

Contract # AR2485**STATE OF UTAH COOPERATIVE CONTRACT**

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

Insight Public Sector, Inc.

Name

6820 S. Harl AvenueTempeAZ85283

City

State

Zip

LEGAL STATUS OF CONTRACTOR

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

Contact Person: Erica Falchetti Phone: 480-760-9488 Email: Erica.Falchetti@Insight.com

Vendor # Commodity Code #920-05

2. GENERAL PURPOSE OF CONTRACT: Contractor is permitted to provide the Cloud Solutions identified in Attachment B to Participating States once a Participating Addendum has been signed.
3. PROCUREMENT PROCESS: This contract is entered into as a result of the procurement process on Bid#CH16012.
4. CONTRACT PERIOD: Effective Date: 09/30/2016 Termination Date: 09/15/2026 unless terminated early or extended in accordance with the terms and conditions of this contract. Pursuant to Solicitation #CH16012, Contractor must re-certify its qualifications each year.
5. Administrative Fee, as described in the Solicitation and Attachment A: The Contractor shall pay to NASPO ValuePoint, or its assignee, a NASPO ValuePoint Administrative Fee of one-quarter of one percent (0.25% or 0.0025) 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.
6. ATTACHMENT A: NASPO ValuePoint Master Terms and Conditions
ATTACHMENT B: Scope of Services Awarded to Contractor
ATTACHMENT C: Pricing Discounts and Pricing Schedule
ATTACHMENT D: Contractor's Response to Solicitation #CH16012
ATTACHMENT E: Service Provider Terms and Conditions

Any conflicts between Attachment A and the other Attachments will be resolved in favor of Attachment A.

8. DOCUMENTS INCORPORATED INTO THIS CONTRACT BY REFERENCE BUT NOT ATTACHED:
- All other governmental laws, regulations, or actions applicable to the goods and/or services authorized by this contract.
 - Utah State Procurement Code and the Procurement Rules.
9. Each signatory below represents that he or she has the requisite authority to enter into this contract.

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

CONTRACTOR

STATE

02/06/2017

Contractor's signature

Date

Spencer Hall
Director, Division of Purchasing2.8.17
DateJohn Carnahan SVP- Business Development

Type or Print Name and Title

Spencer Hall

Division of Purchasing Contact Person

801-538-3307

Telephone Number


801-538-3882

Fax Number

spencerh@utah.gov

Email

(Revision 16 June 2016)



Attachment A: NASPO ValuePoint Master Agreement Terms and Conditions

1. Master Agreement Order of Precedence

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

- (1) A Participating Entity's Participating 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 is non-public information that is designated "confidential" or that a reasonable person should understand to be confidential, including (1) Customer Data; (2) any Purchasing Entity's records, (3) personnel records, and (4) information concerning individuals. Confidential Information does not include information that (a) becomes publicly available without a breach of this agreement, (b) was lawfully known or received by the receiving party without an obligation to keep it confidential, (c) is independently developed, or (d) is a comment or suggestion one party volunteers about the other's business, products or services.

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

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.

Master Agreement.

Customer Data means all data, including all text, sound, software, or image files that are provided to Service Provider by, or on behalf of, a Purchasing Entity through its use of the Online Services. All references to “**Data**” in the Master Agreement shall be deemed to mean Customer 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) shall have the same meaning as the term "protected health information" in 45 CFR § 160.103 of HIPAA, provided that it is limited to such protected health information that is received by Service Provider from Customer, or created, received, maintained, or transmitted by Service Provider on behalf of, Customer.

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 any unlawful access, use, theft or destruction to any Customer Data stored on Service Provider's equipment or in Service Provider's facilities, or unauthorized access to such equipment or facilities resulting in use, theft, loss, disclosure, alteration or destruction of Customer Data. All references to "Data Breach" in the Master Agreement shall be deemed to mean Security Incident.

Service Level Agreement (SLA) means the service levels or service level agreements, if any, set forth in the Service Provider Terms."

Service Provider means a provider of the Cloud Services that are available for resale through Contractor under this Master Agreement.

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.

Terms of Use means the terms and conditions associated with the use of the Cloud Services by the Participating Entity set forth in the Contractor's Cloud Services Order Form.

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

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

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

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

6. Discount Guarantee Period: All discounts must be guaranteed for the entire term of the Master Agreement. Participating Entities and Purchasing Entities shall receive the immediate benefit of price or rate reduction of the services provided under this Master Agreement. A price or rate reduction will apply automatically to the Master Agreement and an amendment is not necessary.

7. Termination: Unless otherwise stated, this Master Agreement may be terminated by either party upon 60 days written notice prior to the effective date of the termination. Further, any Participating Entity may terminate its participation upon 30 days written notice, unless otherwise limited or stated in the Participating Addendum. Termination may be in whole or in part. Any termination under this provision shall not affect the rights and obligations attending orders outstanding at the time of termination, including any right of any Purchasing Entity to indemnification by the Contractor, rights of payment for Services delivered and accepted, data ownership, Contractor obligations regarding Purchasing Entity Data, rights attending default in performance an applicable

Service Level of Agreement in association with any Order, Contractor obligations under Termination and Suspension of Service, and any responsibilities arising out of a Security Incident or Data Breach. Termination of the Master Agreement due to Contractor default may be immediate if defaults cannot be reasonably cured as allowed per Default and Remedies terms.

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. 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 of Contractor's Employees 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 of Contractor's Employees. 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.

e. Notwithstanding the foregoing, damages attributable to Security Incidents shall be subject to the Section of the Master Agreement titled "Limitation of Liability."

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 either party to this Master Agreement or to a Participating State or Purchasing Entity, or the appointment of a receiver or similar officer for any such party or any of such party's 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, the party claiming default shall issue a written notice of default, identifying the nature of the default, and providing a period of

30 calendar days in which the non-defaulting party 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 defaulting party's liability for damages.

c. If a defaulting party is afforded an opportunity to cure and fails to cure the default within the period specified in the written notice of default, the defaulting party shall be in breach of its obligations under this Master Agreement and the non-defaulting party 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) In the event of default by the Contractor, and to the extent permitted by the law of the Participating State or Purchasing Entity, the Lead State shall have the right to 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. **(1) Contract Indemnification:** 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.

(2) Participating Entities Indemnification: The Participating Entities shall defend and indemnify Contractor for, from, and against any losses, damages, penalties, costs, and expenses, including, without limitation, reasonable attorney fees incurred by Contractor in connection with any claims or actions by Service Provider or other third parties arising out of or resulting from (i) Client Data passing through the Cloud Services and/or Service Provider's network to or from the Participating Entity, (ii) unauthorized or misuse of Cloud Services by Client, its employees or agents (excluding any claims that the Cloud Services, as provided by Service Provider, infringe third-party intellectual property rights), (iii) Participating Entity's failure to comply with applicable law, (iv) Participating Entity's failure to pay Contractor for the full Term, regardless of Service Provider performance issues, and/or (v) Participating Entity's failure to comply with these Terms of Sale.

b. Contractor Indemnification. 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, 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 Indemnifying Party within a reasonable time after receiving notice of a Claim. Even if the Indemnified Party fails to provide

reasonable notice, the Indemnifying Party shall not be relieved from its obligations unless the Indemnifying Party can demonstrate that it was prejudiced in defending the Claim resulting in increased expenses or loss to the Indemnifying Party and then only to the extent of the prejudice or expenses. If the Indemnifying Party promptly and reasonably investigates and defends any 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 Indemnify Party's reasonable request and expense, information and assistance necessary for such defense. If the Indemnifying Party fails to vigorously pursue the defense or settlement of the Claim, the Indemnified Party may assume the defense or settlement of it and the Indemnifying Party 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 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

Coverage shall be written on an occurrence basis, cyber liability and professional liability is written on a claims made 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, bodily injury (including death), advertising liability, and property damage, with a limit of not less than \$1 million per occurrence/\$2 million general aggregate;

(2) Tech E&O Coverage

Technology Errors and Omissions Minimum Insurance Coverage including Professional Liability/Risk/Data Breach and Privacy/Cyber in the amount of \$5,000,000 in the aggregate.

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

c. Contractor shall pay premiums on all insurance policies.

d. Prior to commencement of performance, Contractor shall provide to the Lead State a blanket 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 on a blanket basis, (2) provides that cancellation, non-renewal, or expiration of the coverage contained in such policy shall have be in accordance with policy terms and conditions, 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; 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,); 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 laws (including but not limited to privacy and security related laws) applicable to IT service providers.

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. Financial obligations of Participating States are limited to the orders placed by the departments or other state agencies and institutions having available funds. Participating States incur no financial obligations on behalf of political subdivisions.

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

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

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

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

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.

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

22. Data Access Controls: 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 in the course of data center operations, response to service or technical issues, as necessary for the operation and maintenance of the service, 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 will protect Customer Data against accidental, unauthorized or unlawful access, disclosure, alteration, loss or destruction.

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

25. Purchasing Entity Data: Purchasing Entity retains full right and title to Data provided by it.

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

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 a license to use that API for the duration of that applicable Product subscription.

30. Data Privacy: The Contractor must comply with all applicable laws (including but not limited to data privacy and security related laws) applicable to IT service providers. 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 resell the Services described in this Master Agreement.

b. Contractor will perform materially as described in this Master Agreement 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.

32. Transition Assistance: The Contractor shall assist a Purchasing Entity if requested, by providing guidance, in exporting and extracting a Purchasing Entity's Data. Any transition services requested by a Purchasing Entity involving knowledge transfer or guidance and support shall be subject to a separation transition Statement of Work.

33. Waiver of Breach: Failure of a party 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 a party must be in writing. Waiver by a party of any default, right or remedy under this Master Agreement or Participating Addendum, or by Purchasing Entity with respect to any Purchase Order, or breach of any terms or requirements of this Master Agreement, a Participating Addendum, or Purchase Order shall not be construed or operate as a waiver of any subsequent default or breach of

such term or requirement, or of any other term or requirement under this Master Agreement, Participating Addendum, or Purchase Order.

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

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

36. Performance and Payment Time Frames that Exceed Contract Duration: All maintenance or other agreements for services entered into during the duration of an SLA and whose performance and payment time frames extend beyond the duration of this Master Agreement shall remain in effect for performance and payment purposes (limited to the time frame and services established per each written agreement). No new leases, maintenance or other agreements for services may be executed after the Master Agreement has expired. For the purposes of this section, renewals of maintenance, subscriptions, SaaS subscriptions and agreements, and other service agreements, shall not be considered as "new."

37. Governing Law and Venue

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

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

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

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

38. No Guarantee of Service Volumes: The Contractor acknowledges and agrees that the Lead State and NASPO ValuePoint makes no representation, warranty or condition as to the nature, timing, quality, quantity or volume of business for the Services or any other products and services that the Contractor may realize from this Master Agreement, or the compensation that may be earned by the Contractor by offering the Services. The Contractor acknowledges and agrees that it has conducted its own due diligence prior to entering into this Master Agreement as to all the foregoing matters.

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

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

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

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

41. Government Support: No support, facility space, materials, special access, personnel or other obligations on behalf of the states or other Participating Entities,

other than payment, are required under the Master Agreement.

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

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

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

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

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

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

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

43. Limitation of Liability for Contractor:

a. Direct Damages Limitation. If Contractor, at its option, re-performs the Services or replaces the Product that is the subject of, or gave rise to, the claim, Contractor's total liability will be limited to such re-performance. If (i) the claim or matter cannot be remedied by such re-performance; (ii) re-performance is not an applicable remedy; or (iii) re-performance is not provided, then Contractor's total aggregate liability for any and all claims under this Master Agreement will be limited to and shall not exceed: (i) an amount equal to two (2x) times the total amount paid by Participating Entity to Contractor for the Cloud Services under the Purchas Order giving rise to the party's claim (said amount not to exceed a total of twelve months (12) months charges under the applicable purchase order) including indirect damages or (ii) or \$1,000,000, whichever is greater.

b. Indirect/Special Damages. Except for fraud and Participating Entity's obligations under the subsection titled "Indemnification," neither party will be liable for any indirect, special, incidental or consequential damages, nor damages for loss of business profits, business interruption, loss of business information and the like, arising in any way out of the order, any of the documents referenced in the order (or any addenda or amendment thereto), or the use or inability to use any Cloud Services, even if advised of the possibility of such damages.

44. 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.

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

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

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

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

- a. The Contractor shall implement and maintain appropriate administrative, technical and organizational security measures to safeguard against unauthorized access, disclosure or theft of Personal Data and Non-Public Data. Such security measures shall be in accordance with recognized industry practice and not less stringent than the measures the Contractor applies to its own Personal Data and Non-Public Data of similar kind.
- b. All data obtained by the Contractor in the performance of the Master Agreement shall become and remain the property of the Purchasing Entity.
- c. All Personal Data shall be encrypted at rest and in transit with controlled access. Unless otherwise stipulated, the Contractor is responsible for encryption of the Personal Data. Any stipulation of responsibilities will identify specific roles and responsibilities and shall be included in the service level agreement (SLA), or otherwise made a part of the Master Agreement.
- d. Unless otherwise stipulated, the Contractor shall encrypt all Non-Public Data at rest and in transit. The Purchasing Entity shall identify data it deems as Non-Public Data to the Contractor. The level of protection and encryption for all Non-Public Data shall be identified in the SLA.
- e. At no time shall any data or processes — that either belong to or are intended for the use of a Purchasing Entity or its officers, agents or employees — be copied, disclosed or retained by the Contractor or any party related to the Contractor for subsequent use in any transaction that does not include the Purchasing Entity.
- f. The Contractor shall not use any information collected in connection with the Services issued from this Master Agreement for any purpose other than fulfilling the Services.

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

4. Security Incident or Data Breach Notification:

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

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

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

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

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

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

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

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

7. Termination and Suspension of Service:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- a. The Contractor shall implement and maintain appropriate administrative, technical and organizational security measures to safeguard against unauthorized access, disclosure or theft of Personal Data and Non-Public Data. Such security measures shall be in accordance with recognized industry practice and not less stringent than the measures the Contractor applies to its own Personal Data and Non-Public Data of similar kind.
- b. All data obtained by the Contractor in the performance of the Master Agreement shall become and remain the property of the Purchasing Entity.
- c. All Personal Data shall be encrypted at rest and in transit with controlled access. Unless otherwise stipulated, the Contractor is responsible for encryption of the Personal Data. Any stipulation of responsibilities will identify specific roles and responsibilities and shall be included in the service level agreement (SLA), or otherwise made a part of the Master Agreement.
- d. Unless otherwise stipulated, the Contractor shall encrypt all Non-Public Data at rest and in transit. The Purchasing Entity shall identify data it deems as Non-Public Data to the Contractor. The level of protection and encryption for all Non-Public Data shall be identified in the SLA.
- e. At no time shall any data or processes — that either belong to or are intended for the use of a Purchasing Entity or its officers, agents or employees — be copied, disclosed or retained by the Contractor or any party related to the Contractor for subsequent use in any transaction that does not include the Purchasing Entity.
- f. The Contractor shall not use any information collected in connection with the Services issued from this Master Agreement for any purpose other than fulfilling the Services.

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

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

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

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

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

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

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

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

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

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

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

7. Termination and Suspension of Service:

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

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

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

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

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

8. Background Checks:

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

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

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

9. Access to Security Logs and Reports:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- a. The Contractor shall implement and maintain appropriate administrative, technical and organizational security measures to safeguard against unauthorized access, disclosure or theft of Personal Data and Non-Public Data. Such security measures shall be in accordance with recognized industry practice and not less stringent than the measures the Contractor applies to its own Personal Data and Non-Public Data of similar kind.
- b. All data obtained by the Contractor in the performance of the Master Agreement shall become and remain the property of the Purchasing Entity.
- c. All Personal Data shall be encrypted at rest and in transit with controlled access. Unless otherwise stipulated, the Contractor is responsible for encryption of the Personal Data. Any stipulation of responsibilities will identify specific roles and responsibilities and shall be included in the service level agreement (SLA), or otherwise made a part of the Master Agreement.
- d. Unless otherwise stipulated, the Contractor shall encrypt all Non-Public Data at rest and in transit. The Purchasing Entity shall identify data it deems as Non-Public Data to the Contractor. The level of protection and encryption for all Non-Public Data shall be identified in the SLA.
- e. At no time shall any data or processes — that either belong to or are intended for the use of a Purchasing Entity or its officers, agents or employees — be copied, disclosed or retained by the Contractor or any party related to the Contractor for subsequent use in any transaction that does not include the Purchasing Entity.
- f. The Contractor shall not use any information collected in connection with the Services issued from this Master Agreement for any purpose other than fulfilling the Services.

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

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

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

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

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

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

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

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

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

7. Termination and Suspension of Service:

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

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

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

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

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

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

8. Background Checks:

- a. Upon the request of the Purchasing Entity, the Contractor shall conduct criminal background checks and not utilize any staff, including subcontractors, to fulfill the obligations of the Master Agreement who have been convicted of any crime of dishonesty, including but not limited to criminal fraud, or otherwise convicted of any felony or misdemeanor offense for which incarceration for up to 1 year is an authorized penalty. The Contractor shall promote and maintain an awareness of the importance of securing the Purchasing Entity's information among the Contractor's employees and agents.
- b. The Contractor and the Purchasing Entity recognize that security responsibilities are shared. The Contractor is responsible for providing a secure infrastructure. The Purchasing Entity is responsible for its secure guest operating system, firewalls and other logs captured within the guest operating system. Specific shared responsibilities are identified within the SLA.
- c. If any of the stated personnel providing services under a Participating Addendum is not acceptable to the Purchasing Entity in its sole opinion as a result of the background or criminal history investigation, the Purchasing Entity, in its' sole option shall have the right to either (1) request immediate replacement of the person, or (2) immediately terminate the Participating Addendum and any related service agreement.

9. Access to Security Logs and Reports:

- a. The Contractor shall provide reports on a schedule specified in the SLA to the Contractor directly related to the infrastructure that the Contractor controls upon which the Purchasing Entity's account resides. Unless otherwise agreed to in the SLA, the Contractor shall provide the public jurisdiction a history or all API calls for the Purchasing Entity account that includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters and the response elements returned by the Contractor. The report will be sufficient to enable the Purchasing Entity to perform security analysis, resource change tracking and compliance auditing
- b. The Contractor and the Purchasing Entity recognize that security responsibilities are shared. The Contractor is responsible for providing a secure infrastructure. The Purchasing Entity is responsible for its secure guest operating system, firewalls and other logs captured within the guest operating system. Specific shared responsibilities are identified within the SLA.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Attachment B

Under the terms of the NASPO ValuePoint Cloud Solutions contract Insight is capable of providing the following service models and deployment models.

Service Models

AWS: IaaS, PaaS, SaaS

Microsoft: SaaS, IaaS, and PaaS

Deployment Models

AWS: Public (including Government Community Cloud) and Hybrid

Microsoft: Public (including Government Community Cloud), Hybrid, Private

Insight has identified the data risk categories that our Cloud Service Provider Partners are capable of storing and securing.

AWS: It is the responsibility of the customer to assign risk classification levels to their data. AWS' security features are outlined throughout our proposal response.

Service Model:	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered:
SaaS	X	X	X	Public, Hybrid, Private
IaaS	X	X	X	Public, Hybrid, Private
PaaS	X	X	X	Public, Hybrid, Private

Microsoft: Certain Microsoft Services are capable of handling and storing Moderate Data. This is addressed further in the exemption list.

Contract Service Offerings: Insight has partnered with two of our strongest Cloud Service Provider (CSP) partners in response to the NASPO ValuePoint/State of Utah's RFP. Amazon Web Services (AWS) is our first partner offering which has an expansive portfolio offering that will be made available and is outlined throughout Insight's proposal response. While the majority of AWS' offerings are classified as IaaS, some earn the classification of SaaS and PaaS. The classifications have been identified in our response. AWS' offerings can be delivered via Public and Hybrid deployment models.

To support AWS cloud solution purchases, we have also partnered with a third party services firm that specializes in AWS consulting services around design and deployment. REAN Cloud will assist Participating States and Entities in leveraging AWS' offerings to the fullest advantage possible.

The Insight REAN Cloud team is able to provide the following services:

Strategy Phase - SaaS	Assessment Phase - SaaS	Operations Phase - SaaS	Dev Ops Phase - PaaS
ROI & Business Case Justification (Activity) AWS Calculator (Task) Cloud Rationalization/Adoption strategy DR & Business continuity planning DevOps Strategy Account Management Governance & Compliance	Cloud Architecture Security & Risk Assessment Migration and Implementation Phase Secure Infrastructure Setup Lift & Shift Migration (CloudEndure) DevOps based migration	Managed Services (MGS) Billing as Service (BaaS) AWS Infrastructure (IaaS)	Infrastructure Automation Application Reengineering Native AWS Application Development

Insight's second cloud partnership is with **Microsoft**. Participating States and Entities will have access to IaaS, SaaS, and PaaS solution offerings delivered via Public (including the Government Community Cloud), Hybrid, and Private deployment models. Through the Microsoft partnership Office 365, Azure, Intune, and CRM Dynamics will be made available to the participating entities. Insight services will provide design and deployment capabilities for Office 365, Azure, and CRM Dynamics. Further description of the Online Services available is provided below.

Online Services in DPT – Those Online Services which store or process Customer Data, and are included in scope for the Data Processing Terms (DPT) section of the Microsoft Online Services Terms.	Online Services – Those online services which do not store or process Customer Data, and are merely desktop applications delivered using Microsoft's servers as a delivery mechanism.
Microsoft Dynamics CRM Online Services	Office 365 ProPlus
Office 365 Services	Project Pro for Office 365
Microsoft Azure Core Services	Visio Pro for Office 365
Microsoft Intune Online Services	

Insight will offer Participating States and Entities an array of cloud offerings from Microsoft, wrapped with support services and expert resources to centralize management and control for a diverse range of hosted solutions. Insight's team of cloud certified experts will draw on the experience of helping clients implement and manage a wide range of cloud solutions in their organizations. In the U. S. alone, Insight currently manages more than seven million seats distributed over 5,000 clients.

