to provide services to the tenants?	
	IVS-04.4	Is system performance monitored and tuned in order to continuously meet regulatory, contractual and business requirements for all the systems used to provide services to the tenants?	
Infrastructure & Virtualization Security <i>Management - Vulnerability Management</i>	IVS-05.1	Do security vulnerability assessment tools or services accommodate the virtualization technologies being used (e.g. virtualization aware)?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. Regular internal and external vulnerability scans are performed on the host operating system, web application and databases in the AWS environment utilizing a variety of tools. Vulnerability scanning and remediation practices are regularly reviewed as a part of

Control Group	CID	Consensus Assessment Questions	AWS Response
			AWS continued compliance with PCI DSS and FedRAMP.
Infrastructure & Virtualization Security <i>Network Security</i>	IVS-06.1	For your IaaS offering, do you provide customers with guidance on how to create a layered security architecture equivalence using your virtualized solution?	AWS website provides guidance on creating a layered security architecture in a number of white papers available via the AWS public website - http://aws.amazon.com/documentation/ .
	IVS-06.2	Do you regularly update network architecture diagrams that include data flows between security domains/zones?	Boundary protection devices that employ rule sets, access control lists (ACL), and configurations enforce the flow of information between network fabrics. Several network fabrics exist at Amazon, each separated by devices that control the flow of information between fabrics. The flow of information between fabrics is established by approved authorizations, which exist as access control lists (ACL) which reside on these devices. These devices control the flow of information between fabrics as mandated by these ACLs. ACLs are defined, approved by appropriate personnel, managed and deployed using AWS ACL-manage tool.
	IVS-06.3	Do you regularly review for appropriateness the allowed access/connectivity (e.g., firewall rules) between security domains/zones within the network?	
	IVS-06.4	Are all firewall access control lists documented with business justification?	Amazon's Information Security team approves these ACLs. Approved firewall rule sets and access control lists between network fabrics restrict the flow of information to specific

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>OS Hardening and Base Controls</i>	IVS-07.1	Are operating systems hardened to provide only the necessary ports, protocols and services to meet business needs using technical controls (i.e. antivirus, file integrity monitoring and logging) as part of their baseline build standard or template?	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.
Infrastructure & Virtualization Security <i>Production / Nonproduction Environments</i>	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

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>Segmentation</i>	IVS-09.1	Are system and network environments protected by a firewall or virtual firewall to ensure business and customer security requirements?	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.
	IVS-09.2	Are system and network environments protected by a firewall or virtual firewall to ensure compliance with legislative, regulatory and contractual requirements?	
	IVS-09.3	Are system and network environments protected by a firewall or virtual firewall to ensure separation of production and non-production environments?	
	IVS-09.4	Are system and network environments protected by a firewall or virtual firewall to ensure protection and isolation of sensitive data?	
Infrastructure & Virtualization Security <i>VM Security - vMotion Data Protection</i>	IVS-10.1	Are secured and encrypted communication channels used when migrating physical servers, applications or data to virtual servers?	AWS provides customers the ability to use their own encryption mechanism for nearly all services including S3, EBS and EC2. VPC sessions are also encrypted.
	IVS-10.2	Do you use a network segregated from production-level networks when migrating physical servers, applications or data to virtual servers?	AWS Customers retain control and ownership of their own data. AWS provides customers the ability to maintain and develop production and non-production environments. It is the responsibility of the customer to ensure that their production data is not replicated to non-production environments.

Control Group	CID	Consensus Assessment Questions	AWS Response
Infrastructure & Virtualization Security <i>VMM Security - Hypervisor Hardening</i>	IVS-11.1	Do you restrict personnel access to all hypervisor management functions or administrative consoles for systems hosting virtualized systems based on the principle of least privilege and supported through technical controls (e.g. two-factor authentication, audit trails, IP address filtering, firewalls and TLS-encapsulated communications to the administrative consoles)?	AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. When user accounts are created, user accounts are created to have minimal access. Access above these least privileges requires appropriate authorization. Refer to AWS SOC reports for more information on Access Controls.
Infrastructure & Virtualization Security <i>Wireless Security</i>	IVS-12.1	Are policies and procedures established and mechanisms configured and implemented to protect the wireless network environment perimeter and to restrict unauthorized wireless traffic?	Policies, procedures and mechanisms to protect AWS network environment are in place. AWS security controls are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
	IVS-12.2	Are policies and procedures established and mechanisms implemented to ensure wireless security settings are enabled with strong encryption for authentication and transmission, replacing vendor default settings? (e.g., encryption keys, passwords, SNMP community strings)	