Insight's cloud solutions include the following:

Messaging Solutions	Security Solutions	Infrastructure Solutions	Collaboration Solutions
- Email Security - Hosted Exchange - Hosted Black Berry (BES) - Email Archiving - Email Continuity	- Web Security - Managed Firewall - Theft and Recovery Solutions - Intrusion Detection and Prevention - Vulnerability Management	- Online/ Remote Backup - Hosted VOIP & PBX Solutions - Desktop Management - Managed Co-location/ Hosting	- Instant Messaging - SharePoint Online - Web Conferencing - Hosted CRM

Attachment C – Cost Schedule

Solicitation Number CH16012 NASPO ValuePoint Cloud Solutions RFP

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

Software as a Service

Microsoft	Discount %: 10%
Amazon Web Services	Discount %: 1%

Infrastructure as a Service

Microsoft	Discount %: 10%
Amazon Web Services	Discount %: 1%

Platform as a Services

Microsoft	Discount %: 10%
Amazon Web Services	Discount %: 1%

Insight is offering Participating States and Entities minimum discounts off list price for all AWS and Microsoft cloud products purchased through Insight on the NASPO ValuePoint Cloud Solutions contract. These discounts apply to all product categories and subcategories outlined in our product catalogs submitted as separate attachments. Additional discounts may be extended for specific opportunities.

Value Added Services

Discount % _____

Please see Insight's Value Added Services Pricing in the sections below.

Additional Value Added Services:

Insight is offering value added services for the delivery of Microsoft based Cloud Solutions. These service types are clearly marked as *Insight Delivered*.

Maintenance Services (Insight Delivered)

Onsite Hourly Rate \$165/hr (Con I)
Remote Hourly Rate \$165/hr (Con I) (no travel expense)

Professional Services (Insight Delivered)

- **Deployment Services** **Onsite Hourly Rate \$165/hr (Con II)**
Remote Hourly Rate \$165/hr (no travel expense)

- **Consulting/Advisory Services** **Onsite Hourly Rate \$165/hr (Con II)**
Remote Hourly Rate \$165/hr (no travel expense)

- **Architectural Design Services** **Onsite Hourly Rate \$225/hr (Arch Sr)**
Remote Hourly Rate \$225/hr (no travel expense)

- **Statement of Work Services** **Onsite Hourly Rate \$165-\$225/hr**
Remote Hourly Rate \$165-\$225/hr (no travel)

Partner Services **Onsite Hourly Rate \$ please see below**
Remote Hourly Rate \$ please see below

Outlined below are the rates and costs that have been established for the services our subcontractor, REAN Cloud, will perform specific to AWS solutions. These rates apply to both onsite and remote services.

Labor Category (REAN Cloud)	Final Bid Price
Principal Technical architect	\$281.25
DevOps Architect	\$218.75
Sr. Cloud Engineer	\$187.50
Cloud Security Architect	\$218.75
Configuration Manager	\$156.25
Database Engineer	\$125.00
Developer -FE	\$106.25
PMO/Billing Specialist	\$118.75
Project Director	\$150.00
Project Manager SCRUM Master	\$112.50
Technical architect	\$181.25
Test Engineer	\$106.25
Test Lead	\$131.25

The proposed pricing for REAN Cloud delivered services is covered by the following rates for each of the talent categories required to provide the service.

REAN Cloud and Insight offer a one-time up front analysis, architecture, implementation, and migration. This offering is priced on a case-by-case basis. Initial setup and migration estimates are based on a blended price of license, support, engineering, architecture, project management, and AWS subject matter expertise that REAN/ Insight will bring to bear.

Managed Services (reoccurring monthly costs)

\$625/server/month for managed services fees (above AWS infrastructure costs) inclusive of all the help desk, security tools, management, patching, maintenance, monitoring and reporting as described in our sample SOW, for environments up to \$30,000/month in AWS infrastructure spend.

For environments that have more than \$30,000/month in AWS infrastructure spend, \$12,500 monthly flat fee plus 31.25% of AWS infrastructure spend for the month.

Example:

15 Servers in an AWS customer environment – Managed Services fees would be $625 \times 15 = \$9,375/\text{month}$

Example:

58 Servers in AWS Customer environment costing them about \$37,000/month. For Managed Services fees for this larger environment (over \$30,000/month in AWS spend) it would be \$12,500 plus 31.25% of \$37,000 or \$19,250.

Training Deployment Services (Insight Delivered)

Onsite Hourly Rate \$ 145/hr
Online Hourly Rate \$ <varies>

Insight offers in-person, remote, and third-party training for most of the technologies and transformative changes we deliver to our clients. Services range from technical training for technical staff to process training for end users. The cost of the training is variable to match the need, depth, and complexity of the training desired by our clients.



Insight Public Sector, Inc. Proposal Response

PREPARED FOR

The State of Utah

Request for Proposal # CH16012 for
NASPO ValuePoint Cloud Solutions
Technical Proposal

March 10, 2016 @ 1:00PM MST

SUBMITTED BY:

Insight Public Sector, Inc.

TABLE OF CONTENTS


1.	RFP SIGNATURE PAGE (RFP 5.1) (M)	1-1
2.	EXECUTIVE SUMMARY (RFP 5.4) (M).....	2-2
3.	MANDATORY MINIMUMS (RFP 5) (M)	3-4
	COVER LETTER (RFP 5.2) (M)	3-4
	ACKNOWLEDGE OF AMENDMENTS (RFP 5.3) (M)	3-6
	GENERAL REQUIREMENTS FOR THE SERVICE OFFERINGS (RFP 5.5) (M)	3-7
	RE-CERTIFICATION OF MANDATORY MINIMUMS AND TECHNICAL SPECIFICATIONS (RFP 5.7) (M)	3-8
4.	BUSINESS PROFILE (RFP 6)	4-9
	BUSINESS PROFILE (RFP 6.1) (M) (E)	4-9
	SCOPE OF EXPERIENCE (RFP 6.2) (M) (E)	4-11
	FINANCIALS (RFP 6.3) (M)	4-12
	GENERAL INFORMATION (RFP 6.4) (E)	4-13
	BILLING AND PRICING PRACTICES (RFP 6.5) (E)	4-21
	SCOPE AND VARIETY OF CLOUD SOLUTIONS (RFP 6.6) (E)	4-25
	BEST PRACTICES (RFP 6.7) (E)	4-27
5.	ORGANIZATION PROFILE (RFP 7) (M) (E)	5-34
6.	TECHNICAL RESPONSE (RFP 8) (M) (E)	6-36
7.	CONFIDENTIAL, PROTECTED, OR PROPRIETARY INFORMATION.....	7-149
8.	EXCEPTIONS AND/OR ADDITIONS TO THE STANDARD TERMS AND CONDITIONS	8-150

1. RFP Signature Page (RFP 5.1) (M)

Insight Response: Insight elected to submit our proposal response electronically, and therefore provided an electronic signature via BidSync. Below is a copy of the signed Vendor Information Form to confirm Insight's Authorized Representative signed off on Insight's proposal response.



State of Utah Vendor Information Form

Legal Company Name (include d/b/a if applicable) Insight Public Sector, Inc.		Federal Tax Identification Number 36-3949000		State of Utah Sales Tax ID Number NA	
Ordering Address 6820 S. Harl Avenue		City Tempe	State AZ	Zip Code 85283	
Remittance Address (if different from ordering address) PO Box 731072		City Dallas	State TX	Zip Code 75373-1072	
Type ☐ Proprietorship ☐ Partnership ☐ Government ☒ For-Profit Corporation ☐ Non-Profit Corporation		Company Contact Person Erica Falchetti			
Telephone Number (include area code) 480-333-3071		Fax Number (include area code) 480-760-9488			
Company's Internet Web Address ips.insight.com		Email Address Erica.Falchetti@Insight.com			
Offeror's Authorized Representative's Signature 					
Type or Print Name Brian Hicks					
Position or Title of Authorized Representative Vice President, Profitability					
Date: 3/3/16					

2. Executive Summary (RFP 5.4) (M)

Solicitation Objectives

Insight understands that the State of Utah (“UT”), Division of Purchasing, is in the process of selecting high quality cloud based service providers to serve states, territories, and their authorized political subdivisions. The providers should have the ability to provide a menu of cloud solutions offerings that will ultimately increase the technology department’s overall efficiency, reduce costs, improve operational scalability, provide business continuity, increase collaboration efficiencies, and allow for expanded flexibility in work practices and system improvements. The establishment of a cooperative contract provides a vehicle for authorized contract participants to obtain best value, and achieve more favorable pricing, than is obtainable by an individual state or local government entity because of the collective volume of potential purchases by numerous state and local government entities that is possible under a contract of this nature.

Insight Solutions

Insight would like to thank the State of UT Division of Purchasing and NASPO ValuePoint for the opportunity to submit the enclosed response for providing cloud solutions as described in the RFP, exhibits, and attachments.

NASPO ValuePoint can benefit from a continuing partnership with Insight because our IT solutions are designed with our public sector clients in mind. Our process knowledge, product fulfillment and logistics capabilities along with our management tools, and expertise make managing IT solutions easier while helping Participating States and Entities control their IT costs. Based on the documents in the solicitation and description of requested services, we are prepared to offer NASPO ValuePoint the following:

Contract Service Offerings: Insight has partnered with two of our strongest Cloud Service Provider (CSP) partners in response to the NASPO ValuePoint/State of Utah’s RFP. Amazon Web Services (AWS) is our first partner offering which has an expansive portfolio offering that will be made available and is outlined throughout Insight’s proposal response. While the majority of AWS’ offerings are classified as IaaS, some earn the classification of SaaS and PaaS. The classifications have been identified in our response. AWS’ offerings can be delivered via Public and Hybrid deployment models.

To support AWS cloud solution purchases, we have also partnered with a third party services firm that specializes in AWS consulting services around design and deployment. REAN Cloud will assist Participating States and Entities in leveraging AWS’ offerings to the fullest advantage possible. The Insight REAN Cloud team is able to provide the following services:

Strategy Phase - SaaS	Assessment Phase - SaaS	Operations Phase - SaaS	DevOps Phase - PaaS
ROI & Business Case Justification (Activity) AWS Calculator (Task) Cloud Rationalization/Adoption strategy DR & Business continuity planning DevOps Strategy Account Management Governance & Compliance	Cloud Architecture Security & Risk Assessment Migration and Implementation Phase Secure Infrastructure Setup Lift & Shift Migration (CloudEndure) DevOps based migration	Managed Services (MGS) Billing as Service (BaaS) AWS Infrastructure (IaaS)	Infrastructure Automation Application Reengineering Native AWS Application Development

Insight's second cloud partnership is with **Microsoft**. Participating States and Entities will have access to IaaS, SaaS, and PaaS solution offerings delivered via Public (including the Government Community Cloud), Hybrid, and Private deployment models. Through the Microsoft partnership Office 365, Azure, Intune, and CRM Dynamics will be made available to the participating entities. Insight services will provide design and deployment capabilities for Office 365, Azure, and CRM Dynamics. Further description of the Online Services available is provided below.

Online Services in DPT – Those Online Services which store or process Customer Data, and are included in scope for the Data Processing Terms (DPT) section of the Microsoft Online Services Terms.	Online Services – Those online services which do not store or process Customer Data, and are merely desktop applications delivered using Microsoft's servers as a delivery mechanism.
Microsoft Dynamics CRM Online Services	Office 365 ProPlus
Office 365 Services	Project Pro for Office 365
Microsoft Azure Core Services	Visio Pro for Office 365
Microsoft Intune Online Services	

Insight will offer Participating States and Entities an array of cloud offerings from Microsoft, wrapped with support services and expert resources to centralize management and control for a diverse range of hosted solutions. Insight's team of cloud certified experts will draw on the experience of helping clients implement and manage a wide range of cloud solutions in their organizations. In the U.S. alone, Insight currently manages more than seven million seats distributed over 5,000 clients.

Insight's cloud solutions include the following:

Messaging Solutions	Security Solutions	Infrastructure Solutions	Collaboration Solutions
• Email Security • Hosted Exchange • Hosted BlackBerry (BES) • Email Archiving • Email Continuity	• Web Security • Managed Firewall • Theft and Recovery Solutions • Intrusion Detection and Prevention • Vulnerability Management	• Online/Remote Backup • Hosted VOIP & PBX Solutions • Desktop Management • Managed Co-location/Hosting	• Instant Messaging • SharePoint Online • Web Conferencing • Hosted CRM

Conclusion

It is our belief that the requirements outlined in the SOW and the information Insight has provided in our response make a compelling proposition for NASPO ValuePoint to select Insight to participate in this contract. Insight has continuously evolved and grown as the IT industry has changed. We provide significant value in IT procurement and management assistance to state and local governments and educational entities.

The cloud solutions presented throughout our response give evidence to our commitment and our ability to ensure Participating States and Purchasing Entities get the most value out of their cloud technology investments-while decreasing their Total Cost of Ownership. We have proven our competence to our clients, and believe we can exceed your expectations.

3. Mandatory Minimums (RFP 5) (M)

Cover Letter (RFP 5.2) (M)

March 10, 2016

Christopher Hughes
Contracts Analyst
DAS
T: 801-538-3254
E: Christopherhughes@utah.gov

RE: RFP # CH16012 for NASPO ValuePoint Cloud Solutions

Dear Mr. Hughes:

Insight Public Sector, Inc. ("Insight") is pleased to participate in the State of Utah's RFP for Cloud Solutions in furtherance of the NASPO ValuePoint Cooperative Purchasing Program. Based on the scope of the requirements, Insight has prepared a response that represents a comprehensive effort at meeting the requirements of the RFP to provide services related to cloud solutions for all Participating States and Entities.

Insight Public Sector is solely focused on the needs of local, state and federal governments as well as educational institutions. With an industry-leading selection of products, a complete suite of IT services and a wide range of government contracts, Insight helps organizations streamline procurement, simplify deployment and maximize the value of the IT lifecycle.

We have provided below all of the information required for the cover letter.

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

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

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

Insight Response: The following individuals and/or firms were responsible for contributing to the content of Insight's response.

Insight Public Sector Staff		Cloud Service Providers (CSP)/ Subcontractor Firms
Joanna Crowder Erica Falchetti David Solliday Billy Roberts	Jeromy Siebenaler	Amazon Web Services
	Heather Suchobrus	Microsoft
	Joe Benik Joe Monforton	REAN Cloud

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

Insight Response: Neither Insight nor any key personnel are currently debarred pursuant to an established debarment procedure from bidding or contracting by any public body of this or any other state or agency of the federal government.

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

Insight Response: Insight Public Sector acknowledges that a 0.25% NASPO ValuePoint Administration Fee and any Participating Entity Administrative Fee will apply to total sales for the Master Agreement(s) awarded from the RFP.

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

Insight Response: Under the terms of the NASPO ValuePoint Cloud Solutions contract Insight is capable of providing the following service models and deployment models.

Service Models

AWS: IaaS, PaaS, SaaS

Microsoft: SaaS, IaaS, and PaaS

Deployment Models

AWS: Public (including Government Community Cloud) and Hybrid

Microsoft: Public (including Government Community Cloud), Hybrid, Private

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

Insight Response: Insight has identified the data risk categories that our Cloud Service Provider Partners are capable of storing and securing.

AWS: It is the responsibility of the customer to assign risk classification levels to their data. AWS' security features are outlined throughout our proposal response.

Service Model:	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered:
SaaS	X	X	X	Public, Hybrid, Private
IaaS	X	X	X	Public, Hybrid, Private
PaaS	X	X	X	Public, Hybrid, Private

Microsoft: Certain Microsoft Services are capable of handling and storing Moderate Data. This is addressed further in the exemption list.

Insight Public Sector looks forward to working with the NASPO ValuePoint organization and the State of Utah.

Respectfully,



John Carnahan
Senior Vice President, Operations

Acknowledge of Amendments (RFP 5.3) (M)

If the RFP is amended, the Offeror must acknowledge each amendment with a signature on the acknowledgement form provided with each amendment. Failure to return a signed copy of each amendment acknowledgement form with the proposal may result in the proposal being found non-responsive. Note: Offeror will not need to sign an amendment for a Master Agreement update. A Master Agreement update should be used when the action on the contract is objective and provides factual updates. Examples of when an update should be used in lieu of an Amendment include technical clarifications that do not change the SOW or Ts & Cs, e.g. changes of address, phone number, contact person, etc.

Insight Response: Insight understands and has complied with this requirement. Provided as an attachment – *RFP CH16012_Acknowledgement of Amendments_Signed_Insight*- are the following signed Acknowledgement of Amendment forms.

1. Acknowledgement of Amendments to RFP: February 3, 2016
2. Acknowledgement of Amendments to RFP: February 10, 2016

General Requirements for the Service Offerings (RFP 5.5) (M)

5.5.1 If awarded a contract Insight will provide a Usage Report Administrator responsible for the quarterly sales reporting described in the Master Agreement Terms and Conditions. The Usage Administrator will use Attachment F as the template to report usage under the contract.

Insight Response: If awarded a contract Insight will provide a Usage Report Administrator responsible for quarterly sales reporting described in the Master Agreement Terms and Conditions. The Usage Administrator will use Attachment F of the RFP as the template to report all usage under the contract.

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

Insight Response: Insight will comply with this requirement. We will cooperate with NASPO ValuePoint and SciQuest by providing ordering instructions that will provide information on how to order directly from the Contractor outside of the eMarket Center, as well as provide information about the Contractor.

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

Insight Response: As a Value Added Reseller, this requirement does not apply to Insight, but does apply to our Cloud Service Provider partners, AWS and Microsoft. Responses for each CSP is provided in Section 8.13.

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

Insight Response: Insight understands this requirement and has provided a detailed answer in Section 8.10 of the proposal response.

Re-Certification of Mandatory Minimums and Technical Specifications (RFP 5.7) (M)

Offeror must acknowledge that if it is awarded a contract under the RFP that it will annually certify to the Lead State that it still meets or exceeds the mandatory minimum requirements and technical specifications of the RFP.

Insight Response: Insight acknowledges that if the firm is awarded a contract under the RFP that it will be required to annually certify to the Lead State that it still meets or exceeds the mandatory minimum requirements and technical specifications of the RFP.

4. Business Profile (RFP 6)

Business Profile (RFP 6.1) (M) (E)

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

Insight Response: Insight Enterprises, Inc. (our parent company), founded in 1988, is a leading technology provider of hardware, software and service solutions to commercial and public sector customers in the United States, Canada, Europe, and Asia-Pacific. Our company management structure is broken down by North America, EMEA and APAC and consists of Management, Support Services, Administration, Sales Account Executives and Distribution employees. The highest position in each of these areas is a Senior Vice President who reports either directly into our Chief Executive Officer or into one of our other Executive Officers. Insight Enterprises, Inc. became a publicly traded company in 1995, selling its stock on the NASDAQ under the ticker symbol NSIT. Insight Enterprises, Inc. is ranked number 493 on Fortune Magazine's 2015 "Fortune 500" list.

Insight has 206 SLED and Healthcare sales, support and management teammates located throughout the United States. Insight has 50 office and remote locations, as well as 40 home-based offices. In addition, our clients are supported with a national services team of 1,374 staff members, for a combined total of 1,580 sales and technical resources.

Insight Public Sector (Insight) holds over 180 federal, state, local, education and non-profit contracts. Insight currently maintains federal contracts with agencies such as the General Services Administration, and national contracts like U.S. Communities. In addition, our participation in 25 state wide contracts gives us a solid market share of government technology sales. Insight also holds local government and education contracts for computer equipment and services in 33 states. Highly specialized teams are dedicated to each market offering customized solutions that range from initial consulting, procurement and product delivery to maintenance and support.

The combined Insight companies and their subsidiaries represent a \$5.4 billion, in 2015, global enterprise. While remaining small enough to service our public sector clients with personal attention, Insight Public Sector has the resources of the Insight family of companies behind us to support our efforts. Insight's Public Sector business has steadily grown. In 2015, our Public Sector business (Fed/State & Local/Education) grew more than 30%. Cumulatively, as a company, Insight's revenue has remained steady, and in most cases has grown year-over-year (YOY), as demonstrated below.

2015: \$5.4 Billion	2014: \$5.3 Billion	2013 :\$5.1 Billion	2012: \$5.3 Billion	2011: \$5.3 Billion
---------------------	---------------------	---------------------	---------------------	---------------------

Insight's average worldwide headcount over the past three (3) years is as follows:

- 2015 – As of December 31, 2015, we employed 5,761 teammates
- 2014 - As of December 31, 2014, we employed 5,406 teammates.

- 2013 - As of December 31, 2013, we employed 5,202 teammates.

Per the requirements of the RFP, we have provided Insight Public Sector teammate retention statistics for the year 2015.

- 2015: 77.99% Retention Rate

The organization chart presented below outlines Insight's organizational structure and identifies the various Insight teams that will be responsible for ensuring successful delivery of services and providing oversight on the contract.




Figure 1: Insight Public Sector NASPO Cloud Solutions Organization Chart

Experience

Clients looking for a partner to connect them with cloud solutions would be hard-pressed to find a more qualified partner than Insight. With a decade of experience in providing cloud solutions, Insight's own portfolio of best-in-class partners offers NASPO Value Point Purchasing Entities diverse advantages to meet their comprehensive needs. It's why Insight is one of the few organizations anywhere that can truly act as a one-stop source for cloud – from consultative advice early in the process, to delivery and go-live, all the way through management.

In the U.S. alone, Insight currently manages more than 7 million seats distributed over 5,000 clients. It's all possible as a result of having a robust roster of best-in-class solutions and partners – a roster that's consistently being re-assessed for consistency in quality and relevance. In 2015, Insight delivered almost \$20 million of cloud services to nearly 1,000 clients across the commercial and public sector markets including but not limited to: healthcare, state/local government, K-12 education, higher education, federally funded financial institutions, banking and financial services, manufacturing, retail, pharmaceutical, and hospitality.

Independent School District: Due to a lack of internal resources, the client needed a solution that was cost efficient and required little maintenance. Insight helped identify a solution which was easier to manage than the School's PBX on-prem system, provided greater capabilities, and improved communication. Insight designed and installed a Cloud VOIP solution for 750 end users.

Retail Chain: After supporting this national chain's Microsoft Office 365 transition, Insight was quickly engaged to help develop a strategy for the company's current security strategy, particularly as it pertained to their archiving needs. Within two weeks, Insight was able to help the client introduce a cloud-based security/archiving solution into their environment.

American Cable and Satellite News Channel: With an accelerated timeline for the migration of 3,500 users to Microsoft Office 365 while maintaining a limited impact to user's access, Insight was engaged to design a migration solution, including ensuring the client's infrastructure was properly set up in order to support the migration. Insight created a High Availability solution that allowed the client to continue to leverage the features of a hybrid solution with minimal to no downtime.

Scope of Experience (RFP 6.2) (M) (E)

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

Insight Response: Insight has over ten years of experience with major contract implementation efforts. We are an existing NASPO ValuePoint Software VAR contract holder and as such have previously gone through the process to implement a new NASPO ValuePoint contract. And given our staff's familiarity with the contract organization, we expect a seamless implementation of a new contract. In addition to our NASPO ValuePoint experience, Insight holds and manages some of the largest SLED contracts in the country, including state government contracts in over 20 states, a U.S. Communities contract, and a GSA Schedule. Through these experiences, we have developed a robust contract implementation and management program based on accepted project management standards (ref. Project Management Body of Knowledge, PMBOK Guide).

This methodology is widely accepted as the best way to ensure requirements are determined, dependencies are investigated and planned for, and timelines are acknowledged and agreed to by all parties involved in the project. This process helps ensure deliverables meet or exceed expectations. As with any project there is the need to regularly evaluate progress and make adjustments in order to successfully complete the project. The project team has a responsibility to complete all of the process steps required, to adhere to the methodology outlined, and complete the project deliverables on time.

Over the past ten years Insight's Cloud Services team has focused on developing and bringing to market robust Cloud Services offerings to help our clients in the public and commercial markets take advantage of the benefits of cloud computing. Because a high percentage of public sector entities are in the investigative phase of adopting Cloud technology solutions, the opportunity to participate in a cloud services focused contract(s) has been minimal. However, through the large, broad scoped IT products and services contracts Insight holds, we have been able to introduce cloud computing solutions to customers purchasing under those contracts.

Provided in the table below are the five (5) contracts through which Insight has introduced the most cloud solutions.

Contract Name	Contract Value (Annual)	Contract Term
U.S. Communities	\$100M+	May 2009 - Current
NASPO ValuePoint Software VAR	\$60M+	June 2011- Current
State of California SCA for Microsoft Enterprise	\$18M+	January 2012 – Current
County of Riverside for Microsoft Enterprise	\$10M+	November 2011 - Current
State of Florida ACS Agreement	\$7M+	May 2009 - Current

Financials (RFP 6.3) (M)

Offeror must provide audited financial statements to the State and should meet a minimum Dun and Bradstreet (D&B) credit rating of 4A2 or better, or a recognized equivalent rating. Please provide the Respondent's D&B Number and the composite credit rating. The State reserves the right to verify this information. If a branch or wholly owned subsidiary is bidding on this RFP, please provide the D&B Number and score for the parent company that will be financially responsible for performance of the agreement.

Insight Response: Insight Public Sector's Dun and Bradstreet (D&B) number is 88-434-7568. The D&B rating for Insight Public Sector is a 1R3, which means IPS is a subsidiary whose financial information rolls up to our parent company, Insight Enterprises, Inc. Insight Enterprises, Inc. has a D&B rating of 5A2, which is the best/highest rating that can be obtained from D&B.

The State of Utah State Procurement Office and Purchasing States can access electronic versions of our financial information, past annual reports, as well as other audited financial statements via the link below to our website. <http://nsit.client.shareholder.com/financials.cfm>

General Information (RFP 6.4) (E)

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

Insight Response: Insight offers our clients an array of cloud offerings from industry leaders, wrapped with support services and expert resources to centralize management and control for a diverse range of hosted solutions. Insight's team of cloud certified experts draws on the experience of helping clients implement and manage a wide range of cloud solutions in their organizations. The result is more choice and more control of their cloud computing initiative. In the U.S. alone, Insight currently manages more than seven million seats distributed over 5,000 clients.

Insight's cloud solutions include the following:

Messaging Solutions	Security Solutions	Infrastructure Solutions	Collaboration Solutions
• Email Security • Hosted Exchange • Hosted BlackBerry (BES) • Email Archiving • Email Continuity	• Web Security • Managed Firewall • Theft and Recovery Solutions • Intrusion Detection and Prevention • Vulnerability Management	• Online/Remote Backup • Hosted VOIP & PBX Solutions • Desktop Management • Managed Co-location/Hosting	• Instant Messaging • SharePoint Online • Web Conferencing • Hosted CRM

Cloud Computing is a growing trend within the technology industry. Insight's Software-as-a-Service (SaaS)/Cloud Services solutions are designed to provide 24x7 assistance with offerings such as: hosted Exchange email, backup, and recovery for designated critical business data; email security that helps control e-mail threats or to gain greater control; and management of all the PCs on a network.

Insight cloud solutions are designed for affordability and efficiency to provide:

- Minimal to no hardware or software investment required
- No additional IT staff or increased management time necessary to maintain the services
- Easy to implement, easy to use, easy to manage
- Quick on-boarding of new users
- Compatible with multiple platforms
- 24/7 support
- Less software management; updates and upgrades are automatic
- Enterprise-class features at low per-user costs
- Full virus and spam protection

Insight also offers additional services to provide a turn-key onboarding experience migrating information and data into the cloud. Insight offers pre-sales assessments, onboarding project planning and project management, coupled with post-sales data migration and integration services.

As we look forward, Insight is well positioned for continued success. Insight will continue to invest in and apply best practices to our business model, incorporate strategies, and streamline processes. Our clients benefit from being able to obtain all the products, services, and expertise

they require—all from a single, reliable source. Insight's success and growth over the past 27 years demonstrates our commitment to the significant investments required to develop an IT and operating infrastructure. It is Insight's strategy to continue utilizing this business model to support our current and future client base.

Forrester Reports: In a December 2015 report from Forrester Research, Inc., Insight was positioned as one of the "thought leaders" interviewed for the report entitled, "The State of the Cloud: Migration, Portability, and Interoperability." Forrester interviewed over 25 enterprise and vendor thought leaders to discuss the state of migration, portability, and interoperability in the cloud. Insight was listed as a resource multiple places in the report as well as serving as a thought leader. The full report can be found at http://www.insight.com/en_US/learn/whitepapers/forrester/state-of-the-cloud.html




Figure 2: 2014 Forrester Report Value-Added Reseller Findings

According to a 2014 report from Forrester Research, Inc., Insight was recognized for offering the most comprehensive capabilities among 14 leading value-added resellers (VARs) based on a comprehensive study of these VARs in the global marketplace. The VARs were evaluated on the services they provide to clients, the regions in which they serve, and their partnerships with original equipment manufacturers (OEMs). Forrester's report summarizes their perspective on the shifting market, taking an in-depth look at how the VARs continue to adapt to a highly volatile landscape, and how sourcing and vendor management professionals can best take advantage of these changes.

REAN Cloud: Insight has partnered with REAN Cloud ("REAN"), a Premier Consulting Partner of the Amazon Web Services ("AWS") Partner Network, to provide full-service Cloud IT solutions to Participating Entities through the NASPO ValuePoint Cloud Solutions contract. REAN Cloud is a cloud-native firm with deep experience supporting legacy enterprise IT infrastructures and applications. The REAN Cloud team has served Enterprise IT clients to successful cloud adoption, specializing in helping enterprises use the cloud to become agile, realize cost savings, and enhance security and performance.



REAN was the fastest to achieve elite Premier Consulting Status (one amongst the 46 companies worldwide out of 22,000 + AWS partners) and is being recognized by an independent editors' consortium as one of the top 7 AWS consulting partners worldwide.

REAN Cloud's background and experience is in providing a Secure, Compliant Systems Architecture framework for clients in highly regulated industries. Their experience supporting financial services, healthcare, education, and government clients has been a market driver for them to build security and compliance into their offerings from the outset.

REAN Cloud provides Consulting Services around

- Cloud Strategy,
- Cloud Systems Architecture, Migration, Custom Cloud-Based Solutions, DevOps and
- Managed Cloud Services (MCS).

REAN offers a Secure Managed Services framework which handles all end-user requirements in the so-called Shared Responsibility Model. REAN's Managed Services enables clients to confidently migrate their workloads to the cloud.

AWS: Insight is partnering with AWS because of their clear differentiators. Below are some of the features and benefits of AWS that set their cloud infrastructure services apart in the marketplace.

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



Figure 3: Magic Quadrant for Cloud IaaS, Worldwide

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

Partner and Software Ecosystem: According to the 2015 Gartner, Inc. report referenced above, AWS has attracted "a very large technology partner ecosystem that includes software vendors that have licensed and packaged their software to run on AWS, as well as many vendors that have integrated their software with AWS capabilities. It also has an extensive network of

partners that provide application development expertise, managed services, and professional services such as data center migration.”

AWS Cloud Security Authorizations and Experience: AWS offers customers a powerful cloud security capability based on cutting-edge security experience and backed by an extensive repertoire of accreditations and authorizations. In The Forrester Wave™: Public Cloud Platform Service Providers' Security, Q4 2014 report, Forrester Research named AWS as the only provider in the Leader category. Forrester stated, "AWS leads the pack. AWS demonstrated not only a broad set of security capabilities in data center security, certifications, and network security, but also excelled in customer satisfaction, security services partnerships, and a large installed base.”

AWS Pricing: As AWS's cloud computing infrastructure grows, it gains efficiency and economies of scale, which Amazon passes on to their customers in the form of lower prices. The 2015 Gartner, Inc. report referenced above states that AWS has “over 10 times more cloud IaaS compute capacity in use than the aggregate total of the other 14 providers,” demonstrating how AWS's massive economies of scale make it possible to lead the cloud market in lowering prices.

Business Benefits of AWS Cloud Services

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

- **Almost Zero Upfront Infrastructure Investment:** AWS allows customers to access a large-scale system without having to invest in the real estate, physical security, hardware (racks, servers, routers, backup power supplies), hardware management (power management, cooling), and operations personnel.
- **Just-In-Time Infrastructure:** By deploying applications in the AWS cloud with just-in-time self-provisioning, customers do not have to worry about pre-procuring capacity for large-scale systems. AWS's cloud increases agility, lowers risk, and lowers operational cost, because customers can scale cloud resources as they grow and only pay for what they use.
- **More Efficient Resource Utilization:** With AWS, System Administrators can manage resources more effectively and efficiently by having the applications request and relinquish resources on-demand.
- **Usage-Based Costing:** With utility-style pricing, AWS customers are billed only for the infrastructure that has been used. AWS customers do not pay for allocated but unused infrastructure.
- **Reduced Time to Market:** Parallelization is the one of the great ways to speed up processing. If one compute-intensive or data-intensive job that can be run in parallel takes 500 hours to process on one machine, with cloud architectures, it would be possible to spawn and launch 500 instances and process the same job in 1 hour. Having available an elastic infrastructure provides the application with the ability to exploit parallelization in a cost-effective manner reducing time to market.

Analyst Reports: Gartner, Inc., a leading information technology research company, reported in its 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS), Worldwide report that “AWS is a thought leader; it is extraordinarily innovative, exceptionally agile, and very responsive to the market. It has the richest array of IaaS features and PaaS-like capabilities. It continues to rapidly expand its service offerings and offer higher-level solutions.” The Gartner Magic Quadrant for May 2015 (**Figure 3**) depicts AWS in the Leaders Quadrant.




Figure 5: Forrester Wave: Public Cloud Platform Service Providers' Security

in data center security, certifications, and network security, but also excelled in customer satisfaction, security services partnerships, and a large installed base. AWS led with the size of its development and technical support staff as well."

Microsoft: From virtual machines and storage to media services, a broad range of Azure services are available in Azure Government. Their commitment to innovation goes beyond basic services, extending to government-specific solutions.

The Microsoft Cloud for Government allows state and local governments to select the best tools to solve unique problems, whether it be for a large agency or small town government. Plus, Microsoft's massive global investment in data centers, and dedicated to U.S. federal and state policies, mandates, and compliance means they have everyone covered.



Figure 4: Magic Quadrant for Public Cloud Storage Services

Additionally, Gartner positions AWS in the Leaders Quadrant of the new Magic Quadrant for Public Cloud Storage Services (Figure 4). Gartner defines leaders as offering innovative storage offerings built on a hardened platform, with global data centers and established credibility as a business.

The Forrester Wave: Public Cloud Platform Service Providers' Security, Q4 2014 report evaluated four of the leading public clouds along 15 key security criteria. Forrester's evaluation states "AWS leads the pack. AWS demonstrated not only a broad set of security capabilities

Offerings in the Government Cloud

Microsoft Office 365 U.S. Government provides organizations with easy-to-use productivity and collaboration tools that allow them to spend more time serving their community and less time sifting through paperwork. The secure and compliant platform lets departments seamlessly work together from anywhere with an Internet connection on nearly any device.

Microsoft Dynamics CRM Online Government is the solution that equips organization's employees with data, along with reporting, modeling, and powerful workflows, while also offering security features that can limit access to sensitive data. Dynamics can also free data trapped in outdated systems and automate monotonous tasks, allowing employees to focus on more important work.

Microsoft Azure Government increases the agility of federal, state, and local government organizations and partners with hyperscale computing, storage, networking, and identity management services. Its integration of on-premise apps and data with cloud computing breaks down office walls, allowing governments to collaborate with citizens in their communities.

Azure Government is a *government-community cloud (GCC)* designed to support strategic government scenarios that require speed, scale, security, compliance and economics for U.S. government organizations. It was developed based on Microsoft's extensive experience delivering software, security, compliance, and controls in other Microsoft cloud offerings such as Azure public, Office 365, O365 GCC, Microsoft CRM Online etc.

In addition, Azure Government is designed to meet the higher level security and compliance needs for sensitive, dedicated, U.S. Public Sector workloads found in regulations such as United States Federal Risk and Authorization Management Program (FedRAMP), Department of Defense Enterprise Cloud Service Broker (ECSB), Criminal Justice Information Services (CJIS) Security Policy and Health Insurance Portability and Accountability Act (HIPAA).

Industry Recognition

Gartner has positioned Microsoft in the Leaders Quadrant in the 2015 Magic Quadrant for Cloud Infrastructure as a Service (IaaS) based on its completeness of vision and ability to execute in the IaaS market. Gartner defines cloud IaaS as a standardized, highly-automated offering where compute resources, complemented by storage and networking capabilities, are owned by a service provider and offered to the customer on-demand. Microsoft is currently the only vendor to be positioned as a Leader in Gartner's Magic Quadrants for Cloud Infrastructure as a Service, Application Platform as a Service, Cloud Storage Services and Server Virtualization, and they believe this validates Microsoft's strategy to enable the power of choice as

Figure 1. Magic Quadrant for Cloud Infrastructure as a Service, Worldwide



Figure 6: Gartner Magic Quadrant for Cloud IaaS, Worldwide- 2015

they deliver industry-leading infrastructure services, platform services and hybrid solutions.



Figure 7: Gartner Magic Quadrant for Enterprise Application Platform as a Service, Worldwide

Microsoft is currently the only vendor to be positioned as a Leader in Gartner's **Magic Quadrants for Cloud Infrastructure as a Service, Application Platform as a Service, Cloud Storage Services** and **Server Virtualization**. Their strategy is driving significant usage and growth for Azure with more than 90,000 new Azure customer subscriptions every month and over 57% of Fortune 500 companies using Azure.



Figure 8: Gartner Magic Quadrant for Public Cloud Storage Services

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

Insight Response: Due to Insight's business status as a Value Added Reseller this requirement is not applicable to Insight. However, it does apply to our CSP partners that we are proposing and we have provided their compliance below.

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

Microsoft: For Data Processing Term (DPT) Services (each, as defined in the Microsoft contract terms attached with Insight's Proposal), each such DPT Service follows a written data security policy ("Information Security Policy") that complies with the control standards and frameworks of SSAE 16 SOC 1 (Type II) and SOC 2 (Type II). Additional standards are listed in the Microsoft Online Services Terms, which also include terms and conditions pursuant to which the audit findings may be provided to Purchasing Entities under non-disclosure agreement.

Billing and Pricing Practices (RFP 6.5) (E)

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

Insight Response: Insight will provide price lists for both Microsoft and AWS showing both the provider's MSRP as well as the contract price. These price lists will be available on our dedicated NASPO ValuePoint Cloud Solutions website and updated as often as we receive updates from the provider. Participating Entities can always be assured that the price offered to them will be at or below the not-to-exceed prices listed on the price lists. Additionally, Insight has provided details pertaining to the ordering process for Amazon Web Services and Microsoft.

Amazon Web Services

Insight quotes AWS using AWS's Simple Monthly Calculator to generate an appropriate solution. The results from the generator are then added to our Customer Facing Agreement (CFA). The Insight CFA contains several key parts:

- The Insight Public Sector remit to address
- Insight's Terms of Sale
- AWS's Terms of Service
- AWS's Terms of Use
- Customer's Legal signature and optional PO field

With this order form we should have enough customer information to set up a working AWS solution and allow the customer to be operating in the AWS environment quickly.

Figure 9: AWS Simple Monthly Calculator

We have included a copy of the Customer Facing Agreement as an attachment to our response.

Microsoft

Provided below is an explanation of the ordering process for the Microsoft services Insight is offering in our proposal response.

For Office 365, Exchange Online, and CRM Online, the customer initially makes a three-year commitment for a quantity of products that is paid for annually. This is referred to as their annual payment. If the customer wishes to add additional quantities, they can pay a pro-rated

amount to the anniversary date and then that quantity is added to their annual payment. If the customer wishes to reduce their quantity, they can do this on the anniversary date only.

For Azure, the customer has two options to procure—upfront or in arrears. With the upfront option, the customer commits to a three-year agreement for an annual dollar amount of Azure services they want to consume. This is referred to as their annual payment. The customer can add more commitment dollars at any time during the year for additional services. If the customer exhausts their commitment dollars before the anniversary date, they will be billed quarterly in arrears for the services they use. The customer can make adjustments to their annual commitment upward or downward at their anniversary date. With the pay in arrears option, the customer commits no upfront dollars. They are simply billed quarterly in arrears for the Azure services they use.

Provided below is information pertaining to AWS and Microsoft pricing and billing practices.

AWS: With AWS, customers can incorporate a utility-style pricing model, only paying for the resources consumed. AWS continues to lower the cost of cloud computing for its customers. In 2014, AWS reduced the cost of compute by an average of 30%, storage by an average of 51%, relational databases by an average of 28%, and they continue to drive down the cost of customer IT infrastructure. AWS's utility-style pricing model is explained below to provide NASPO ValuePoint and the State of Utah with further understanding of how AWS services are charged.

- **Pay as You Go:** No minimum commitment or long-term contract is required. Customers can turn off cloud resources and stop paying for them when they are not needed, maximizing Return on Investment (ROI) through full utilization.
- **Pay Less When You Reserve:** For certain AWS products, customers can invest in reserved capacity, paying a low up-front fee to receive a significant discount. This results in overall savings of up to 60% (depending on the type of instance reserved) over equivalent on-demand capacity.
- **Pay Even Less Per Unit by Using More:** AWS pricing is tiered for storage and data transfer, so the more customers use, the less they pay per gigabyte.
- **Pay Even Less as AWS Grows:** Amazon continually focuses on reducing their data center hardware costs, improving our operational efficiencies, lowering their power consumption, and passing savings back to customers. AWS has a history of continually lowering prices and has reduced prices 51 times since AWS launched in 2006.
- **Transparency:** AWS provides transparent, publicly available, and up-to-date pricing, as well as tools that allow customers to evaluate AWS pricing against other CSPs. AWS's Simple Monthly Calculator is available online.

Potential Business Value of Running Applications on AWS

- Five year ROI: 560%
- Payback period: 5.5 months
- \$1.54M average five-year discounted business benefits per application
- 64.3% lower TCO
- 68.1% more efficient IT staff operations
- \$76,800 additional revenue per year per application
- 118.4% more applications delivered
- 81.7% less downtime

Source: IDC Whitepaper, sponsored by AWS, "Quantifying the Business Value of Amazon Web Services," May 2015.

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

The AWS Total Cost of Ownership (TCO) Calculator allows organizations to compare AWS to the cost of running applications in an on-premises or traditional hosting environment.

Microsoft: Microsoft pricing for government customers is updated monthly on the first of every month. Prices don't always change, nor do they always go up, but the changes can be significant, especially when new products and services are released. Insight will keep the Purchasing Entities updated throughout the RFP process, and if selected, through the term of the engagement. In order to do this, Insight will work with each Purchasing Entity to ensure that the updated pricing is passed effectively.

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

Insight Response: This requirement does not apply to Insight because we are not the host of the solution's infrastructure.

AWS: Per the response provided in Section 6.4.1., clients make an almost zero upfront infrastructure investment when they choose Amazon Web Services. AWS allows customers to access a large-scale system without having to invest in the real estate, physical security, hardware (racks, servers, routers, backup power supplies), hardware management (power management, cooling), and operations personnel.

Microsoft: Listed below are costs that a Purchasing Entity might need to consider when implementing a Microsoft cloud solution. In addition to implementation costs, included in the list are common costs associated with the solution post-implementation.

Office 365 subscription costs: These are per-user per-month costs for users subscribing to Microsoft business productivity services available through the Office 365 program.

Azure Compute costs: These are per-minute charges for virtual machines running in the Azure Cloud.

Azure Storage costs: These are per-GB of storage for data stored in the Azure Cloud.

Azure ExpressRoute: This is an option that provides the Purchasing Entity with a private, dedicated connection to the Microsoft Cloud, instead of connecting over the public Internet. ExpressRoute is priced based on the speed of the outbound data. Inbound data is free.

- **Migration Costs.** This is the cost of the time and manpower used to plan, execute and test the migration of Purchasing Entity data from their premises (or hosted environment) to the Microsoft Cloud. This would include stored data, as well as emails, documents and web content that will be maintained in the Cloud. Insight can be contracted to provide these migration services, and will offer either a packaged offering for a standard set of services, or a full offering based on the time and materials necessary to deliver the migration.

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

Insight Response: Due to Insight's status as a Value Added Reseller, this requirement is not applicable to Insight. However, we have provided responses explaining how our chosen CSP partner's solutions are NIST compliant. Confirmation of compliancy is provided in the answer to Section 8.1.2.

Scope and Variety of Cloud Solutions (RFP 6.6) (E)

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

Insight Response: Insight has partnered with two of our strongest cloud partners in response to the NASPO ValuePoint/State of Utah's RFP. Amazon Web Services (AWS) is our first partner offering which has an expansive SaaS, PaaS, and IaaS portfolio offering that will be made available and is outlined throughout Insight's proposal response. AWS' offerings can be delivered via Public and Hybrid deployment models.

We have also partnered with a 3rd party services firm who specializes in AWS consulting services around design and deployment that will assist in leveraging these services to the fullest advantage possible.