Control Group	CID	Consensus Assessment Questions	AWS Response
	IVS-12.3	Are policies and procedures established and mechanisms implemented to protect wireless network environments and detect the presence of unauthorized (rogue) network devices for a timely disconnect from the network?	
Infrastructure & Virtualization Security <i>Network Architecture</i>	IVS-13.1	Do your network architecture diagrams clearly identify high-risk environments and data flows that may have legal compliance impacts?	AWS Customers retain responsibility to manage their own network segmentation in adherence with their defined requirements. Internally, AWS network segmentation is aligned with the ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	IVS-13.2	Do you implement technical measures and apply defense-in-depth techniques (e.g., deep packet analysis, traffic throttling and black-holing) for detection and timely response to network-based attacks associated with anomalous ingress or egress traffic patterns (e.g., MAC spoofing and ARP poisoning attacks) and/or distributed denial-of-service (DDoS) attacks?	AWS Security regularly scans all Internet facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. In addition, external vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment. AWS security controls are reviewed by independent external auditors during audits for our SOC, PCI DSS, ISO 27001 and FedRAMP compliance.
Interoperability & Portability <i>APIs</i>	IPY-01	Do you publish a list of all APIs available in the service and indicate which are	Details regarding AWS APIs can be found on the AWS website at https://aws.amazon.com/documentation/ .

Control Group	CID	Consensus Assessment Questions	AWS Response
		standard and which are customized?	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 <i>Data Request</i>	IPY-02	Is unstructured customer data available on request in an industry-standard format (e.g., .doc, .xls, or .pdf)?	
Interoperability & Portability <i>Policy & Legal</i>	IPY-03.1	Do you provide policies and procedures (i.e. service level agreements) governing the use of APIs for interoperability between your service and third-party applications?	
	IPY-03.2	Do you provide policies and procedures (i.e. service level agreements) governing the migration of application data to and from your service?	Customer retain control and ownership of their content. Customers can choose how they migrate applications and content both on and off the AWS platform at their discretion.
Interoperability & Portability <i>Standardized Network Protocols</i>	IPY-04.1	Can data import, data export and service management be conducted over secure (e.g., non-clear text and authenticated), industry accepted standardized network protocols?	AWS allows customers to move data as needed on and off AWS storage. Refer to http://aws.amazon.com/choosing-a-cloud-platform for more information on Storage options.
	IPY-04.2	Do you provide consumers (tenants) with documentation detailing the relevant interoperability and portability network protocol standards that are involved?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Interoperability & Portability <i>Virtualization</i>	IPY-05.1	Do you use an industry-recognized virtualization platform and standard virtualization formats (e.g., OVF) to help ensure interoperability?	Amazon EC2 currently utilizes a highly customized version of the Xen hypervisor. The hypervisor is regularly assessed for new and existing vulnerabilities and attack vectors by internal and external penetration teams, and is well suited for maintaining strong isolation between guest virtual machines. The AWS Xen hypervisor security is regularly evaluated by independent auditors during assessments and audits. Refer to the AWS Cloud Security Whitepaper for additional details - available at http://aws.amazon.com/security .
	IPY-05.2	Do you have documented custom changes made to any hypervisor in use, and all solution-specific virtualization hooks available for customer review?	
Mobile Security <i>Anti-Malware</i>	MOS-01	Do you provide anti-malware training specific to mobile devices as part of your information security awareness training?	AWS' program, processes and procedures to managing antivirus / malicious software is in alignment with ISO 27001 standards. Refer to ISO 27001 standard, Annex A, domain 12 for additional information.
Mobile Security <i>Application Stores</i>	MOS-02	Do you document and make available lists of approved application stores for mobile devices accessing or storing company data and/or company systems?	AWS has established an information security framework and policies and has effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.1 and the National Institute of Standards and Technology (NIST) Publication 800-53 (Recommended Security Controls for Federal Information Systems). Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.
Mobile Security <i>Approved Applications</i>	MOS-03	Do you have a policy enforcement capability (e.g., XACML) to ensure that only approved applications and those from approved application stores be loaded onto a mobile device?	
Mobile Security <i>Approved Software for BYOD</i>	MOS-04	Does your BYOD policy and training clearly state which applications and applications stores are	