The Insight REAN Cloud team is able to provide the following services:

Strategy Phase - SaaS	Assessment Phase - SaaS	Operations Phase - SaaS	DevOps Phase - PaaS
ROI & Business Case Justification (Activity) AWS Calculator (Task) Cloud Rationalization/Adoption strategy DR & Business continuity planning DevOps Strategy Account Management Governance & Compliance	Cloud Architecture Security & Risk Assessment Migration and Implementation Phase Secure Infrastructure Setup Lift & Shift Migration (CloudEndure) DevOps based migration	Managed Services (MGS) Billing as Service (BaaS) AWS Infrastructure (IaaS)	Infrastructure Automation Application Reengineering Native AWS Application Development

Insight's second cloud partnership is with Microsoft. Participating Entities will have access to IaaS, SaaS, and PaaS solution offerings delivered via Public (including the Government Community Cloud), Hybrid, and Private deployment models. Through the Microsoft partnership Office 365, Azure, Intune, and CRM Dynamics will be made available to the Participating Entities. Insight services will provide design and deployment capabilities for Office 365, Azure, and CRM Dynamics. Further description of the Online Services available is provided below.

Those Online Services which **do not store or process** Customer Data, and are merely desktop applications delivered using Microsoft's servers as a delivery mechanism. As of the date of Insight's submission, there are three such Online Services: (1) Office 365 ProPlus; (2) Project Pro for Office 365; and (3) Visio Pro for Office 365.

Those Online Services which **store or process** Customer Data, and are included in scope for the Data Processing Terms (DPT) section of the Microsoft Online Services Terms. As of the date of Insight's submission, that list exclusively includes the following:

Online Services in DPT	
Microsoft Dynamics CRM Online Services	Microsoft Dynamics CRM Online services made available through volume licensing or the Microsoft online services portal, excluding (1) Microsoft Dynamics CRM for supported devices, which includes but is not limited to Microsoft Dynamics CRM Online services for tablets and/or smartphones and (2) any separately-branded service made available with or connected to Microsoft Dynamics CRM Online, such as Microsoft Social Engagement, Parature, from Microsoft, and Microsoft Dynamics Marketing.
Office 365 Services	The following services, each as a standalone service or as included in an Office 365-branded plan or suite: Exchange Online, Exchange Online Archiving, Exchange Online Protection, Advanced Threat Protection, SharePoint Online, OneDrive for Business, Project Online, Skype for Business Online, Sway, Office Online, Delve Analytics, Customer Lockbox, and Yammer Enterprise. Office 365 Services do not include Office 365 ProPlus, any portion of PSTN Services that operate outside of Microsoft's control, any client software, or any separately branded service made available with an Office 365-branded plan or suite, such as a Bing or a service branded "for Office 365."
Microsoft Azure Core Services	Cloud Services (web and worker roles), Virtual Machines (including with SQL Server), Storage (Blobs, Tables, Queues), Virtual Network, Traffic Manager, Batch, Web Sites, BizTalk Services, Media Services, Mobile Services, Service Bus, Notification Hub, Workflow Manager, Express Route, Scheduler, Multi-Factor Authentication, Active Directory, Rights Management Service, SQL Database, and HDInsight.
Microsoft Intune Online Services	The cloud service portion of Microsoft Intune such as the Microsoft Intune Add-on Product or a management service provided by Microsoft Intune such as Mobile Device Management for Office 365.

Best Practices (RFP 6.7) (E)

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

Insight Response: Insight has provided responses explaining how our CSP partners and services partner, REAN Cloud, ensures visibility, compliance, data security and threat protection for cloud-delivered services.

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

In addition, network traffic between AWS regions, Availability Zones, and individual data centers travels over private network segments by default. These private network segments are fully isolated from the public Internet and not routable externally. AWS resources can be configured to reside only on isolated AWS network segments and to avoid utilizing any public IP addresses or routing over the public Internet.

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

Built-In Security Features

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

- **Secure Access** – Customer access points, also called Application Programming Interface (API) endpoints, allow secure HTTP access (HTTPS) so that customers can establish secure communication sessions with their AWS cloud services using Secure Socket Layer (SSL)/Transport Layer Security (TLS).
- **Built-In Firewalls** – Customers can control how accessible their instances are by configuring built-in firewall rules—from totally public to completely private or somewhere in between. And when instances reside within an **Amazon Virtual Private Cloud (Amazon VPC)** subnet, customers can control egress as well as ingress.
- **Unique Users** – The **AWS Identity and Access Management (IAM)** tool allows AWS customers to control the level of access their own users have to AWS infrastructure services. With AWS IAM, each user can have unique security credentials, eliminating the need for

shared passwords or keys and allowing the security best practices of role separation and least privilege.

- **Multi-Factor Authentication (MFA)** – AWS provides built-in support for **MFA** for use with AWS accounts as well as individual AWS IAM user accounts.
- **Private Subnets**–The Amazon VPC service allows customers to add another layer of network security to instances by creating private subnets and even adding an Internet Protocol Security (IPsec) Virtual Private Network (VPN) tunnel between a home network and Amazon VPC.
- **Encrypted Data Storage** – Customers can have the data and objects they store in Amazon Elastic Block Store (Amazon EBS), Amazon Simple Storage Service (Amazon S3), Amazon Glacier, Amazon Redshift, and Amazon Relational Database Service (Amazon RDS) on Oracle and SQL Server encrypted automatically using Advanced Encryption Standard (AES) 256, a secure symmetric-key encryption standard using 256-bit encryption keys.
- **Dedicated Connection Option** –The **AWS Direct Connect** service allows customers to establish a dedicated network connection from their premises to AWS. Using industry-standard 802.1q VLANs, this dedicated connection can be partitioned into multiple logical connections to enable access to both public and private IP environments within the AWS cloud.
- **Dedicated, Hardware-Based Crypto Key Storage Option** – For customers who must use Hardware Security Module (HSM) appliances for cryptographic key storage, **AWS CloudHSM** provides a highly secure and convenient way to store and manage keys.
- **Centralized Key Management**– For customers who use encryption extensively and require strict control of their keys, the **AWS Key Management Service (KMS)** provides a convenient management option for creating and administering the keys used to encrypt data at rest.
- **Perfect Forward Secrecy**– For even greater communication privacy, several AWS cloud services such as **Elastic Load Balancing** and **Amazon CloudFront** offer newer, stronger cipher suites. These cipher suites allow SSL/TLS clients to use Perfect Forward Secrecy, a technique that uses session keys that are ephemeral and not stored anywhere. This prevents the decoding of captured data, even if the secret long-term key itself is compromised.

AWS Multi-Factor Authentication (MFA)

(MFA) is a simple best practice that adds an extra layer of protection on top of your user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their user name and password (the first factor—what they know), as well as for an authentication code from their AWS MFA device (the second factor—what they have). Taken together, these multiple factors provide increased security for your AWS account settings and resources.

AWS GovCloud (US) is an isolated AWS region designed to host sensitive workloads in [the cloud](#), ensuring that this work meets the US government's regulatory and compliance requirements. The AWS GovCloud (US) region adheres to United States International Traffic in Arms Regulations (ITAR) as well as Federal Risk and Authorization Management Program (FedRAMP) requirements. It provides special endpoints that utilize only Federal Information Processing Standard (FIPS) 140-2 encryption. AWS GovCloud (US) is available to US government agencies, government contractors, private and public commercial entities, educational institutions, non-profits and research organizations that meet GovCloud (US) requirements for access.

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

- **AWS Trusted Advisor** – Provided automatically when AWS customers sign up for premium support, the **AWS Trusted Advisor** service is a convenient way for customers to see where they could use a little more security. It monitors AWS resources and alerts customers to security configuration gaps such as overly permissive access to certain Amazon Elastic Compute Cloud (Amazon EC2) instance ports and Amazon S3 storage buckets, minimal use of role segregation using AWS IAM, and weak password policies.
- **Amazon CloudWatch** – **Amazon CloudWatch** enables customers to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by a customer's applications and services and any log files their applications generate. Customers can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health, using these insights to react intelligently and keep applications running smoothly.
- **AWS CloudTrail** – **AWS CloudTrail** provides logs of all user activity within an AWS account. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS cloud service. The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.
- **AWS Config** – With the **AWS Config** service, customers can immediately discover all of their AWS resources and view the configuration of each. Customers can receive notifications each time a configuration changes as well as dig into the configuration history to perform incident analysis.

Third-Party Security Tools

Amazon also offers additional third-party security tools to complement and enhance their customers' operations in the AWS cloud. **AWS Partner Network (APN)** partners offer hundreds of familiar and industry-leading products that are equivalent to, identical to, or integrate with existing

controls in a customer's on-premises environments. Customers can browse and purchase APN partner products on the **AWS Marketplace**. These products complement existing AWS cloud services to enable customers to deploy a comprehensive security architecture and a more seamless experience across their cloud and on-premises environments. The APN partner security products cover multiple areas of security, including application security, policy management, identity management, security monitoring, vulnerability management, and endpoint protection. The figure below is a snapshot of the APN partners and categories of products available under the security category in the AWS Marketplace.

Several of the security products that AWS offers are provided only by APN partners that are prequalified by the **APN Partner Competency Program**, which confirms their technical proficiency and proven customer success in specialized solution areas. AWS's **Security Competency Partners** can also provide demos and consulting services that are not always available through the AWS Marketplace.



Figure 10: AWS Marketplace Security Offerings

**Value Added Security, Visibility, Compliance and Threat Protection Best Practices
Insight/ REAN Provides:**

Amazon Virtual Private Cloud (Amazon VPC) lets customers provision a private, isolated section of the Amazon Web Services (AWS) Cloud where members can launch AWS resources in a virtual network that they define. With Amazon VPC, users can define a virtual network topology that closely resembles a traditional network that they might operate in their own data center. REAN will help Participating Entities 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.

REAN will help Participating Entities customize the network configuration for their Amazon VPC. For example, Participating Entities may need a public-facing subnet for their web servers that have access to the Internet, and place their backend systems such as databases or application servers in a private-facing subnet with no Internet access. REAN will help Participating Entities leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Additionally, REAN will help Participating Entities create a Hardware VPN connection between their corporate data center and their VPC and leverage the AWS cloud as an extension of their corporate data center. Figure below shows a notional picture of the AWS VPC infrastructure offering.

A variety of connectivity options exist for Participating Entities to connect to their Amazon VPC: Participating Entities can connect their VPC to the Internet, to their datacenter, or both, based on the AWS resources that they want to expose publicly and those that they want to keep private.

- Connect directly to the Internet (public subnets) – Participating Entities can launch instances into a publicly accessible subnet where they can send and receive traffic from the Internet.
- Connect to the Internet using Network Address Translation (private subnets) – Private subnets can be used for instances that Participating Entities do not want to be directly addressable from the Internet. Instances in a private subnet can access the Internet without exposing their private IP address by routing their traffic through a Network Address Translation (NAT) instance in a public subnet.
- Connect securely to a corporate datacenter – All traffic to and from instances in a Participating Entity's VPC can be routed to their corporate datacenter over an industry standard, encrypted IPsec hardware VPN connection.
- Combine connectivity methods to match the needs of the application – Customers can connect a VPC to both the Internet and their corporate datacenter and configure Amazon VPC route tables to direct all traffic to its proper destination.




Figure 11: AWS VPC Infrastructure Offering

Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level. In addition, Participating Entities can store data in Amazon S3 and restrict access so that it's only accessible from instances in their VPC. Optionally, Participating Entities can also choose to launch Dedicated Instances that run on hardware dedicated to a single customer for additional isolation.

REAN Secure VPC

REAN has devised a secure virtual private cloud (S-VPC) framework that provides assurance of information protection with additional security controls to ensure the confidentiality, integrity and availability of information. The S-VPC wraps the customer application in a secure shell to meet the internal governance and ensure compliance with regulations like SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 type II), PCI DSS Level 1, ISO 27001, HIPAA, HITECH, and FedRAMP.

Figure below shows the high level architecture for REAN S-VPC. The following sections explain virtual network, server, storage, access control, and audit controls in further detail.




Figure 12: High Level Architecture for REAN S-VPC

Network Protection

REAN S-VPC protects the network perimeter by creating a Demilitarized Zone (DMZ) with a unified threat management suite. The suite provides firewall services, intrusion protection/detection services, secure Virtual Private Network (VPN) connectivity, packet filtering, and web application firewall protection not available via AWS standard offerings. This front-end protects against denial-of-service attacks, worms, and hacker exploits.

Server Protection

REAN S-VPC offers comprehensive server security designed to protect all the AWS instances in the customer environment from data breaches and business disruptions, and achieve cost-effective compliance across these environments.

Tightly integrated modules including anti-malware, web reputation, firewall, host based intrusion prevention, integrity monitoring, and log inspection expand the security posture to ensure server, application, and data security across physical, virtual, and cloud environments. The solution also features FIPS 140-2 certification to support high security standards.

Storage Protection

REAN S-VPC provides distinctive data protection for information stored on elastic block store volumes using encryption with key management system that enables policy based restrictions to determine where and when encrypted data can be accessed. In addition, server validation applies identity and integrity rules when servers request access to secure storage volumes. Solution ensures that encryption keys are delivered to valid devices without the need to deploy an entire file system and management infrastructure. This solution protects sensitive information from theft, unauthorized exposure, or unapproved geographic migration to other data centers.

Access Control

REAN S-VPC environment provides various convenient options to the end users to access the environment and initiate their VPN connections. These include:

- HTML5 based remote access VPN that they can initiate from any HTML5 compatible browser with requiring any plug-in.
- SSL remote access VPN that provides additional security by a double authentication using X.509 certificates and username/password.
- IPSec based VPN using native Windows or Mac VPN clients.
- Mobile VPN using native iPhone VPN client to securely connect to VPC.

System administrator access control is provided through the integration of GU identity and access management solution. This suite supplements the AWS Management Console by vaulting administrator's credentials, enforcing separation of duties, and recording all accesses and actions.

Logging and Auditing

REAN S-VPC ensures that the customer environment is continuously monitored using auditing at the network, server, and application levels to help meet all the forensics and compliance requirements. In case of server and infrastructure access, the solution not only provides system logs but could optionally provide full video stream of an administrator session into Amazon S3. By providing such video stream that is tied back to customer Identity and Access Management (IAM), enterprises can maintain full accountability for any changes performed on the service. The entire above audit data is fed into a Security Information and Event Management (SIEM) system that provides full contextual awareness of the events that can be summarized in a simple dashboard.

Availability

Customer environment is architected to take full advantage of highly available AWS infrastructure. All the components (application servers and files stores) of the solution are deployed in a redundant fashion across multiple fault isolated AWS Availability Zones. 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. The file store uses Amazon Simple Storage Service (S3) that provides eleven 9s SLA on durability of the data.

REAN S-PVC Value

REAN S-VPC has successfully passed security testing and auditing by a leading auditor that provides the services to the Department of Defense. Participating States and Entities can adopt a proven and working framework and save time and money.

Auditing and Compliance

REAN firmly believes in the separation of duties between architecting and implementing security solutions and auditing and accreditation process to ensure compliance. Hence, REAN will work with any independent third party vendor a Participating Entity recommends to help them through the certification and accreditation process.

5. Organization Profile (RFP 7) (M) (E)

7.1 Contract Manager

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

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

Contract Manager Information
Name: Pam Potter
Phone: (630) 924-6810
Email: Pam.Potter@Insight.com
Work Hours: 8:00am – 5:00pm (CST)

Insight Response: While Pam Potter will serve as the Contract Manager and primary point of contact, she will be supported by an Insight Product Manager who is well versed in the AWS and Microsoft Cloud offerings and associated contract terms. A professional bio has been provided below.

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

Insight Response: Pam Potter is a seasoned professional in the IT industry having spent the past 15 years with Insight, of which four have been spent as a member of the Insight Contract Compliance team. She is responsible for managing all IPS SLED contracts, as well as SLED-related client and nationwide non-services partner agreements. Pam's primary responsibilities include new contract roll-outs, contract lifecycle/ change management, and ongoing training efforts. She further serves as the primary contact for the client contract managers of Insight's 160+ contracts. She has extensive experience managing large multi-state contracts, including the NASPO Software Value Added Reseller and U.S. Communities contracts that Insight holds.

David Solliday is a seasoned professional in the IT industry with over 25 years of experience in the channel, with a focus on Cloud for the past six (6) years. He serves as the Business Development Specialist for Insight's Cloud Practice. His current responsibilities are around Managing Infrastructure partners as well as the responsibility of Onboarding new Cloud business partners and how we operationalize them into Insight's business systems. David will serve as Pam's deputy in all matters that pertain to the solution-specific offering terms and technical details.

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

Contract Manager

Insight Response: Ms. Pam Potter will have the authority and responsibility for the overall success of the NASPO ValuePoint Cloud Solutions contract within Insight's organization. As the Manager of Compliance for IPS' SLED contracts, Ms. Potter is responsible for working directly with the NASPO ValuePoint organization and state procurement offices to ensure that Insight is properly administering the Master Agreement and Participating Addendums. Areas of responsibility include managing the following contract items:

✓ Renewals	✓ Extensions	✓ Product Catalog	✓ Update Schedules
✓ Issues	✓ Daily Contract Management	✓ Reporting	✓ Contract Management Website

Additionally, Ms. Potter will work directly with state procurement offices to educate them on Insight's NASPO ValuePoint contract. She will assist participating states and individual entities throughout the addendum signing process, ensuring each PA is executed smoothly. She will be a dedicated point of contact for states and entities interested in and going through the sign up process. Additionally, she works closely with Insight sales and provides appropriate contract training and compliance oversight.

6. Technical Response (RFP 8) (M) (E)

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

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

8.1 TECHNICAL REQUIREMENTS

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

Insight Response: Insight's proposed solutions that we will provide through the NASPO ValuePoint Cloud Solutions contract are outlined below.

Vendor Name	Service Model(s)	Deployment Model(s)
Amazon Web Services	IaaS, PaaS, SaaS	Public, Hybrid, Private
Microsoft	IaaS, SaaS, PaaS	Public, Hybrid, Private
Insight Public Sector	Not Applicable	Not Applicable
REAN Cloud	Not Applicable	Not Applicable

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

Insight Response: This requirement is not applicable to Insight but is relevant to our CSP partners. Provided below is how their solutions meet the NIST Characteristics.

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

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

- **Management Console:** The AWS Management Console is the self-service, on-demand destination for managing all AWS resources, from Amazon Elastic Compute Cloud (Amazon EC2) instances to Amazon DynamoDB tables. The AWS Management Console is used 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 allows 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).

Microsoft: Microsoft Windows Azure is an internet-scale, high-availability cloud fabric operating on globally-distributed Microsoft data centers. Windows Azure and related tools support the development and deployment of applications into a hosted environment that extends the on-premises data center. On-demand self-service refers to the service provided by Microsoft Azure that enables the provision of resources on demand whenever required. On-demand services can be enabled from the HTML Portal, using Azure API, using CLI for Mac, Linux, and Windows with Azure Service Management. Azure on-demand self-service resource sourcing programmatically is a prime feature allowing the user to scale the infrastructure.

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

AWS: 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 users provision and use what they need via a web application, mobile client, or programmatically through published and well documented APIs. AWS products and solutions that support broad network access include Amazon Route 53, a scalable domain name system, Virtual Private Cloud to isolate network resources, AWS Direct Connect, a dedicated network connection to non-AWS resources, and Auto Scaling to respond to significant changes to network resource demands.

Microsoft: Microsoft Azure enables the enterprise to create an on-premises network route from on-premises VPN device and the Azure virtual network. Configure on-premises hardware or software VPN device to terminate the VPN tunnel, which uses Internet Protocol security (IPsec). Gateway connects on-premises to Microsoft Azure through many connection bandwidth using software only or hardware based connectivity, including Basic VPN connection, Standard VPN, and ExpressRoute gateway connections. Create internal Microsoft Azure virtual network to support separation of site resources or sharing of information and subscription to subscription connection.

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

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

Microsoft: Azure resource pooling supports scalable systems involved in cloud computing and software as a service (SaaS), enabling “near” infinite growth with immediately availability. The kinds of services that can apply to a resource pooling strategy include data storage services, processing services, bandwidth provided services and other Azure related compute elasticity.

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

AWS: AWS provides a massive global cloud infrastructure that allows users to quickly innovate, experiment, and iterate. Instead of waiting weeks or months for hardware, users can instantly deploy new applications, instantly scale up as 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 an unexpected spike in demand and then scale those resources down as demand decreases.

Microsoft: Microsoft Azure supports rapid elasticity allowing automated requests for additional resources (i.e. compute, disk space, connectivity and other types of services). Microsoft Azure services are allocated and de-allocated resources irrelevant to the client or user's side. Microsoft Azure provides resources that appear to be “nearly” infinite with automatic availability.

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

AWS: AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

In addition, the Insight/REAN team provides a variety of value added managed services in our delivery of Cloud Solutions. We provide monitoring and management across a spectrum of metrics that define measured service. Our team also provides consolidated billing services to assist Participating Entities in monitoring and measuring utilization and spend.

REAN reviews customer’s current AWS environments, collaboratively develops relevant metrics, deploys REAN monitoring agents, works with customer team to identify alerting thresholds, notification groups and make necessary changes to take over management of systems.

Sample Summary of REAN Activities in Measuring, Monitoring and Managing Customer's AWS Environment:

- Acquire access to customer's AWS environment.
- Assessment report of current customer environment including servers, OS versions, software tools, data volumes with size and user accounts.
- Determine Relevant Metrics for customer environment to provide measured service.
- Deploy REAN monitoring agents and tools.
- Setup access for customer team members to REAN ticketing system.
- Work with customer team to finalize alerting thresholds and notification groups.
- Start monitoring customer's AWS environments as per the agreed service levels defined below
- Provide monthly managed services reports for infrastructure.
- Regular patching and other software updates to resources in the environment.

Deliverables

The following are deliverables to customer:

- Executive summary report (Monthly)
- Security reports (Monthly)
- Compliance reports (HIPAA, Windows Hardening, Patches, etc.. as applicable) (Monthly)
- Operations reports (Monthly)
- Budget policy recommendations (Monthly)
- Uptime reports (if applicable) (Monthly)
- Usage/traffic reports (if applicable) (Monthly)
- AMI Backups (daily, weekly and monthly backups)
- Patching systems (as needed with approval from customer team)
- Software updates (as needed for tools used for managed services)

Microsoft: Microsoft Azure supports NIST principles in areas of measured services such as measured service setup allowing Azure systems to control a system, user or tenant's usage of resources with metering capability. Azure supports automated remote services measurement tools to provide auditing and accountability of utilization. Azure measured service ensures that even when there is no specific interaction for a service change, that service change is still audited to support billing cycles.

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

Insight Response: Through the NASPO ValuePoint Cloud Solutions contract Insight is offering solutions for each of the service models – IaaS, PaaS, and SaaS. Provided in the responses below is a description of the various service model offerings we will be able to provide.