Control Group	CID	Consensus Assessment Questions	AWS Response
		approved for use on BYOD devices?	
Mobile Security <i>Awareness and Training</i>	MOS-05	Do you have a documented mobile device policy in your employee training that clearly defines mobile devices and the accepted usage and requirements for mobile devices?	
Mobile Security <i>Cloud Based Services</i>	MOS-06	Do you have a documented list of pre-approved cloud based services that are allowed to be used for use and storage of company business data via a mobile device?	
Mobile Security <i>Compatibility</i>	MOS-07	Do you have a documented application validation process for testing device, operating system and application compatibility issues?	
Mobile Security <i>Device Eligibility</i>	MOS-08	Do you have a BYOD policy that defines the device(s) and eligibility requirements allowed for BYOD usage?	
Mobile Security <i>Device Inventory</i>	MOS-09	Do you maintain an inventory of all mobile devices storing and accessing company data which includes device status (os system and patch levels, lost or decommissioned, device assignee)?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Mobile Security <i>Device Management</i>	MOS-10	Do you have a centralized mobile device management solution deployed to all mobile devices that are permitted to store, transmit, or process company data?	
Mobile Security <i>Encryption</i>	MOS-11	Does your mobile device policy require the use of encryption for either the entire device or for data identified as sensitive enforceable through technology controls for all mobile devices?	
Mobile Security <i>Jailbreaking and Rooting</i>	MOS-12.1	Does your mobile device policy prohibit the circumvention of built-in security controls on mobile devices (e.g., jailbreaking or rooting)?	
	MOS-12.2	Do you have detective and preventative controls on the device or via a centralized device management system which prohibit the circumvention of built-in security controls?	
Mobile Security <i>Legal</i>	MOS-13.1	Does your BYOD policy clearly define the expectation of privacy, requirements for litigation, e-discovery and legal holds?	Customers retain the control and responsibility of their data and associated media assets. It is the responsibility of the customer to manage mobile security devices and the access to the customer's content.

Control Group	CID	Consensus Assessment Questions	AWS Response
	MOS-13.2	Do you have detective and preventative controls on the device or via a centralized device management system which prohibit the circumvention of built-in security controls?	
Mobile Security <i>Lockout Screen</i>	MOS-14	Do you require and enforce via technical controls an automatic lockout screen for BYOD and company owned devices?	
Mobile Security <i>Operating Systems</i>	MOS-15	Do you manage all changes to mobile device operating systems, patch levels and applications via your company's change management processes?	
Mobile Security <i>Passwords</i>	MOS-16.1	Do you have password policies for enterprise issued mobile devices and/or BYOD mobile devices?	
	MOS-16.2	Are your password policies enforced through technical controls (i.e. MDM)?	
	MOS-16.3	Do your password policies prohibit the changing of authentication requirements (i.e. password/PIN length) via a mobile device?	
Mobile Security <i>Policy</i>	MOS-17.1	Do you have a policy that requires BYOD users to perform backups of specified corporate data?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	MOS-17.2	Do you have a policy that requires BYOD users to prohibit the usage of unapproved application stores?	
	MOS-17.3	Do you have a policy that requires BYOD users to use anti-malware software (where supported)?	
Mobile Security <i>Remote Wipe</i>	MOS-18.1	Does your IT provide remote wipe or corporate data wipe for all company-accepted BYOD devices?	
	MOS-18.2	Does your IT provide remote wipe or corporate data wipe for all company-assigned mobile devices?	
Mobile Security <i>Security Patches</i>	MOS-19.1	Do your mobile devices have the latest available security-related patches installed upon general release by the device manufacturer or carrier?	
	MOS-19.2	Do your mobile devices allow for remote validation to download the latest security patches by company IT personnel?	
Mobile Security <i>Users</i>	MOS-20.1	Does your BYOD policy clarify the systems and servers allowed for use or access on the BYOD-enabled device?	
	MOS-20.2	Does your BYOD policy specify the user roles that are allowed	