The Insight/REAN Cloud team provides a number of SaaS offerings across a broad spectrum of functionality and industries. Examples include:

1. REAN Insights - Analyzing rapidly growing social media data requires flexible computing resources. The fast-growing social media landscape creates terabytes of data per day, with unpredictable volumes. Participating States and Entities need to be able to scale their peak compute capacity, coordinate secure access for multiple analysts, and share results. REAN Insights is crucial in understanding the pulse of the customers and their views on latest trends and policies. The solution listens in to relevant social intelligence, analyzes them and provides insights to help achieve goals and become highly relevant, proactive and market-oriented. Public Sector customers, such as state and local governments, can benefit from this capability by being able to analyze constituent social media activity.
2. REAN Sites (Large Scale Content Management) – Automated, secure, highly scalable, turnkey Drupal Solution.
3. REAN Genomics - Analyzing huge amounts of sequencing data requires formidable computing resources. The expanding scale of Genomics Research creates analytical challenges like accommodating peak compute demand, coordinating secure access for multiple analysts, and sharing validated tools and results. REAN Genomics leverages AWS cloud to provide a secure platform to help users scale research cost-effectively. It supports high performance workloads to derive powerful insights with zero capital investments. Large public research institutions will find this to be a valuable tool.
4. HIPAA Compliant SaaS Solutions - Covered entities and their business associates subject to the U.S. Health Insurance Portability and Accountability Act (HIPAA) can leverage the built-in security offered by REAN Secure Virtual Private Cloud (S-VPC) on AWS to process, maintain, and store protected health information. Any agency charged with protecting health related data in compliant manner will benefit from this capability.
5. REAN Migrate and Manage - REAN Cloud provides solutions that empower innovation and help safe and seamless cloud transition for organizations, without compromising productivity, security, customer service or budget. Our mature Migration Methodology and custom cloud solutions help improve the efficiency and secure applications and data in Participating States and Entity's organization, along with leveraging the flexibility, scalability, elasticity and cost-savings of the cloud. REAN's management of cloud systems follows AWS best practices that include a highly available infrastructure, failover protection, and auto scaling. This offer is public sector market agnostic, and entities ranging from State and Local Governments to school systems can benefit from this capability.
6. REAN Open Learning - REAN Open Learning, built on the agile and secure AWS cloud platform, helps build and launch the customized Online Education platform quickly and cost-effectively. REAN Open Learning platform based on edX is used by universities in various countries to develop innovative online, on-campus, and blended teaching and

learning models. Our team of experts works on the following principles: Performance, Security and Availability. This end-to-end managed solution enables K-12 and Higher Education institutions to launch online courses in a secure, scalable and cost effective way.

PaaS

The Insight REAN Team also provide PaaS solutions in the form of custom DevOps pipeline solutions. REAN Cloud delivers an automated CI/CD SecDevOps pipeline on AWS that goes from code to release automatically. REAN Cloud can implement a continuous integration and delivery pipeline on AWS and instill a DevOps culture for a user's dev teams. REAN provides a combination of DevOps and AWS expertise while also delivering managed services through CloudOps & SecOps.

Amazon Web Services and Microsoft: The tables below outline Microsoft and AWS's product offerings broken down by service model categories and subcategories.

Amazon Web Services Product Categories, Subcategories, and Service Models			
IaaS	AWS Product	Description	Service Model
Compute			
	Amazon EC2	Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable computing capacity—literally, servers in Amazon's data centers—that you use to build and host your software systems.	Public
	Amazon EC2 Container Service	Amazon EC2 Container Service is a highly scalable, high-performance container management service that supports Docker containers and allows you to easily run distributed applications on a managed cluster of Amazon EC2 instances.	Public
	AWS Lambda	AWS Lambda is a compute service that runs your code in response to events and automatically manages the compute resources for you, making it easy to build applications that respond quickly to new information. AWS Lambda starts running your code within milliseconds of an event such as an image upload, in-app activity, website click, or output from a connected device.	Public

	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.	Public
	Elastic Load Balancing	Elastic Load Balancing automatically distributes your incoming application traffic across multiple Amazon EC2 instances. It detects unhealthy instances and reroutes traffic to healthy instances until the unhealthy instances have been restored. Elastic Load Balancing automatically scales its request handling capacity in response to incoming traffic.	Public
IaaS			
Networking			
	Amazon VPC	Amazon Virtual Private Cloud (Amazon VPC) enables you to launch AWS resources into a virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS.	Private
	Amazon Route 53	Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service.	Public

		AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 GB or 10 GB Ethernet fiberoptic cable. One end of the cable is connected to your router and the other to an AWS Direct Connect router. With this connection, you can create virtual interfaces directly to the AWS cloud and Amazon VPC, bypassing Internet service providers in your network path.	
	AWS Direct Connect		Public
IaaS			
Storage and Content Delivery			
	Amazon S3		
		Amazon Simple Storage Service (Amazon S3) is storage for the Internet. You can use Amazon S3 to store and retrieve any amount of data, at any time, from anywhere on the web. You can accomplish these tasks using the simple and intuitive web interface of the AWS Management Console.	Public
	Amazon Glacier		
		Amazon Glacier is a storage service optimized for infrequently used data, or "cold data." The service provides secure, durable, and extremely low-cost storage for data archiving and backup. With Amazon Glacier, you can store your data cost effectively for months, years, or even decades. Amazon Glacier enables you to offload the administrative burdens of operating and scaling storage to AWS, so you don't have to worry about capacity planning, hardware provisioning, data replication, hardware failure detection and recovery, or time-consuming hardware migrations.	Public
	Amazon EBS		

		Amazon Elastic Block Store (Amazon EBS) provides block-level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are highly available and reliable storage volumes that can be attached to any running instance that is in the same Availability Zone. Amazon EBS volumes that are attached to an Amazon EC2 instance are exposed as storage volumes that persist independently from the life of the instance. With Amazon EBS, you only pay for what you use.	Public
	Amazon CloudFront		
		Amazon CloudFront is a content delivery web service. It integrates with other AWS cloud services to give developers and businesses an easy way to distribute content to end users with low latency, high data transfer speeds, and no commitments.	Public
	AWS Import/Export		
		AWS Import/Export accelerates transferring large amounts of data between the cloud and portable storage devices that you mail to us. AWS transfers data directly onto and off of your storage devices using Amazon's high-speed internal network. Your data load typically begins the next business day after your storage device arrives at AWS. After the data export or import completes, we return your storage device. For large data sets, AWS Import/Export is significantly faster than Internet transfer and more cost effective than upgrading your connectivity.	Public
	AWS Storage Gateway		

		AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and AWS's storage infrastructure.	Hybrid
IaaS			
Databases			
	Amazon RDS		
		Amazon Relational Database Service (Amazon RDS) is a web service that makes it easier to set up, operate, and scale a relational database in the cloud. It provides cost-efficient, resizable capacity for an industry-standard relational database and manages common database administration tasks. Database engines available through Amazon RDS include Amazon Aurora, MySQL, Oracle, Microsoft SQL Server, and PostgreSQL.	Public
	Amazon DynamoDB		
		Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. You can use Amazon DynamoDB to create a database table that can store and retrieve any amount of data, and serve any level of request traffic. Amazon DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent and fast performance.	Public
	Amazon Redshift		

		Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse solution that makes it simple and cost-effective to efficiently analyze all your data using your existing business intelligence tools. You can start small for just \$0.25 per hour with no commitments or up-front costs and scale to a petabyte or more for \$1,000 per terabyte per year, less than a tenth of most other data warehousing solutions.	
	Amazon ElastiCache		
		Amazon ElastiCache is a web service that makes it easy to set up, manage, and scale distributed, in-memory cache environments in the cloud. It provides a high-performance, resizable, and cost-effective in-memory cache, while removing the complexity associated with deploying and managing a distributed cache environment.	Public
IaaS			
Analytics			
	Amazon EMR		
		Amazon Elastic MapReduce (Amazon EMR) is a web service that makes it easy to process large amounts of data efficiently. Amazon EMR uses Hadoop processing combined with several AWS products to perform such tasks as web indexing, data mining, log file analysis, machine learning, scientific simulation, and data warehousing.	Public
	Amazon Kinesis		

		Amazon Kinesis is a managed service that scales elastically for real-time processing of streaming big data. The service takes in large streams of data records that can then be consumed in real time by multiple data processing applications that can be run on Amazon EC2 instances. The data processing applications use the Amazon Kinesis Client Library and are called "Amazon Kinesis applications."	Public
	AWS Data Pipeline		
		AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premises data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to AWS cloud services such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon EMR.	Public
	Amazon Mobile Analytics		
		Amazon Mobile Analytics is a service that lets you easily collect, visualize, and understand application usage data at scale. Many mobile application analytics solutions deliver usage data several hours after the events occur. Amazon Mobile Analytics is designed to deliver usage reports within 60 minutes of receiving data from an application so that you can act on the data more quickly.	Public
IaaS			
Administration & Security			
	AWS Identity & Access Management		

		AWS Identity and Access Management (IAM) is a web service that enables AWS customers to manage users and user permissions in AWS. The service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console. With AWS IAM, you can centrally manage users, security credentials such as access keys, and permissions that control which AWS resources users can access.	Public
	AWS Directory Service		
		AWS Directory Service is a managed service that allows you to connect your AWS resources with an existing on-premises Microsoft Active Directory or to set up a new, stand-alone directory in the AWS cloud. Connecting to an on-premises directory is easy, and once this connection is established, all users can access AWS resources and applications with their existing corporate credentials.	Public
	AWS Service Catalog		
		AWS Service Catalog is a service that allows administrators to create and manage approved catalogs of resources that end users can then access via a personalized portal. You can control which users have access to which applications or AWS resources to enable compliance with your business policies, while users can easily browse and launch products from the catalogs you create.	Public
	AWS Config		

		AWS Config is a fully managed service that provides you with an AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. With AWS Config you can discover existing AWS resources, export a complete inventory of your AWS resources with all configuration details, and determine how a resource was configured at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.	Public
	AWS CloudHSM		
		AWS CloudHSM provides secure cryptographic key storage to customers by making Hardware Security Modules (HSMs) available in the AWS cloud.	
	AWS Key Management Service		
		AWS Key Management Service (KMS) is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data and uses Hardware Security Modules (HSMs) to protect the security of your keys. AWS KMS is integrated with other AWS cloud services including Amazon EBS, Amazon S3, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.	Public
	AWS CloudTrail		

		With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS Software Development Kits (SDKs), the command line tools, and higher-level AWS cloud services. You can also identify which users and accounts called AWS APIs for services that support AWS CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate AWS CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn AWS CloudTrail logging on and off.	Public
	Amazon CloudWatch		
		Amazon CloudWatch is a web service that enables you to collect, view, and analyze metrics. Amazon CloudWatch lets you programmatically retrieve your monitoring data, view graphs, and set alarms to help you troubleshoot, spot trends, and take automated action based on the state of your cloud environment.	Public
IaaS			
Deployment & Management			
	AWS Management Console		
		Access and manage Amazon cloud services through a simple and intuitive web-based user interface. You can also use the AWS Console mobile app to quickly view resources on-the-go.	Public
	AWS Command Line Interface		

		The AWS Command Line Interface (CLI) is a unified tool used to manage your AWS cloud services. With just one tool to download and configure, you can control multiple AWS cloud services from the command line and automate them through scripts.	Public
	APIs		
		AWS provides API-based cloud computing services with multiple interfaces to those services, including SDKs, IDE Toolkits, and Command Line Tools for developing and managing AWS resources.	Public
	AWS Elastic Beanstalk		
		With AWS Elastic Beanstalk, you can quickly deploy and manage applications in the AWS cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.	Public
	AWS CloudFormation		
		AWS CloudFormation gives developers and system administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.	
		You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.	Public
	AWS CodeDeploy		

		AWS CodeDeploy is a service that automates code deployments to Amazon EC2 instances. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate deployments, eliminating the need for error-prone manual operations, and the service scales with your infrastructure so you can easily deploy to one Amazon EC2 instance or thousands.	Public
	AWS CodeCommit		
		AWS CodeCommit is a secure, highly scalable, managed source control service that hosts private Git repositories. AWS CodeCommit eliminates the need for you to operate your own source control system or worry about scaling its infrastructure. You can use AWS CodeCommit to store anything from code to binaries, and it supports the standard functionality of Git, allowing it to work seamlessly with your existing Git-based tools.	Public
	AWS CodePipeline		
		AWS CodePipeline is a continuous delivery and release automation service that aids smooth deployments. You can design your development workflow for checking in code, building the code, deploying your application into staging, testing it, and releasing it to production. You can integrate third-party tools into any step of your release process or you can use AWS CodePipeline as an end-to-end solution.	Public
	AWS OpsWorks		

		AWS OpsWorks provides a simple and flexible way to create and manage stacks and applications. With AWS OpsWorks, you can provision AWS resources, manage their configuration, deploy applications to those resources, and monitor their health.	Public
IaaS			
Enterprise Applications			
	Amazon WorkDocs		
		Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity. Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments.	Public
	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.	Public
PaaS			
Application Services			
	Amazon AppStream		

		The Amazon AppStream web service deploys your application on AWS infrastructure and streams input and output between your application and devices such as personal computers, tablets, and mobile phones. Your application's processing occurs in the cloud, so it can scale to handle vast computational loads. Devices need only display output and return user input, so the client application on the device can be lightweight in terms of file size and processing requirements.	Public
	Amazon CloudSearch		
		Amazon CloudSearch is a fully managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. With Amazon CloudSearch, you can quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.	Public
	Amazon SWF		

		Amazon Simple Workflow Service (Amazon SWF) makes it easy to build applications that coordinate work across distributed components. In Amazon SWF, a task represents a logical unit of work that is performed by a component of your application. Coordinating tasks across the application involves managing intertask dependencies, scheduling, and concurrency in accordance with the logical flow of the application. Amazon SWF gives you full control over implementing tasks and coordinating them without worrying about underlying complexities such as tracking their progress and maintaining their state.	Public
	Amazon SQS		
		Amazon Simple Queue Service (Amazon SQS) is a messaging queue service that handles messages or workflows between other components in a system.	Public
	Amazon SES		
		Amazon Simple Email Service (Amazon SES) is an outbound-only email-sending service that provides an easy, cost-effective way for you to send email.	Public
	Amazon SNS		
		Amazon Simple Notification Service (Amazon SNS) is a web service that enables applications, end users, and devices to instantly send and receive notifications from the cloud.	Public
	Amazon Elastic Transcoder		

		Amazon Elastic Transcoder lets you convert media files that you have stored in Amazon S3 into media files in the formats required by consumer playback devices. For example, you can convert large, high-quality digital media files into formats that users can play back on mobile devices, tablets, web browsers, and connected televisions.	Public
	Amazon Cognito		
		Amazon Cognito is a simple user identity and data synchronization service that helps you securely manage and synchronize application data for your users across their mobile devices. You can create unique identities for your users through a number of public login providers (Amazon, Facebook, and Google) and also support unauthenticated guests.	Public
	Amazon FPS		
		Amazon Flexible Payments Service facilitates the digital transfer of money between any two entities, humans or computers.	Public
Support			
	AWS Support		
		AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced and technical support engineers to help customers of all sizes and technical abilities successfully utilize the products and features provided by AWS.	Public
SaaS	AWS Trusted Advisor		

		AWS Trusted Advisor acts like your customized cloud expert, and it helps you provision your resources by following best practices. AWS Trusted Advisor inspects your AWS environment and finds opportunities to save money, improve system performance and reliability, or help close security gaps. Since 2013, customers have viewed over 1.7 million best-practice recommendations and realized over \$300 million in estimated cost reductions.	Public
	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.	Public

Microsoft Product Categories, Subcategories, and Service Models		
	Microsoft Product	Deployment Model
SaaS		
Analytics	Microsoft Azure Core Services	Public Cloud
Data Analytics	Microsoft Azure Core Services	Public Cloud
Business Intelligence	Office 365 Services	Public Cloud
Business Continuity/Disaster Recovery	Microsoft Azure Core Services	Public Cloud
Cloud and Infrastructure Management Tools	Microsoft Azure Core Services	Public Cloud
Collaboration	Office 365 Services	Public Cloud
Customer Relationship Management	CRM Online	Public Cloud
Data Management	Microsoft Azure Core Services	Public Cloud
E-Discovery	Office 365 Services	Public Cloud
Electronic Records Management	Microsoft Azure Core Services	Public Cloud
Office Productivity	Office 365 Services	Public Cloud

Message Filtering	Office 365 Services	Public Cloud
Meeting Planning, hosting, conferencing	Office 365 Services	Public Cloud
Mobile Data Management	Microsoft Intune Online Services	Public Cloud
Security	Microsoft Intune Online Services	Public Cloud
IaaS		
Computer/Infrastructure Services	Microsoft Azure Core Services	Public Cloud
Operating systems	Microsoft Azure Core Services	Public Cloud
Hypervisors	Microsoft Azure Core Services	Public Cloud
Disaster Recovery	Microsoft Azure Core Services	Public Cloud
Business Continuity	Microsoft Azure Core Services	Public Cloud
High Availability / Failover	Microsoft Azure Core Services	Public Cloud
GIS		
Storage	Microsoft Azure Core Services	Public Cloud
File	Microsoft Azure Core Services	Public Cloud
Block	Microsoft Azure Core Services	Public Cloud
Object	Microsoft Azure Core Services	Public Cloud
Archive	Microsoft Azure Core Services	Public Cloud
Cache	Microsoft Azure Core Services	Public Cloud
	Microsoft Azure Core Services	Public Cloud
Content Delivery Networks (CDN)	Microsoft Azure Core Services	Public Cloud
Litigation Hold	Office 365 Services	Public Cloud
Network	Microsoft Azure Core Services	Public Cloud
Virtual network	Microsoft Azure Core Services	Public Cloud
Load balancer	Microsoft Azure Core Services	Public Cloud
DNS	Microsoft Azure Core Services	Public Cloud

Gateway (e.g. VPN or Application)	Microsoft Azure Core Services	Public Cloud
Firewall		
Traffic manager	Microsoft Azure Core Services	Public Cloud
Direct link		
PC/Desktop "aaS"	Microsoft Azure Core Services	Public Cloud
Security	Microsoft Azure Core Services	Public Cloud
Identity & Access Management	Microsoft Azure Core Services	Public Cloud
	Microsoft Azure Core Services	Public Cloud
Encryption		
Data Loss Prevention (DLP)	Exchange Online	Public Cloud
Web Security	Microsoft Azure Core Services	Public Cloud
Email Security	Exchange Online	Public Cloud
Network Security	Microsoft Azure Core Services	Public Cloud
Security Information and Event Management (SIEM)		
Intrusion Management	Microsoft Azure Core Services	Public Cloud
DDOS Monitoring / Management	Microsoft Azure Core Services	Public Cloud
Multi-factor Authentication	Microsoft Azure Core Services	Public Cloud
PaaS		
Analytics	Microsoft Azure Core Services	Public Cloud
Hadoop	Microsoft Azure Core Services	Public Cloud
Business Intelligence	Microsoft Azure Core Services	Public Cloud
Data Warehouse	Microsoft Azure Core Services	Public Cloud

Database	Microsoft Azure Core Services	Public Cloud
Relational	Microsoft Azure Core Services	Public Cloud
NoSQL	Microsoft Azure Core Services	Public Cloud
Development, Testing and Deployment	Microsoft Azure Core Services	Public Cloud
Containers	Microsoft Azure Core Services	Public Cloud
Services and APIs	Microsoft Azure Core Services	Public Cloud
Mobile	Microsoft Intune Online Services	Public Cloud
Internet of Things	Microsoft Azure Core Services	Public Cloud
Tools	Microsoft Azure Core Services	Public Cloud
Runtime environments		
Electronic Records Management	Microsoft Azure Core Services	Public Cloud
E-Discovery	Office 365 Services	Public Cloud
GIS		
Integration (iPaaS)	Microsoft Azure Core Services	Public Cloud
Open Source		
Other (identify additional sub-categories and/or descriptors)		
Project Management	Project Professional for Office 365	Public Cloud
Business Flowcharting and Diagrams	Visio Professional for Office 365	Public Cloud

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

Insight Response: Per Insight's proposal, we are willing to provide Cloud Solutions that comply with the requirements of Attachments C and D.

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

Insight Response: Below Insight has highlighted how AWS and Microsoft's offerings adhere to the services, definitions, and deployment models identifies in the Scope of Services in Attachment D.

AWS: AWS provides NIST compliant cloud infrastructure services. AWS' compliance is validated by two Agency Authority to Operate (ATOs) achieved based on testing performed against the stringent set of FedRAMP requirements (NIST 800-53 Rev. 4 – Moderate baseline requirements, plus additional FedRAMP security controls). They provide federal security personnel with their security documentation as a means of verifying the security and compliance of AWS in accordance with applicable NIST controls as defined by 800-53 rev4 and the DoD Cloud Computing Security Requirements Guide (SRG).

AWS NIST compliant infrastructure services follow the NIST definition of cloud computing and adheres to the five essential characteristics of On-Demand Self Service, Broad Network Access, Rapid Elasticity, Resource Pooling and Measured Service. Details of each characteristic are provided in Section 8.1.2.

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 the capabilities, without giving up the investments that have already been made, Participating States and Entities can accelerate their 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 of 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 existing on-premises equipment. With Amazon VPC, users can extend their on-premises network configuration into virtual

private networks on the AWS cloud. AWS resources can operate as if they are part of the existing corporate network. 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.

- **Integrated Cloud Backups:** AWS helps simplify the backup and recovery environment for the enterprise. Users can leverage the on-demand nature of the cloud and automate their 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 users can feel confident their backups are protected.

- **Integrated Network Connection:** On-premises connection with AWS is best accomplished with AWS Storage Gateway, a software appliance installed in the 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 users 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 of the data encrypted in the Amazon Simple Storage Service (Amazon S3) or Amazon Glacier.

The **AWS Storage Gateway** is a service connecting an on-premises software appliance with cloud-based storage to provide seamless and secure integration between an organization's on-premises IT environment and AWS's storage infrastructure. The service allows users to securely store data in the AWS cloud for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all of the data encrypted in 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 AWS cloud resources. It's likely that many of the tools an organization is using to manage its on-premises environments can be extended to include AWS as well. Integrating the AWS environment can provide a simpler and quicker path for cloud adoption, because an operations team does not need to learn new tools or develop completely new processes.

Microsoft:

Infrastructure-as-a-Service (IaaS)

Microsoft offerings for SaaS offers a variety of online services to address an agency's most pressing needs:

Microsoft Business Productivity Online Suite delivers a suite of services for hosted communication and collaboration. Dedicated cloud offerings for U.S. government organizations can deliver integrated communications with high availability, comprehensive security, and simplified IT management.

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

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

Platform-as-a-Service (PaaS)

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

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

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

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

Software-as-a-Service (SaaS)

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

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

Microsoft Office Skype for Business delivers hosted web conferencing.

Microsoft Exchange Hosted Services are attached services that include filtering, archiving, encryption, and continuity.

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

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

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

Many scenarios need to implement a blend of Azure offerings to meet the needs of their organization and application requirements.

Private Cloud: A private cloud delivers cloud services on resources dedicated to clients, either on-premises, such as within their own data center, or in a partner's hosting facility.

Control and customization. Dedicated resources offer more control over the level of security, privacy, customization, and governance of the software and services than does a public cloud.

Government Community Cloud: Data segregation for Government Community Cloud, when provisioned as part of Office 365 Government, the following services are offered in accordance with the National Institute of Standards and Technology (NIST) Special Publication 800-145:

- Exchange Online
- Exchange Online Archiving
- SharePoint Online (includes Project Online, Access Online and Office Delve)
- Skype for Business Online

Microsoft refers to this offer as the Government Community Cloud.

In addition to the logical separation of customer content at the application layer, each of these Office 365 services provides an organization with a secondary layer of physical segregation for customer content by using infrastructure that is separate from the infrastructure used for commercial Office 365 customers, including by using Azure services in Azure's Government Cloud.

Public Cloud: Microsoft Azure public cloud is a growing collection of integrated cloud services, analytics, computing, database, mobile, networking, storage, and web.

Hybrid Cloud: Microsoft Azure supports Hybrid cloud and it combines on-premises hosting with applications on demand. Implementers use it to build solutions that keep federal agency data behind their firewall but that also allow access to computing, storage, and application services via the cloud

8.2 (E) SUBCONTRACTORS

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

Insight Response: Insight intends to provide cloud solutions through a combination of directly delivered services via Insight resources and through the use of Subcontractors.

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

Insight Response: Insight is partnering with Amazon Web Services and Microsoft to provide the cloud solutions.

Insight is partnering with REAN Cloud to provide services in support of the AWS cloud infrastructure. REAN will be engaged to provide the following services if requested by the Participating Entity.

Strategy Phase - SaaS	Assessment Phase - SaaS	Operations Phase - SaaS	DevOps Phase - PaaS
ROI & Business Case Justification (Activity) AWS Calculator (Task) Cloud Rationalization/Adoption strategy DR & Business continuity planning DevOps Strategy Account Management Governance & Compliance	Cloud Architecture Security & Risk Assessment Migration and Implementation Phase Secure Infrastructure Setup Lift & Shift Migration (CloudEndure) DevOps based migration	Managed Services (MGS) Billing as Service (BaaS) AWS Infrastructure (IaaS)	Infrastructure Automation Application Reengineering Native AWS Application Development

REAN has a rich talent of engineers who cover broad range of skills from Software Development, Network and Security Architecture, AWS and DevOps Architecture.

Below is a breakdown of REAN Employee profile:

- Total Number of Employees – 85 (US – 55, India 30)
- Engineers/Technical team – 73
 - o Architects - 19
 - o AWS and DevOps Engineers – 46
 - o PM/Scrum Masters/Technical Writers – 8

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

Insight Response:

REAN since its inception in 2013 has been working with SME's and Enterprises alike leveraging the AWS Cloud technology focusing especially on regulated markets including Financial Services and Healthcare industries.

REAN has quickly grown (in less than 2 years) to an AWS APN Premier Consulting Partner and a leading provider of end-to-end Cloud IT Solutions ranging from business case justification, ROI analysis, migration services, native cloud implementations to 24x7 managed services. REAN specializes in supporting highly regulated industries including solutions for the following verticals: Financial Services, Healthcare/Life Sciences, Education, and Government. REAN Cloud is a cloud-native firm with deep experience supporting legacy enterprise IT infrastructures and applications. REAN Cloud provides Consulting Services around Strategy, Systems Architecture, Cloud Migration, Custom Cloud-Based Solutions, DevOps and Managed Services (MGS). REAN Cloud offers a Secure Managed Services framework, which handles end-user requirements in the AWS Shared Responsibility Model.

Over the course of the 2+ years, REAN has not only become an AWS Premier Consulting Partner but also gained few sought after AWS-designated competencies in DevOps and Life Sciences. REAN has also been awarded one of the AWS prestigious Learn and be Curious awards highlighting their prowess in adapting to the changing technology trends and growing their employees to 100+ highly skilled and talented teams in a short period.

One of their key differentiators is their experience in implementing complex and highly scalable cloud architectures creating secure, compliant operations in highly regulated industries.

REAN management team comes with an enterprise background across a wide range of industry verticals including former AWS employees, Government, Life Sciences, Telecom, ISV's, Financial Services and Big5 Consulting. Following is a quick snapshot of REAN Management team.

- **Sri Vasireddy**, Managing Partner, REAN, was the first public sector solutions architect for AWS. In this capacity, he has helped the first AWS public sector customers such as Recovery.gov and Treasury.gov go through their FISMA and FedRAMP programs, which has paved the path for many government and enterprise customers to meet their compliance needs on AWS. Prior to joining AWS, Sri has supported Centers for Medicare/Medicaid, Defense Information Systems Agency and General Services Administration on their cloud security programs.
- **Sekhar Puli**, Managing Partner, responsible for leading Sales, Marketing and Global Operations, is a seasoned business leader with 20+ years of experience, operating across multiple cultures and geographically dispersed teams in Europe, Australia and Asia; driving coordination and alignment between global teams to deliver business success. Sekhar has effectively built Consulting Practices as well as managed and delivered Enterprise class solutions during his 20+ years' career spanning Financial, Information Technology, Healthcare, Non-Profit and Telecom domains. Most recently Sekhar was with Amdocs for 10+ years holding several senior executive level positions managing P&L's of \$150M+.

-
- **Sean Finnerty**, Executive Director of Life Sciences and Compliance, led the cloud and security initiatives at Merck. Sean brings an immense knowledge in the compliance and validation arena that are unique to life sciences industry such as CFR 11 and GXP.
 - **Ben Butler** serves as the Vice President of Business Development and Solutions Architecture for REAN Cloud and has a passion for helping organizations of all sizes drive innovation into their products and processes by enabling customers to take advantage of the cloud through REAN's professional and managed services. Prior to REAN, Ben Butler was the Global Senior Marketing Manager for Big Data and High Performance Computing solutions at AWS, executed on the Amazon strategy of building broad use of cloud computing for big data and HPC workloads through speaking events, customer and partner engagements, marketing campaigns, and sales enablement tools. Ben also was a Senior Solutions Architect in the World Wide Public Sector team supporting customers with big data projects such as the NIH, SEC, FINRA, and the Department of Health and Human Services, winning Solution Architect of the Year for Worldwide Public Sector in 2012.

8.3 (E) WORKING WITH PURCHASING ENTITIES

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

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

Insight Response: Insight has provided detail on how our CSP partners will work with Purchasing Entities before, during, and after a data breach.

AWS: 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 to ensure 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:
 - a) 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.
 - b) Trouble ticket entered by an AWS employee
 - c) 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 the AWS Event Management Tool system to start the engagement and page relevant program resolvers (e.g. Security team). The resolvers will perform an analysis of the incident to determine if additional resolvers should be engaged and to determine the approximate root cause.
- 2) **Recovery Phase** - the relevant resolvers will perform break fix to address the incident. Once troubleshooting, break fix and affected components are addressed, the call leader will assign next steps in terms of follow-up documentation and follow-up actions and end the call engagement.
- 3) **Reconstitution Phase** - Once the relevant fix activities are complete the call leader will declare that the recovery phase is complete. Post mortem and deep root cause analysis of the incident will be assigned to the relevant team. The results of the post mortem will be reviewed by relevant senior management and relevant actions such as design changes etc. will be captured in a Correction of Errors (COE) document and tracked to completion.

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

available and maintained by the customer support team to alert customers to any issues that may be of broad impact.

The AWS incident management program is reviewed by independent external auditors during audits for their 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).

Microsoft: As the provider of the Cloud services that make up this solution, and operator of the datacenters that provide those services, Microsoft responds to data breaches, and provides resolution directly to their subscribers. Insight is not involved, except in an advisory role. Microsoft describes how such events are managed on its Trust Center websites for Azure and Office 365. This may change from time to time as Microsoft refines its processes and service levels.




Figure 13: AWS Online Service Health Dashboard

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

Insight Response: This requirement is not applicable to Insight as a Value Added Reseller. Provided below is how AWS addresses this requirement.

AWS: AWS services are provisioned on-demand by the customer; this is the passive nature of IaaS. The customer controls how it uses its account and what content moves onto and off of its account. AWS SOC reports (available under AWS NDA) provide additional details on the specific control activities executed by AWS to prevent unauthorized access to AWS resources.

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

Microsoft: Offeror's subcontractor, Microsoft, currently, as of the date of the Proposal, has a mechanism by which 30-day Trial subscriptions may be ordered for some, but not all, of the cloud services offered hereunder. Microsoft will provide additional information about this upon request of Lead State, Participating States, or any Purchasing Entity.

AWS: Participating States or Entities can get started quickly, with processes that are easy to repeat, through the ability to create a custom Amazon Machine Image (AMI) in Amazon Web Services. This makes sure that every developer and tester can be working with the same configuration. In addition, they can use AWS CloudFormer to take an image of the entire cloud infrastructure and create a template so they can start up exact replicas of that infrastructure for development and test.

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

Insight Response: This requirement is not applicable to Insight because the computer applications that the Participating State or Entity will access for administration of the solution is done via the CSP's website.

Microsoft: Offeror's subcontractor, Microsoft, complies with all laws applicable to it as IT service provider, but not laws applicable to a Purchasing Entity's own operations. Microsoft's research indicates that most if not all State accessibility laws (and the Federal ADA) applies to their customers (and not to Microsoft, as service provider), so Offeror respectfully takes exception with this clause, as written. Microsoft supports the government's obligation to provide accessible technologies to its citizens with disabilities as required by Section 508 of the Rehabilitation Act of 1973, and its state law counterparts (including applicable California provisions). Offeror encourages Purchasing Entities to judiciously compare product accessibility performance. The Voluntary Product Accessibility Templates ("VPATs") for the Microsoft technologies used in providing the online services can be found at Microsoft's VPAT page.

AWS: 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. 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 Voluntary Product Accessibility Template (VPAT) upon request.

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

Insight Response: This requirement is not applicable to Insight because the computer applications that the Participating State or Entity will access for administration of the solution is done via the CSP's website.

Microsoft: For each of the Microsoft cloud services offered by Insight, to the extent the services deliver content through Web browsers, our subcontractor Microsoft, generally endeavors to ensure compatibility with the latest versions of the most popular browsers including Internet Explorer, Firefox, Chrome and Safari. As of the date of Insight's Proposal, each of these are supported. However, Insight respectfully declines to commit to any requirement that would constrain Microsoft's ability to evolve its services to meet market needs. Over the 10 year term of the Master Agreement, it is likely that browser technology will change, and Microsoft will make decisions (independent of contractual commitments) as to how it will support future versions of these browsers. Insight addresses this in the exceptions to the contract terms and conditions.

AWS: An end customer can access the AWS console via all current releases of browsers so long as that is how Insight makes the console available to the end user.

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

Insight Response: Microsoft and AWS are the providers of the Cloud services that make up this solution, and while they do provide information on how its customers' information is stored, moved and kept secure, they do not hold meetings with customers for this purpose. Insight will be happy to hold a meeting to discuss how Microsoft and AWS will manage the Purchasing State or Entity's information, but we do not store, migrate, or use this information ourselves.

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

Insight Response: When Insight service resources are involved in the implementation of the cloud solution, our team of certified professionals adhere to the following project delivery methodology.

Assessment

- Current State
- Optimization recommendations
- AD Readiness
- Topology and Distribution
- Client Configuration
- Application dependencies

Design

- Change Management and Planning
- HA / DR
- Security Delegation
- Provisioning
- Client Access Architecture
- Coexistence

Build and Test

- Automation for efficiency and uniformity of builds
- Unit Tests: validate performance and capacity estimates
- Integration Tests: Validate configuration

Pilot

- User Acceptance Testing
- Validate processes

Production

- Project Coordination
- Field presence
- Onsite services
- Desk side
- Help Desk

Operate

- Standard Operating Procedures
- Maintenance

- Monitoring
- Knowledge transfer

8.4 (E) CUSTOMER SERVICE

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

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

Insight Response: During the contract transition/implementation phase, Insight will work with the NASPO ValuePoint organization and the State of Utah, the Lead State, to define an escalation process, including escalation paths, and related communications plan as they pertain to the terms and conditions of the Master Agreement. Upon the signature of every new Participating Addendum with Participating Entities, Insight will work with the individual Entity to define an escalation process, including escalation paths, and related communications plan as they pertains to the terms and conditions of the Participating Addendum.

Insight has provided an overview of our standard escalation processes and methodologies below.

Escalation Process: Insight's escalation process incorporates personnel from all areas of our business. Our goal in doing business with Participating Entities is to see that all their business requirements are being addressed across their organization. Insight is flexible in working with our clients to create programs that eliminate concerns regarding the Insight-NASPO VP Cloud Solutions partnership. We work diligently to minimize any issues that may arise to ensure we meet our service level goals. Service level issues are addressed promptly through our escalation path and issue resolution process. Insight will work with NASPO ValuePoint, the State of Utah, and Participating Entities to define and implement mutually agreed upon issue escalation and resolution procedures and processes based on the business awarded from this RFP initiative.

Escalation Path: The following escalation path has been established should a Participating Entity experience a lack of expected service. NASPO ValuePoint is encouraged to contact Pam Potter, Contract Manager, so the proper resolution can be achieved in a timely manner. Issues that are not resolved in a suitable timeframe will be escalated to the appropriate Insight Management team and a resolution plan with timetables and measurable improvement targets will be created as needed. Insight's Sales Operations Management team tracks client concerns regarding Insight Account Team personnel. Insight's Sales Operations Managers conduct regular meetings to discuss and resolve serious topics related to team personnel, client issues, etc.

Issue Resolution Process: Insight has implemented a client service initiative that is monitored by our Sales Operations Management Team and other internal Operations departments to collaborate and quickly and effectively correct any issues that may arise and ensure ongoing client satisfaction. Departments within Insight continuously measure critical customer service factors and recognize individuals and teams based on stringent quality and client service measurements. All issues are tracked and discussed during regularly scheduled meetings to minimize repeat occurrences. In the event SLAs are missed, Insight will perform a root cause analysis and institute a corrective action plan to rectify the issue including but not limited to the following processes:

- (i) Investigate and report on the causes of the problem, including performing a root cause analysis of the problem

- (ii) Advise our clients of the status of remedial efforts being undertaken with respect to such problem
- (iii) Minimize the impact of and correct the problem and begin meeting the Performance Standard
- (iv) Take appropriate preventive measures so that the problem does not recur

Problem Prevention: Insight recommends establishing a standard agenda-driven recurring meeting at a pre-determined time (i.e., weekly, bi-weekly, or monthly) between the Insight Account team and the day-to-day NASPO ValuePoint/State of Utah stakeholders. This meeting will cover all on-going activities around Cloud Solutions, open or upcoming projects, and our performance and product standards. Insight also recommends establishing a Quarterly Business review calendar to review activities with the extended Insight and NASPO ValuePoint/State of Utah management teams.

Service Level Agreements: Due to the vast scope of offerings potentially available through the NASPO ValuePoint Cloud Solutions, there are many SLAs available depending upon the solution chosen by the Participating Entity. Insight understands the importance of establishing and maintaining SLA objectives. During the signing of the Participating Addendum and initial planning phase of the contract, our Contract Manager and the Entity will define the performance standards that will be required in order to provide the services contracted to deliver. Our team will remain dedicated to ensuring we maintain any and all established SLA objectives.

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

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

Insight Response: Insight will be able to fully meet this requirement. When any end user becomes an IPS client, the Participating State or Entity is assigned a dedicated account team to manage its technology needs. This team approach ensures that someone who is familiar with the account is always available for personalized attention and service.

IPS account teams consist of highly tenured and technically proficient people, dedicated to the markets they serve. As a Participating State or Entity's trusted advisor, IPS will work closely with Participating Entities in the field to examine the issues face to face. Insight's account management model integrates an expansive network of field sales representatives with inside sales personnel in strategically located operations centers around the country. Our account teaming approach ensures our clients have the support of the experts they need for hardware, software and services.

Outlined below are support options offered by Insight's CSP partners.

Microsoft: Microsoft offers Microsoft Azure Premier Support in markets where Microsoft Azure is supported. Some specific services may not be covered in all regions immediately after General Availability (GA). Unlimited 24x7 technical support, Unlimited 24x7 billing & subscription support, Escalation management, minimum response times.

Another option for customers to consider is Azure Rapid Response in addition to the Azure Premier Support Contract.

Business applications and platforms which leverage public cloud services require around the clock availability, lightning fast user response, and resiliency to handle unforeseen events. Microsoft Cloud includes these capabilities today.

AWS: Partnering with Insight for a Participating State or Entity's AWS solutions gives access to AWS Support, 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 utilize the products and features provided by AWS.

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

Insight Response: Customer Service Representatives will be available by phone, email, and web during the minimum mandated days and time listed in the RFP. Our national sales presence allows Insight to provide support to our customers from 5AM – 8PM. However, due to the critical nature of some solutions, Insight will identify support who will be available 24/7 for emergencies.

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

Insight Response: We answer 90% of calls within 60 seconds, respond to email and voicemail within 2 business hours, and respond to 90% of quotes or answer questions for standard products within 4 business hours. The dedicated Participating State/Insight Account Team will address each inquiry on a case-by-case basis and engage the appropriate resources to assist in timely responses. Resources that may be engaged include our Field Services group, our Technology Practices experts, on-site manufacturers' representatives, and/or sales management and executive management. Insight's goal is to ensure that all Participating Entities' inquiries and business requirements are met to the satisfaction of the Entity.

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

Insight Response: Insight also offers additional services to provide a turn-key on-boarding experience migrating information, data, and workloads into the cloud. Insight offers envisioning workshops, pre-sales assessments, on-boarding, project planning and project management, pilot/POC engagements coupled with migration, integration, and greenfield deployment services.

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

Insight Response: Insight also offers additional services to provide a turn-key on-boarding experience migrating information, data, and workloads into the cloud. Insight offers envisioning workshops, pre-sales assessments, on-boarding, project planning and project management, pilot/POC engagements coupled with migration, integration, and greenfield deployment services.

8.5 (E) SECURITY OF INFORMATION

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

Insight Response: Insight has provided responses describing the measures our CSP partners and services partner take to protect data.

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

1. Customers continue to own their data.

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

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

Data Recovery/ Transfer

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

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

Deleting Data

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

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

Archiving Data

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

AWS Storage Device Decommissioning

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

decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.

REAN Cloud:

REAN Secure Virtual Private Cloud (S-VPC) on AWS

REAN has devised a secure virtual private cloud (S-VPC) framework that provides assurance of information protection with additional security controls to ensure the confidentiality, integrity and availability of information.

REAN S-VPC provides distinctive data protection for information stored on elastic block store volumes using encryption with key management system that enables policy based restrictions to determine where and when encrypted data can be accessed.

In addition, server validation applies identity and integrity rules when servers request access to secure storage volumes. This solution ensures that encryption keys are delivered to valid devices without the need to deploy an entire file system and management infrastructure. This solution protects sensitive information from theft, unauthorized exposure, or unapproved geographic migration to other data centers.

Microsoft: Microsoft believes that their customers should control their own data whether stored on their premises or in a cloud service. Accordingly, they will not disclose Customer Data to a third party (including law enforcement, other government entities or civil litigants) except as their customers direct them or as required by law.

Should a third party contact them with a demand for Customer Data, they will attempt to redirect the third party to request it directly from their customers. As part of that, they may provide customers' basic contact information to the third party. They require a court order or warrant before they will consider disclosing content to law enforcement. If compelled to disclose Customer Data to a third party, they will promptly notify the customer and provide a copy of the demand to them, unless legally prohibited from doing so.

Microsoft also publishes a Law Enforcement Requests Report that provides insight into the scope and number of requests.

In the Microsoft Cloud, Participating States and Entities are the owner of their customer data. Customer data is defined as all data, including text, sound, video, or image files and software that is provided to Microsoft, or is provided on the Participating State or Entity's behalf, through use of the enterprise online services that make up the Microsoft Cloud.

Microsoft will use the customer data only to provide the services have agreed upon, and for purposes that are compatible with providing those services. They do not share the data with their advertiser-supported services, nor do they mine it for marketing or advertising.

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

Insight Response: This requirement is not applicable to Insight since we do not own the technology infrastructure; however, we have provided responses explaining how our CSP partners and services partner can comply with all applicable laws and related to data privacy and security.

REAN Cloud: REAN provides secure, compliant Cloud services for the most highly regulated industries including Healthcare, Life Sciences and Financial sectors. These customers often require HIPAA and PCI compliance audit support from REAN to help them comply with laws related to data privacy and security.

REAN has devised a secure virtual private cloud (S-VPC) framework that provides assurance of information protection with additional security controls to ensure the confidentiality, integrity and availability of information. The S-VPC wraps the customer application in a secure shell to meet the internal governance and ensure compliance with regulations like SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 type II), PCI DSS Level 1, ISO 27001, HIPAA, HITECH, and FedRAMP.

Figure below shows the high level architecture for REAN S-VPC. The following sections explain virtual network, server, storage, access control, and audit controls in further detail.




Figure 14: High Level Architecture for REAN S-VPC

Network Protection

REAN S-VPC protects the network perimeter by creating a Demilitarized Zone (DMZ) with a unified threat management suite. The suite provides firewall services, intrusion protection/detection services, secure Virtual Private Network (VPN) connectivity, packet filtering, and web application firewall protection not available via AWS standard offerings. This front-end protects against denial-of-service attacks, worms, and hacker exploits.

Server Protection

REAN S-VPC offers comprehensive server security designed to protect all the AWS instances in the customer environment from data breaches and business disruptions, and achieve cost-effective compliance across these environments. Tightly integrated modules including anti-malware, web reputation, firewall, host based intrusion prevention, integrity monitoring, and log inspection expand the security posture to ensure server, application, and data security across physical, virtual, and cloud environments. The solution also features FIPS 140-2 certification to support high security standards.