Control Group	CID	Consensus Assessment Questions	AWS Response
		access via a BYOD-enabled device?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Contact / Authority Maintenance</i>	SEF-01.1	Do you maintain liaisons and points of contact with local authorities in accordance with contracts and appropriate regulations?	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 <i>Incident Management</i>	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. The AWS Cloud Security Whitepaper (available at http://aws.amazon.com/security) provides additional details.
	SEF-02.2	Do you integrate customized tenant requirements into your security incident response plans?	
	SEF-02.3	Do you publish a roles and responsibilities document specifying what you vs. your tenants are responsible for during security incidents?	
	SEF-02.4	Have you tested your security incident response plans in the last year?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Reporting</i>	SEF-03.1	Does your security information and event management (SIEM) system merge data sources (app logs, firewall logs, IDS logs, physical access logs, etc.) for granular analysis and alerting?	

Control Group	CID	Consensus Assessment Questions	AWS Response
	SEF-03.2	Does your logging and monitoring framework allow isolation of an incident to specific tenants?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Response</i> <i>Legal Preparation</i>	SEF-04.1	Does your incident response plan comply with industry standards for legally admissible chain-of-custody management processes and controls?	
	SEF-04.2	Does your incident response capability include the use of legally admissible forensic data collection and analysis techniques?	
	SEF-04.3	Are you capable of supporting litigation holds (freeze of data from a specific point in time) for a specific tenant without freezing other tenant data?	
	SEF-04.4	Do you enforce and attest to tenant data separation when producing data in response to legal subpoenas?	
Security Incident Management, E-Discovery & Cloud Forensics <i>Incident Response Metrics</i>	SEF-05.1	Do you monitor and quantify the types, volumes and impacts on all information security incidents?	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 <i>Data Quality and Integrity</i>	STA-01.1	Do you inspect and account for data quality errors and associated risks, and work with your cloud supply-chain partners to correct them?	Customers retain control and ownership over the quality of their data and potential quality errors that may arise through their usage of AWS services. Refer to AWS SOC report for specific details in relation to Data Integrity and Access Management (including least privilege access)
	STA-01.2	Do you design and implement controls to mitigate and contain data security risks through proper separation of duties, role-based access, and least-privileged access for all personnel within your supply chain?	
Supply Chain Management, Transparency and Accountability <i>Incident Reporting</i>	STA-02.1	Do you make security incident information available to all affected customers and providers periodically through electronic methods (e.g. portals)?	AWS' incident response program, plans and procedures have been developed in alignment with ISO 27001 standard. The AWS SOC reports provides details on the specific control activities executed by AWS. The AWS Cloud Security Whitepaper (available at http://aws.amazon.com/security) provides additional details.
Supply Chain Management, Transparency and Accountability <i>Network / Infrastructure Services</i>	STA-03.1	Do you collect capacity and use data for all relevant components of your cloud service offering?	AWS manages capacity and utilization data in alignment with ISO 27001 standard. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	STA-03.2	Do you provide tenants with capacity planning and use reports?	
Supply Chain Management, Transparency and Accountability <i>Provider Internal Assessments</i>	STA-04.1	Do you perform annual internal assessments of conformance and effectiveness of your policies, procedures, and supporting measures and metrics?	AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 15 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.

Control Group	CID	Consensus Assessment Questions	AWS Response
Supply Chain Management, Transparency and Accountability <i>Third Party Agreements</i>	STA-05.1	Do you select and monitor outsourced providers in compliance with laws in the country where the data is processed, stored and transmitted?	Personnel security requirements for third-party providers supporting AWS systems and devices are established in a Mutual Non-Disclosure Agreement between AWS' parent organization, Amazon.com, and the respective third-party provider. The Amazon Legal Counsel and the AWS Procurement team define AWS third party provider personnel security requirements in contract agreements with the third party provider. All persons working with AWS information must at a minimum, meet the screening process for pre-employment background checks and sign a Non-Disclosure Agreement (NDA) prior to being granted access to AWS information. AWS does not generally outsource development of AWS services to subcontractors.
	STA-05.2	Do you select and monitor outsourced providers in compliance with laws in the country where the data originates?	
	STA-05.3	Does legal counsel review all third-party agreements?	
	STA-05.4	Do third-party agreements include provision for the security and protection of information and assets?	
	STA-05.5	Do you provide the client with a list and copies of all sub processing agreements and keep this updated?	
Supply Chain Management, Transparency and Accountability <i>Supply Chain Governance Reviews</i>	STA-06.1	Do you review the risk management and governance processes of partners to account for risks inherited from other members of that partner's supply chain?	AWS maintains formal agreements with key third party suppliers and implements appropriate relationship management mechanisms in line with their relationship to the business. AWS' third party management processes are reviewed by independent auditors as part of AWS ongoing compliance with SOC and ISO 27001.