Storage Protection

REAN S-VPC provides distinctive data protection for information stored on elastic block store volumes using encryption with key management system that enables policy based restrictions to determine where and when encrypted data can be accessed. In addition, server validation applies identity and integrity rules when servers request access to secure storage volumes. The solution ensures that encryption keys are delivered to valid devices without the need to deploy an entire file system and management infrastructure. This solutions protects sensitive information from theft, unauthorized exposure, or unapproved geographic migration to other data centers.

Access Control

REAN S-VPC environment provides various convenient options to the end users to access the environment and initiate their VPN connections. These include:

- HTML5 based remote access VPN that they can initiate from any HTML5 compatible browser with requiring any plug-in.
- SSL remote access VPN that provides additional security by a double authentication using X.509 certificates and username/password.
- IPSec based VPN using native Windows or Mac VPN clients
- Mobile VPN using native iPhone VPN client to securely connect to VPC

System administrator access control is provided through the integration of GU identity and access management solution. This suite supplements the AWS Management Console by vaulting administrator's credentials, enforcing separation of duties, and recording all accesses and actions.

Logging and Auditing

REAN S-VPC ensures that the customer environment is continuously monitored using auditing at the network, server, and application levels to help meet all the forensics and compliance requirements. In case of server and infrastructure access, the solution not only provides system logs but could optionally provide full video stream of an administrator session into Amazon S3. By providing such video stream that is tied back to customer Identity and Access Management (IAM), enterprises can maintain full accountability for any changes performed on the service. All the above audit data is fed into a Security Information and Event Management (SIEM) system that provides full contextual awareness of the events that can be summarized in a simple dashboard.

Availability

The customer environment is architected to take full advantage of highly available AWS infrastructure. All the components (application servers and files stores) of the solution are deployed in a redundant fashion across multiple fault isolated AWS Availability Zones. 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. The file store uses Amazon Simple Storage Service (S3) service that provides eleven 9s SLA on durability of the customer's data.

REAN S-PVC Value

REAN S-VPC has successfully passed security testing and auditing by a leading auditor that provider that servers the Department of Defense. Customer can adopt a proven and working framework and save time and money.

Microsoft: Refer to the Microsoft answer above.

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

Insight Response: Insight will not have access to a Purchasing Entity's user accounts or data.

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

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

Microsoft: Refer to the Microsoft answer above.

8.6 (E) PRIVACY AND SECURITY

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

Insight Response: Insight has described how the CSP partners and our service partner are committed to complying with NIST.

AWS: AWS provides NIST compliant cloud infrastructure services. AWS's compliance is validated by two Agency Authority to Operate (ATOs) achieved based on testing performed against the stringent set of FedRAMP requirements (NIST 800-53 Rev. 4 – Moderate baseline requirements, plus additional FedRAMP security controls). AWS provides federal security personnel with their security documentation as a means of verifying the security and compliance of AWS in accordance with applicable NIST controls as defined by 800-53 rev4 and the DoD Cloud Computing Security Requirements Guide (SRG).

REAN Cloud: REAN has devised a secure virtual private cloud (S-VPC) framework that provides assurance of information protection with additional security controls to ensure the confidentiality, integrity and availability of information. The S-VPC wraps the customer application in a secure shell to meet the internal governance and ensure compliance with regulations like SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 type II), PCI DSS Level 1, ISO 27001, HIPAA, HITECH, and FedRAMP.

REAN S-VPC offers comprehensive server security designed to protect all the AWS instances in the customer environment from data breaches and business disruptions, and achieve cost-

effective compliance across these environments. Tightly integrated modules including anti-malware, web reputation, firewall, host based intrusion prevention, integrity monitoring, and log inspection expand the security posture to ensure server, application, and data security across physical, virtual, and cloud environments.

The solution also features FIPS 140-2 certification to support high security standards.

In order to provide end-to-end security and end-to-end privacy, REAN architects and delivers its services on the highly compliant AWS infrastructure. AWS builds infrastructure in accordance with security best practices, provides the appropriate security features in those services and documents how to use those features. The AWS cloud infrastructure has been designed and managed in alignment with regulations, standards and best-practices including:

- Federal Risk and Authorization Management Program (FedRAMP)
- Service Organization Controls (SOC) 1/Statement on Standards for Attestation Engagements (SSAE) 16/International Standard on Assurance Engagements (ISAE) 3402 (formerly Statement on Auditing Standards [SAS] No. 70)
- SOC 2
- SOC 3
- Payment Card Industry Data Security Standard (PCI DSS)
- International Organization for Standardization (ISO) 27001
- ISO 9001
- ISO 27001
- Department of Defense Risk Management Framework (DoD RMF) Cloud Security Model (CSM)
- Federal Information Security Management Act (FISMA)
- International Traffic in Arms Regulations (ITAR)
- Federal Information Processing Standard (FIPS) 140-2
- Family Educational Rights and Privacy Act (FERPA)
- IRAP (Australia)

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

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

Insight Response: Insight has provided an answer based on the government or standard organization certifications that our CSP partners hold.

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

Microsoft: The National Institute of Standards and Technology (NIST) 800-53 controls is the standard, and FedRAMP is the program that certifies that a CSP meets that standard. FedRAMP, ISO/IEC 27001 and ISO/IEC 27018, SOC 1 and SOC2. Microsoft Azure and Microsoft Azure Government have earned a Provisional Authority to Operate (P-ATO) from the FedRAMP Joint Authorization Board; Microsoft Dynamics CRM Online Government has received an Agency ATO from HUD; and Microsoft Office 365 U.S. Government has received an Agency ATO from DHHS.

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

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

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

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

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

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

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

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

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

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

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

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

Service Organization Controls (SOC) are a series of accounting standards that measure the control of financial information for a service organization. Azure's SOC 1 and SOC 2 Type 2 audit reports attest to the effectiveness of the design and operation of its security controls. Other country-specific standards and contractual commitments, including the EU Model Clauses, UK G-Cloud, Singapore MTCS, and Australia CCSL (IRAP).

Microsoft has a responsibility to process their customers' information in a trustworthy manner, many customers have a responsibility to comply with national, regional, and industry-specific requirements governing the collection and use of individuals' data.

To give customer's the foundation to achieve that compliance, Microsoft takes a two-pronged approach to help ensure that compliance controls are current and that we build and maintain a dynamic compliance framework.

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

Insight Response: As a Value Added Reseller, this requirement is not applicable to Insight. However, we have described how our CSP partners meet this requirement.

AWS: AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools have the ability to set custom performance metrics thresholds for unusual activity.

Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to automatically notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used so personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel.

Documentation is maintained to aid and inform operations personnel in handling incidents or issues. If the resolution of an issue requires collaboration, a conferencing system is used which supports communication and logging capabilities. Trained call leaders facilitate communication and progress during the handling of operational issues that require collaboration. Post-mortems are convened after any significant operational issue, regardless of external impact, and Cause of Error (COE) documents are drafted so the root cause is captured and preventative actions are taken in the future. Implementation of the preventative measures is tracked during weekly operations meetings.

AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant telecommunication providers at each region as well as additional capacity protect against the possibility of DoS attacks.

The AWS network provides significant protection against traditional network security issues, and Participating States and Entities 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. Participating States

and Entities can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. AWS encourages Participating States and Entities to use SSL for all of their interactions with AWS.

- **IP Spoofing.** Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.
- **Port Scanning.** Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on their website at: <http://aws.amazon.com/contact-us/report-abuse/>. When unauthorized port scanning is detected by AWS, it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by customers. The strict management of security groups can further mitigate the threat of port scans. If Participating States and Entities 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, Participating States and Entities 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. Participating States and Entities may request permission to conduct vulnerability scans as required to meet their specific compliance requirements. These scans must be limited to their own instances and must not violate the AWS Acceptable Use Policy.
- **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 Participating States and Entities can place their 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 Participating States and Entities should encrypt sensitive traffic.

REAN Cloud: REAN S-VPC protects the network perimeter by creating a Demilitarized Zone (DMZ) with a unified threat management suite. The suite provides firewall services, intrusion protection/detection services, secure Virtual Private Network (VPN) connectivity, packet filtering, and web application firewall protection not available via AWS standard offerings. This front-end protects against denial-of-service attacks, worms, and hacker exploits.

REAN Security Framework includes:

Amazon Virtual Private Cloud (VPC)

Amazon Virtual Private Cloud (Amazon VPC) lets customers provision a private, isolated section of the Amazon Web Services (AWS) Cloud where members can launch AWS resources in a virtual network that they define. With Amazon VPC, users can define a virtual network topology that closely resembles a traditional network that they might operate in their own data center. REAN will help NASPO 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.

REAN will help NASPO customize the network configuration for their Amazon VPC. For example, NASPO may need a public-facing subnet for their web servers that have access to the Internet, and place their backend systems such as databases or application servers in a private-facing

subnet with no Internet access. REAN will help NASPO leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Additionally, REAN will help NASPO create a Hardware VPN connection between their corporate data center and their VPC and leverage the AWS cloud as an extension of their corporate data center. Figure below shows a notional picture of the NASPO AWS VPC infrastructure offering.




Figure 15: REAN AWS Infrastructure

AWS VPC Infrastructure Offering

A variety of connectivity options exist for Participating States or Entities to connect to their Amazon VPC: NASPO can connect their VPC to the Internet, to their datacenter, or both, based on the AWS resources that they want to expose publicly and those that they want to keep private.

- Connect directly to the Internet (public subnets) – States or Entities can launch instances into a publicly accessible subnet where they can send and receive traffic from the Internet.
- Connect to the Internet using Network Address Translation (private subnets) – Private subnets can be used for instances that the State or Entity do not want to be directly addressable from the Internet. Instances in a private subnet can access the Internet without exposing their private IP address by routing their traffic through a Network Address Translation (NAT) instance in a public subnet.
- Connect securely to a corporate datacenter – All traffic to and from instances in the State or Entity's VPC can be routed to their corporate datacenter over an industry standard, encrypted IPSec hardware VPN connection.

- Combine connectivity methods to match the needs of the State or Entity's application – Customers can connect a VPC to both the Internet and their corporate datacenter and configure Amazon VPC route tables to direct all traffic to its proper destination.

Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level. In addition, States and Entities can store data in Amazon S3 and restrict access so that it's only accessible from instances in their VPC. Optionally, NASPO can also choose to launch Dedicated Instances that run on hardware dedicated to a single customer for additional isolation.

Microsoft: Microsoft has made major investments in cloud security in the following areas.

- Design and operational security

Microsoft Cloud security begins with a trustworthy technology foundation. Microsoft designs its software for security from the ground up and helps ensure that the cloud infrastructure is resilient to attack. Microsoft uses an “assume breach” stance as a security strategy, and their global incident-response team works around the clock to mitigate the effects of any attacks against the Microsoft Cloud. These practices are backed by centers of excellence that fight digital crime, respond to security incidents and vulnerabilities in Microsoft software, and combat malware.

- Encryption

Technological safeguards, such as encrypted communications and operational processes, enhance the security of our customers' data. For data in transit, the Microsoft Cloud uses industry-standard encrypted transport protocols between user devices and Microsoft datacenters, and within datacenters themselves. For data at rest, the Microsoft Cloud offers a wide range of encryption capabilities up to AES-256, giving the flexibility to choose the solution that best meets the client's needs.

- Identity and access management

Azure Active Directory is a comprehensive identity and access management cloud solution that helps secure access to data and on-premises and cloud applications, and simplifies the management of users and groups. It combines core directory services, advanced identity governance, security, and application access management, and is a key component of Microsoft Cloud services, including Microsoft Azure, Office 365, Microsoft Dynamics CRM Online, and Intune, as well as thousands of third-party SaaS apps. Azure Active Directory also makes it easy for developers to build policy-based identity management into their applications.

- Security Development Lifecycle

Microsoft recognizes that focusing on security as a core component in the software development process can reduce the risk of costly issues, improve the security and privacy of infrastructure and applications, and protect data in the Microsoft Cloud. The SDL is composed of proven security practices that consist of multiple phases in which core software assurance activities are defined.

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

- Identity and access management – Azure Active Directory helps ensure that only authorized users can access the environments, data, and applications, and provides multi-factor authentication for highly secure sign-in.
- Encryption – Azure uses industry-standard protocols to encrypt data as it travels between devices and Microsoft datacenters, and crosses within datacenters

- Secure networks – Azure infrastructure relies on security practices and technologies to connect virtual machines to each other and to on-premises datacenters, while blocking unauthorized traffic. Azure Virtual Networks extend the on-premises network to the cloud via a site-to-site virtual private network (VPN). Participating States and Entities can also use ExpressRoute to create a cross-premises connection when needing to use the Internet.
- Threat management – Microsoft Antimalware protects Azure services and virtual machines. Microsoft also uses intrusion detection, denial-of-service (DDoS) attack prevention, penetration testing, data analytics, and machine learning to constantly strengthen its defense and reduce risks.
- Compliance – Microsoft complies with both international and industry-specific compliance standards and participate in rigorous third-party audits, which verify their security controls.

Customers maintain full ownership and control over their own data. They are a leader in providing transparency about their privacy practices—one reason they have adopted the world's first code of practice for cloud privacy, ISO/IEC 27018.

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

Insight Response: Provided below are descriptions of the data confidentiality standards and practices that are in place to ensure data confidentiality.

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

REAN Cloud: Refer the response provided for 8.6.3 for information describing REAN Cloud's data confidentiality standards and practices.

Microsoft: Where Azure data is physically stored is very important to most customers. If the organization is restricted by any government regulations or internal company policies about data storage and location, this needs to be transparent. Many times there are restrictions about data export and Government Regulatory Compliance (GRC) for some data sets. This information needs to be understood before deploying any applications or services.

Within each datacenter, the racks of equipment are built to be fault tolerant with respect to networking, physical host servers, storage, and power. The physical host servers are placed in high availability units called a cluster. The cluster configurations are spread across multiple server racks.

A single rack is referred to as a Fault Domain (FD), and it can be viewed as a vertical partitioning of the hardware. The fault domain is considered the lowest common denominator within the datacenter for fault tolerance. Microsoft Azure can lose a complete rack, and the hosted services can continue unaffected.

A second partition within the datacenter is called the Upgrade Domain (UD) and it can be viewed as a set of horizontal stripes passing through the vertical racks of fault domains. Upgrade domains are used to deploy updates (security patches) within Azure without affecting the availability of the running services within the Azure fabric. The following diagram shows a high-level relationship between fault domains and update domains in the Azure datacenters.

Microsoft has been a leader in creating robust online solutions that protect the privacy of their customers for twenty years. Today, they operate more than 200 cloud and online services that serve hundreds of millions of customers across the globe. Their enterprise cloud services, such as Office 365 and Windows Azure, serve millions of end users whose companies entrust their mission-critical data to Microsoft.

Their experience has enabled them to develop industry-leading business practices, privacy policies, compliance programs, and security measures that we apply across the cloud computing ecosystem. Driven by a commitment to empower organizations to control the collection, use, and distribution of their data, their time-tested approach to privacy provides a solid foundation for addressing customer privacy requirements and enabling greater trust in cloud computing.

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

Insight Response: Provided below are lists of third-party attestations, reports, security credentials, etc. for each of the CSP solutions represented in Insight's proposal response.

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

The figure outlines Microsoft's compliance and adherence to other standards, such as CJIS, IRS 1075, HIPAA, FERPA, ISO/IEC 27001 and 27018, SOC1 and 2, and others. Please note that some of these standards apply only to certain services (e.g. CJIS applies only their Government Community Cloud services) and that some of them require special Amendments and/or Agreements (e.g. CJIS requires that a State's CJIS Systems Agency must execute a special agreement with Microsoft, before Microsoft will provide an FBI CJIS Addendum for use in each such state).

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



Figure 16: Microsoft Cloud Services Certifications

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

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

Insight Response: Insight describes the logging process for each of our CSP partners and service partner.

AWS: 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.

AWS CloudTrail: Features and Benefits

Some of the many features of AWS CloudTrail include:

- o **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?
- o **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.
- o **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.
- o **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.
- o **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.
- o **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.

Customer 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 break down 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 (<http://www.cascading.org>) and is an example of how to construct an Amazon Elastic MapReduce application. Customers can also customize reports generated by the LogAnalyzer.

Third Party Tools

Many third-party log monitoring and analysis tools are available on AWS Marketplace.

REAN Cloud:

Logging and Auditing

REAN S-VPC ensures that the customer environment is continuously monitored using auditing at the network, server, and application levels to help meet all the forensics and compliance requirements. In case of server and infrastructure access, the solution not only provides system logs but could optionally provide full video stream of an administrator session into Amazon S3. By providing such video stream that is tied back to customer Identity and Access Management (IAM), enterprises can maintain full accountability for any changes performed on the service. The entire above audit data is fed into a Security Information and Event Management (SIEM) system that provides full contextual awareness of the events that can be summarized in a simple dashboard.

REAN Cloud MGS

For on-going support, REAN MGS includes monitoring, alerting, and automated trouble ticketing solutions to ensure timely reporting and response to fixing unhealthy infrastructure and application errors. REAN Cloud configures all applicable resources to ship logs, including Amazon CloudWatch metrics, to the central logging system backed by Splunk Enterprise. REAN Cloud MGS can also provide proactive monthly reports to check for cost optimizations, security improvement recommendations, and any remediation recommendations.

Microsoft: Customers can enable or disable the following kinds of logs:

- Detailed Error Logging - Detailed error information for HTTP status codes that indicate a failure (status code 400 or greater). This may contain information that can help determine why the server returned the error code.
- Failed Request Tracing - Detailed information on failed requests, including a trace of the IIS components used to process the request and the time taken in each component. This can be

useful if the client is attempting to increase site performance or isolate what is causing a specific HTTP error to be returned.

- Web Server Logging - Information about HTTP transactions using the W3C extended log file format. This is useful when determining overall site metrics such as the number of requests handled or how many requests are from a specific IP address.

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

Insight Response: As a Value Added Reseller, this requirement does not apply to Insight. However, we have explained how our CSP partner's and services partner are able to restrict the visibility of cloud hosted data.

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

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

REAN Cloud:

Access Control

REAN S-VPC environment provides various convenient options to the end users to access the environment and initiate their VPN connections. These include:

- HTML5 based remote access VPN that they can initiate from any HTML5 compatible browser with requiring any plug-in.
- SSL remote access VPN that provides additional security by a double authentication using X.509 certificates and username/password.
- IPSec based VPN using native Windows or Mac VPN clients
- Mobile VPN using native iPhone VPN client to securely connect to VPC

System administrator access control is provided through the integration of GU identity and access management solution. This suite supplements the AWS Management Console by vaulting administrator's credentials, enforcing separation of duties, and recording all accesses and actions.

Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level. In addition, NASPO can store data in Amazon S3 and restrict access so that it's only accessible from instances in their VPC. Optionally, NASPO can also choose to launch Dedicated Instances that run on hardware dedicated to a single customer for additional isolation.

Microsoft: Because each virtual network is run as an overlay, only virtual machines and services that are part of the same network can access each other. Services outside the virtual network have no way to identify or connect to services hosted within virtual networks. This provides an added layer of isolation to the services.

Customers can join virtual machines in Azure to the domain running on-premises. Customers can access and leverage all on-premises investments for monitoring and identity for the services hosted in Azure.

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

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

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

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

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

The lock level identifies the type of enforcement for the policy, which presently has two values – CanNotDelete and ReadOnly. CanNotDelete means authorized users can still read and modify resources, but they can't delete any of the restricted resources. ReadOnly means authorized users can only read from the resource, but they can't modify or delete any of the restricted resources

Locks can be applied using ARM templates, ARM REST API, or ARM Azure PowerShell. To create or delete management locks, the customer must have access to Microsoft.Authorization/* or Microsoft.Authorization/locks/* actions. Of the built-in roles, only Owner and User Access Administrator are granted those actions.

Every request made to an Azure Storage account must be authenticated, unless it is an anonymous request against a public container or its blobs. There are two ways to authenticate a request against the storage accounts:

- Use the shared key or shared key lite authentication schemes for the Blob, Queue, Table, and File services.
- Create a shared access signature. A shared access signature includes the credentials required for authentication and the address of the resource being accessed. Because the shared access signature includes all data needed for authentication, it can be used to grant access to a Blob, Queue, or Table service, and it can be distributed separately from any code.

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

Insight Response: Insight has described how our CSP partners and service partner provide notifications in the event of a security incident.

Microsoft: Operational Security Assurance (OSA) is an important process that Microsoft uses to make its networks more resilient to attack and increase the security of its cloud-based services. OSA helps Microsoft achieve this increased resilience and security by extending the foundation of Microsoft cloud-based services to protect against Internet-based security threats and by incorporating best practices and methodology to continuously update services to improve security and resolve incidents as quickly as possible.

The Security Development Lifecycle (SDL) is a software development process that helps developers build more secure software and address security compliance requirements while reducing development cost

AWS: 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. The "[AWS Security Center](#)" is available to provide the Participating State or Entity with security and compliance details about AWS. They can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

REAN Cloud: Please see description below provided from a recent REAN customer SOW. REAN will review customer's current AWS environments, deploy REAN monitoring agents, work with customer team to identify alerting thresholds, notification groups and make necessary changes to take over management of systems. REAN will provide support for the agreed upon customer's AWS environments at the defined support levels.

Procedures for Security Incidents

An incident is an unplanned interruption to an IT Service or reduction in the Quality of an IT Service. Failure of any Item, software or hardware, used in the support of a system that has not yet affected service is also an Incident. For example, the failure of one component of a redundant high availability configuration is an incident even though it does not interrupt service.

An incident occurs when the operational status of a production item changes from working to failing or about to fail, resulting in a condition in which the item is not functioning as it was designed or implemented. The resolution for an incident involves implementing a repair to restore the item to its original state.