Control Group	CID	Consensus Assessment Questions	AWS Response
Supply Chain Management, Transparency and Accountability <i>Supply Chain Metrics</i>	STA-07.1	Are policies and procedures established, and supporting business processes and technical measures implemented, for maintaining complete, accurate and relevant agreements (e.g., SLAs) between providers and customers (tenants)?	
	STA-07.2	Do you have the ability to measure and address non-conformance of provisions and/or terms across the entire supply chain (upstream/downstream)?	
	STA-07.3	Can you manage service-level conflicts or inconsistencies resulting from disparate supplier relationships?	
	STA-07.4	Do you review all agreements, policies and processes at least annually?	
Supply Chain Management, Transparency and Accountability <i>Third Party Assessment</i>	STA-08.1	Do you assure reasonable information security across your information supply chain by performing an annual review?	
	STA-8.2	Does your annual review include all partners/third-party providers upon which your information supply chain depends?	

Control Group	CID	Consensus Assessment Questions	AWS Response
Supply Chain Management, Transparency and Accountability <i>Third Party Audits</i>	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 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.
	STA-09.2	Do you have external third party services conduct vulnerability scans and periodic penetration tests on your applications and networks?	
Threat and Vulnerability Management <i>Antivirus / Malicious Software</i>	TVM-01.1	Do you have anti-malware programs that support or connect to your cloud service offerings installed on all of your systems?	AWS' program, processes and procedures to managing antivirus / malicious software is in alignment with ISO 27001 standards. Refer to AWS SOC reports provides further details. In addition, refer to ISO 27001 standard, Annex A, domain 12 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification standard.
	TVM-01.2	Do you ensure that security threat detection systems using signatures, lists or behavioral patterns are updated across all infrastructure components within industry accepted time frames?	
Threat and Vulnerability Management <i>Vulnerability / Patch Management</i>	TVM-02.1	Do you conduct network-layer vulnerability scans regularly as prescribed by industry best practices?	Customers retain control of their own guest operating systems, software and applications and are responsible for performing vulnerability scans and patching of their own systems. Customers can request permission to conduct scans of their cloud infrastructure as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. AWS Security regularly scans all Internet-facing service endpoint IP addresses for vulnerabilities. AWS Security notifies the appropriate parties to remediate any identified vulnerabilities. AWS' own maintenance and system patching generally do not impact customers. Refer to AWS Cloud Security Whitepaper for further information - available at
	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	

Control Group	CID	Consensus Assessment Questions	AWS Response
		by industry best practices?	
	TVM-02.4	Will you make the results of vulnerability scans available to tenants at their request?	
	TVM-02.5	Do you have a capability to rapidly patch vulnerabilities across all of your computing devices, applications and systems?	
	TVM-02.6	Will you provide your risk-based systems patching time frames to your tenants upon request?	
Threat and Vulnerability Management <i>Mobile Code</i>	TVM-03.1	Is mobile code authorized before its installation and use, and the code configuration checked, to ensure that the authorized mobile code operates according to a clearly defined security policy?	AWS allows customers to manage client and mobile applications to their own requirements.
	TVM-03.2	Is all unauthorized mobile code prevented from executing?	

Further Reading

For additional information, see the following sources:

- [AWS Risk and Compliance Overview](#)
- [AWS Certifications, Programs, Reports, and Third-Party Attestations](#)
- [AWS Answers to Key Compliance Questions](#)

Document Revisions

Date	Description
January 2017	Migrated to new template.
January 2016	First publication

Attachment E
Service Offering EULAs, SLAs

None