Incident Management Process Flow Steps:

Role	Step	Description
Requesting Customer	➤	Incidents can be reported by the customer or technical staff through various means, i.e., phone, email, or a self service web interface. Incidents may also be reported through the use of REAN-Checks REAN-Drop Agent on specific servers
REAN Support Service Desk	➤	Incident identification Work cannot begin on dealing with an incident until it is known that an incident has occurred. As far as possible, all key components should be monitored so that failures or potential failures are detected early so that the incident management process can be started quickly.
	➤	Incident logging All incidents must be fully logged and date/time stamped, regardless of whether they are raised through a Service Desk telephone call or whether automatically detected via an event alert. All relevant information relating to the nature of the incident must be logged so that a full historical record is maintained – and so that if the incident has to be referred to other support group(s), they will have all relevant information at hand to assist them.
	➤	Incident categorization All incidents will relate to one of the published services listed in the Service Catalogue. If the customer is calling about an issue they have that is not related to one of the services in the catalogue, then it will be put into a general bucket and deemed if it is or is not an incident.
	➤	Is this actually a Service Request incorrectly categorized as an incident? If so, update the case to reflect that it is a Service Request and follow the appropriate Service Request process.
	➤	Has this issue already been reported by others?
	➤	If this is another person reporting the same issue, relate the issue to the cases already reported. More people reporting the

		same issue means the impact of the issue is broader than what might have been reported at first. The impact needs to be recorded base upon current knowledge of the impact.
	➤	Incident prioritization Before an incident, priority can be set, the severity and impact need to be assessed. Once the severity and impact are set, the priority can be derived using the prescriptive table.
	➤	Is this a priority 1 (major) incident?
	➤	If this is a priority 1 incident meaning that a service is unavailable in part or whole, all mid level and senior REAN Support management should be alerted to make certain any resources necessary to the resolution will be immediately made available.
	➤	Initial diagnosis If the incident has been routed via the Service Desk, the Service Desk analyst must carry out initial diagnosis, using diagnostic scripts and known error information to try to discover the full symptoms of the incident and to determine exactly what has gone wrong. The Service Desk representative will utilize the collected information on the symptoms and use that information to initiate a search of the Information Base to find an appropriate solution. If possible, the Service Desk Analyst will resolve the incident and close the incident if the resolution is successful.
	➤	Is the necessary information in the Information Base to resolve the incident? If not, the case should then be assigned to the provider group that supports the service.
	➤	If the necessary information to resolve the incident is not in the Information Base, the incident must be immediately assigned to an appropriate provider group for further support. The assignee will then research the issue to determine cause and remediation options.

	➤	After a possible resolution has been determined either from the Information Base or through research, attempt the resolution.
	➤	Verify with the customer that the resolution was satisfactory and the customer is able to perform their work. An incident resolution does not require that the underlying cause of the incident has been corrected. The resolution only needs to make it possible for the customer to be able to continue their work.
REAN Support Service Desk	➤	If the customer is satisfied with the resolution, proceed to closure, otherwise continue investigation and diagnosis.
	➤	Incident Closure The Service Desk should check that the incident is fully resolved and that the users are satisfied and willing to agree the incident can be closed. The Service Desk should also check the following: Closure categorization. Check and confirm that the initial incident categorization was correct or, where the categorization subsequently turned out to be incorrect, update the record so that a correct closure categorization is recorded for the incident – seeking advice or guidance from the resolving group(s) as necessary. User satisfaction survey. Carry out a user satisfaction call-back or e-mail survey for the agreed percentage of incidents. Incident documentation. Chase any outstanding details and ensure that the Incident Record is fully documented so that a full historic record at a sufficient level of detail is complete. Ongoing or recurring problem? Determine (in conjunction with resolver groups) whether it is likely that the incident could recur and decide whether any preventive action is necessary to avoid this. In conjunction with Problem Management, raise a Problem Record in all such cases so that preventive action is initiated. Formal closure. Formally close the Incident Record.

Escalation

According to ITIL standards, although assignment may change, ownership of incidents always resides with the Service Desk. As a result, the responsibility of ensuring that an incident is escalated when appropriate also resides with the Service Desk.

The Service Desk will monitor all incidents, and escalate them based on the following guidelines:

Priority	Time Limit before Escalation	
3 - Low	3 business days	Manager
2 - Medium	4 hours	Manager
	If on-call contact cannot be reached during non-business hours	Manager
	If neither on-call contact or their manager cannot be reached during non-business hours	Senior Mgt
	48 hours	Senior Mgt
1 - High	Immediate	Manager
	Immediate	Senior Mgt

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

Insight Response: Insight has described whether our CSP partners and service partner has any security controls.

AWS: Amazon Virtual Private Cloud (Amazon VPC) lets the Participating State or Entity provision a logically isolated section of the Amazon Web Services (AWS) Cloud where they can launch AWS resources in a virtual network that they define. They 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.

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

REAN Cloud:

Administrative Access, Control of Provisioning

This section describes the access controls REAN helps implement for AWS customers to meet compliance requirements. These include:

1. Cloud infrastructure access
2. Privileged User (OS/DB admin) access

3. End user (application) access

Cloud Infrastructure Access

Access to Cloud infrastructure entails access to AWS resources that include virtual server (EC2 instance), virtual storage (EBS volume), virtual network (routing tables and firewall rules), and other AWS resources. AWS provides two types of access to provision and manage these resources.

1. AWS Console based access
2. AWS API based access

The following subsections describe the two methods of access and how REAN helps customers secures the access.

AWS Console Access

AWS provides a web console based access control to provision cloud resources. This access is protected by using a two-factor authentication method that includes a password and a soft-token (Google Authenticator) generated one time access key.

The picture below shows how REAN further secures this access by using another level of two factor authentication mechanism to access AWS console.




Figure 17: AWS Console Access Diagram

REAN provisions an AWS console (AWS IAM user) login account for each user that has admin privileges and one read-only account for service desk personnel that perform initial triage. These credentials are shared only through REAN credential management tool LastPass. Each REAN support person would have to first access their LastPass account using another two factor authentication (password and Google authenticator based one-time key) and gain access to their

AWS credentials and then login to the AWS account. LastPass access is logged to a LastPass log and AWS/IAM access is logged to AWS CloudTrail, which is then forwarded to Splunk (log monitoring and auditing tool) for analysis and alerting of malicious login attempts. This process is further defined in the logging and auditing procedure.

This will be further enhanced with the additional security controls REAN will implement above the hypervisor as the project progresses.

Microsoft: Complication with attempting to use traditional network-based security controls exclusively is that most of these controls assume the IP address is a good proxy for machine or service identity.

IP addresses are a poor proxy for identity outside of a corporate LAN that is using static assignments, particularly in a globally scaled Internet service such as Azure where IP addresses change rapidly. This typically creates significant challenges for organizations that are overly reliant on network security measures and are using static IP addresses for server and service mapping.

Review the guidance in the Microsoft Azure Security section (specifically the Containment and Segmentation Strategy) for how to design complete security containment strategies that overcome the limitations of networking controls alone.

Virtual Appliances are third-party-based virtual machine solutions that can be selected from the Azure Gallery or Marketplace to provide services like network firewall, application firewall and proxy, load balancing, and logging.

As organizations move workloads to the cloud, they must address threats in new ways and shed legacy security practices that often have proven to be ineffective and burdensome. In some cases, extending to the cloud provides an opportunity to implement security controls and contain adversaries in ways that are more challenging to accomplish in existing on-premises environments. Although containment strategies are not new, the traditional network-centric approach has failed in several ways and needs to be updated.

This section defines the following terminology:

- Containment strategy - High-level strategic approach designed to limit the risk and scope of any given compromise
- Segmentation strategy - Component of the containment strategy that separates computing assets into security zones that reflect significantly different asset valuation, trust levels, and/or risk exposure profiles
- Security zone - Set of computing assets with a common asset valuation, trust level, and/or risk exposure profile.

The notions of containment and segmentation have been around for a long time in IT security, though the interpretations of how to implement them have varied in practice. This document starts with an assume breach mindset and calls for designing security controls to prevent propagation of breaches among enterprise assets.

This requires architects and system designers to look at what a breached system or compromised account means to the environment so as to limit the impact of that breach, to make it detectable, and to enable the organization to respond.

This assume breach approach complements the traditional perimeter approach focused on preventing breaches for a combined approach that results in a more resilient strategy.

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

Insight Response: Insight has provided documentation that will be useful to NASPO ValuePoint and the State of Utah in understanding the CSP partner and service provider's security architectures.

AWS: AWS has two publicly available whitepapers that are helpful in providing answers to this question. Both of these white papers have been submitted with the proposal response. Abstracts for each whitepaper are provided below.

Whitepaper 1: Architecting for the AWS Cloud Best Practices

Abstract:

This whitepaper is intended for solutions architects and developers who are building solutions that will be deployed on Amazon Web Services (AWS). It provides architectural patterns and advice on how to design systems that are secure, reliable, high performing, and cost efficient. It includes a discussion on how to take advantage of attributes that are specific to the dynamic nature of cloud computing (elasticity, infrastructure automation, etc.). In addition, this whitepaper also covers general patterns, explaining how these are evolving and how they are applied in the context of cloud computing.

Whitepaper 2: Managing Your AWS Infrastructure at Scale

Abstract:

Amazon Web Services (AWS) enables organizations to deploy large-scale application infrastructures across multiple geographic locations. When deploying these large, cloudbased applications, it's important to ensure that the cost and complexity of operating such systems does not increase in direct proportion to their size.

This whitepaper is intended for existing and potential customers—especially architects, developers, and sysops administrators—who want to deploy and manage their infrastructure in a scalable and predictable way on AWS.

In this whitepaper, we describe tools and techniques to provision new instances, configure the instances to meet your requirements, and deploy your application code.

We also introduce strategies to ensure that your instances remain stateless, resulting in an architecture that is more scalable and fault tolerant. The techniques we describe allow you to scale your service from a single instance to thousands of instances while maintaining a consistent set of processes and tools to manage them.

For the purposes of this whitepaper, we assume that you have knowledge of basic scripting and core services such as Amazon Elastic Compute Cloud (Amazon EC2).

Provided is a sample DR reference architecture for local applications.

REAN Cloud: REAN Cloud has been reselling AWS IaaS and providing Managed Services since 01 JAN 2014. REAN Management and staff members have been architecting and managing solutions in AWS since 2010.

AWS offers a broad set of global compute, storage, database, analytics, application, and deployment services, all of which are listed at: <http://aws.amazon.com/products/>. The following figure is a simple view of the set of services that AWS offers. AWS offerings are provided with a range of supporting components like management tools, networking services, and application augmentation services, with multiple interfaces to AWS API-based services, including [SDKs](#), [IDE Toolkits](#), and [Command Line Tools](#).

All AWS products are hosted within the AWS' global data center footprint that allows the Participating State or Entity to consume services without having to build or manage facilities or equipment.




Figure 18: AWS DR Reference Architecture Example




Figure 19: AWS Global Data Center

AWS Cloud Services Approach

A core AWS offering is Amazon Elastic Compute Cloud (Amazon EC2), a web service that provides resizable compute capacity in the cloud. Its simple interface allows customers to obtain and configure capacity with minimal friction, providing complete control of computing resources. AWS is a language and operating system agnostic platform and customers receive a virtual environment with the choice of operating system, programming language, web application platform, database, and other services needed.

AWS cloud services are optimized to scale up to the demands of millions of users across the Internet. For example, Amazon S3 holds trillions of objects and regularly peaks at 1.5 million requests per second. In terms of computing capacity, according to the most recent Gartner Magic Quadrant report on Infrastructure as a Service (IaaS), *"It [AWS] is the overwhelming market share leader, with more than five times the cloud IaaS compute capacity in use than the aggregate total of the other 14 providers in this Magic Quadrant."*

Figure below shows the high level architecture for REAN S-VPC. The following sections explain virtual network, server, storage, access control, and audit controls in further detail.




Figure 20: High Level Architecture for REAN S-VPC

As cloud computing customers are building systems on top of cloud infrastructure, the security and compliance responsibilities are shared between the Cloud Service Provider (CSP) and cloud customers or partners. In an Infrastructure as a Service (IaaS) model, customers control how they architect and secure their applications and data put on the infrastructure, while CSPs are responsible for providing services on a highly secure and controlled platform, providing a wide array of additional security features.

Microsoft: Microsoft provides security for the IaaS, SaaS and PaaS services that make up its Cloud Solutions. While it does not provide a public Technical Reference Architecture, it does provide the following information about its methods for keeping these services and customer data secure.

Microsoft uses multiple safeguards to protect customer and enterprise data. These **security practices and technologies** include:

- Identity and access management – Microsoft's Directory Synchronization Services and Azure Active Directory helps ensure that only authorized users can access the Participating State or Entity's environments, data, and applications, and provides multi-factor authentication for highly secure sign-in.
- Encryption – Microsoft uses industry-standard protocols to encrypt data as it travels between devices and Microsoft datacenters, and crosses within datacenters
- Secure networks – Microsoft's Cloud infrastructure relies on security practices and technologies to connect virtual machines to each other and to on-premises datacenters, while blocking unauthorized traffic. Azure Virtual Networks extend their on-premises network to the cloud via a site-to-site virtual private network (VPN). The Participating State or Entity can also use **ExpressRoute** to create a cross-premises connection when needing to use the Internet.
- Threat management – **Microsoft Antimalware** protects Microsoft Cloud services and virtual machines. Microsoft also uses intrusion detection, denial-of-service attack prevention, penetration testing, data analytics, and machine learning to constantly strengthen its defense and reduce risks.
- Compliance – Microsoft complies with both international and industry-specific compliance standards and participate in rigorous third-party audits, which verify its security controls.

Customers maintain full ownership and control over their own data. Microsoft is a leader in providing transparency about its privacy practices—one reason they have adopted the world's first code of practice for cloud privacy, ISO/IEC 27018.

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

Insight Response: Insight requires applicants to undergo a series of steps prior to being made an offer of employment. This includes a completed job application, right to work documentation, background investigation, and drug test. As a standard practice, Insight conducts pre-hire background checks and drug testing on all new employees. Background checks and drug testing are performed after the candidate accepts an offer with us and prior to their official employment start date.

The background checks consist of the following:

- County Criminal Checks for felonies and misdemeanors going back as far as the State allows
- Social Security Trace
- Address Verification
- Motor Vehicle Registration (MVR) Check
- Terrorist Watch List (Office of Foreign Assets Control – OFAC)

Insight will conduct current background checks on any Insight or subcontractor personnel who work in an Entity owned/leased/rented facility, and provide proof and results of those background checks to the Entity.

AWS: AWS operates under a shared security responsibility model, where AWS is responsible for the security of the underlying cloud infrastructure and customers are responsible for securing the workloads they deploy in AWS. AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS conducts pre-employment criminal background checks, as permitted by law, for employees

commensurate with their position and level of access. The AWS SOC reports provides additional details regarding the controls in place for background verification.

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

Insight Response: As a Value Added Reseller, this requirement is not applicable to Insight. However, we have described our CSP partner's security measures.

AWS: AWS offers the Participating State or Entity the ability to add a layer of security to their data at rest in the cloud, providing scalable and efficient encryption features. These include:

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

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

Microsoft: Outlined below is Microsoft's security measures and standards that are in place for data at rest and in transit.

Data at Rest: Performed by the customer by encrypting the virtual hard disk (VHD) files. Microsoft and third-party mechanisms are used.

Workloads (such as SQL Server) also support Transparent Data Encryption (TDE).

Technologies that assist with this are:

- Key Vault
- SQL Server Transparent Data Encryption
- Azure Disk Encryption

Third-party virtual machine volume encryption

Data in Transit: Performed by the customer by using transport encryption of traffic traversing exposed virtual machine network endpoints. Microsoft and third-party mechanisms are used.

Actions performed by Microsoft include disk encryption using BitLocker Drive Encryption for bulk import/export operations and encrypting traffic between Azure datacenters.

Technologies that assist with this are:

- HTTPS/REST API
- Azure endpoints
- Azure Import/Export service

Data Access: Performed by the customer by using native protections within the installed operating system to authenticate and authorize access to the virtual hard disk (VHD) data that is exposed through the operating system and published endpoints (for example, operating system file shares).

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

Insight Response: As a Value Added Reseller, this requirement is not applicable to Insight. However, we have described how provides notifications of data breaches.

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

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

AWS utilizes a three-phased approach to manage incidents:

- Activation and Notification Phase: Incidents for AWS begin with the detection of an event. This can come from several sources including:
 - a) 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.
 - b) Trouble ticket entered by an AWS employee
 - c) 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.
- 4) 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.
- 5) Reconstitution Phase - Once the relevant fix activities are complete the call leader will declare that the recovery phase is complete. Post mortem and deep root cause analysis of the incident will be assigned to the relevant team. The results of the post mortem will be reviewed by relevant senior management and relevant actions such as design changes etc. will be captured in a Correction of Errors (COE) document and tracked to completion.

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

available and maintained by the customer support team to alert customers to any issues that may be of broad impact.

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

8.7 (E) MIGRATION AND REDEPLOYMENT PLAN

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

Insight Response: As a Value Added Reseller, this requirement is not applicable to Insight. However, we have described how our CSP partners and service partner meets this requirement.

Microsoft: As for return of data, Microsoft's approach is to provide self-service access to its customer's administrators to extract data upon termination. With regard to Office 365 Services, Microsoft Azure Core Services, Microsoft Dynamics CRM Online Services, and Microsoft Intune Online Services (as each is defined in the Microsoft Online Service Terms, or "OST"), Microsoft provides Customer administrators access to their Customer Data in the Online Services at all times during the term the subscription, and for at least 90 days thereafter (but for no more than 180 days). Where the modality of the Online Service is applicable and as described in the applicable service documentation and service descriptions at the time, Customer Data in the Online Services will be downloadable by the State in a common industry or published Microsoft format (e.g. MS Outlook PST files, MS Office document files in the then-current format,, MS SQL Database files, CSV format files) , during the term of each subscription and for a 90-day "limited functionality" period following expiration (as set forth in the Online Services Terms). For some Online Services service components (also variously described as workloads, services, or modules in Microsoft documentation) download is not possible (such as when the module provides for functionality to synchronize from primary copies of Customer Data held and maintained by the customer), or Customer is intended by the component design to prepare and develop or configure their own download modality (such as when Microsoft provides a platform for Customers own applications to be run as a cloud service).

REAN Cloud:

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.

Storage Protection

REAN S-VPC provides distinctive data protection for information stored on elastic block store volumes using encryption with key management system that enables policy based restrictions to determine where and when encrypted data can be accessed. In addition, server validation applies

identity and integrity rules when servers request access to secure storage volumes. Solution ensures that encryption keys are delivered to valid devices without the need to deploy an entire file system and management infrastructure. This solution protects sensitive information from theft, unauthorized exposure, or unapproved geographic migration to other data centers.

AWS: 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. Refer to the AWS SOC 1 audit report (available under AWS NDA) for more information and validation of the control testing related to access permissions and data deletion for AWS S3 Services. Refer to the AWS PCI Compliance Package (available under AWS NDA) for testing performed to confirm data deletion. Both the AWS SOC 1 audit report and the AWS PCI Compliance Package can be requested.

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

Insight Response: This requirement does not apply to Insight. As the reseller we are not involved in the returning of data. We have provided how our CSP partners and service partners do so below.

Microsoft: See Microsoft response above.

REAN Cloud:

Getting Customer Data back

REAN facilitates this transition using CloudEndure. Customer content, control and ownership always remain with the customer. Transition assistance to customer on premises due to termination or other reason is written into standard REAN terms and conditions.

Setup Services Warranty

REAN warrants that, for a period of 30 days from completion of on boarding, it has performed the Setup Services in substantial accordance with the SOW. Customer must notify REAN of any breach of this warranty no later than 30 days after completion of the Setup Services.

Customer's exclusive remedy and REAN's sole obligation under this warranty will be for REAN to re-perform any non-conforming portion of the Setup Services, or if REAN cannot remedy the breach within 30 days, then refund the portion of the fee attributable to such non-conforming portion of the Setup Services. This warranty will not apply to the extent Customer, its contractors or agents have modified any item or to the extent Customer's equipment does not meet the specifications provided by REAN.

8.8 (E) SERVICE OR DATA RECOVERY

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

Insight Response: In the event of a situation such as those listed in the RFP, Insight's service partner would provide the first line of response for AWS solutions. We have provided how REAN Cloud would respond to the situations.

Service Level Agreements

Following are the service level agreements (SLAs) offered as part of REAN's Managed Services Offering.

Services Warranty

REAN warrants to Customer that commercially reasonable efforts will be made to maintain the online availability of the Service for a minimum availability in any given month as provided in the chart below (excluding scheduled outages, force majeure, and outages that result from any Customer technology issues or incorrect application configurations)

Definitions

"Monthly Uptime Percentage" is calculated by subtracting from 100% the percentage of minutes during the month in which the service was unavailable.

Availability Warranty	Credit
Less than 99.9% but equal to or greater than 99.0%	10% of the managed monthly fee (beyond the warranty)
Less than 99.0%	20% of the managed monthly fee (beyond the warranty)

Customer's exclusive remedy and REAN's sole obligation for its failure to meet the warranty defined above will be for REAN to provide a credit for the applicable month as provided in the chart above (if this agreement is not renewed, then a refund), for the month; provided that Customer notifies REAN of such breach within 30 days of the end of that month.

"Incident" Service Levels

Incidents are the result of services failure or interruptions that may be impacting Customer's ability to conduct business.

- Incidents are assigned severity levels (e.g. P1, P2, P3) based on the impact to the business.
- Many incidents are automatically detected via monitoring utilities. Additionally, Customer can open Incidents by calling REAN directly. The guidelines below will be used for setting Incident Severity.

Incident Severities	Initial Response / Case Assignment	Incident Follow Up / Updates	Time To Resolution	Defined
---------------------	------------------------------------	------------------------------	--------------------	---------

P1 Urgent Full site Outage:	CALL (571)- 252-9696 15 Minute	60 Minutes	4 Hours	A major system or component is down; direct or imminent business impact; client cannot perform business critical functions.
P2 High Partial site - Outage:	CALL (571)- 252-9696 15 Minute	60 Minutes	8 Hours	A system or component is down; client may be experiencing degradation of service, or loss of resilience.
P3 Low Non- Business Impacting	4 hours	Upon Completion	Upon Completion	A system or component is experiencing minor issues but is not causing degradation of service

Standard "Request" Service Levels

A Request may be submitted via the ticketing tool for changes or additions to the infrastructure that are not associated with resolving a Break/fix issue. Examples of Requests include: adding users, patching software and requests for information. Requesting this constitutes approval for REAN to conduct the work. Requests are assigned severity levels (e.g. P1, P2, P3) based on the urgency of the need to support the business.

Request Priorities	Initial Response / Case Assignment	Request Follow Up / Updates	Time To Fulfillment	Defined
P1 (Emergency) Service Request	CALL for Immediate Response (571) 252- 9696	60 Minutes	60 Minutes	Emergency change to avoid or cure potential business impact Service Requests that are included as Emergency include: Emergency access revocation Certain firewall changes designated by Customer as

				Emergency based on the impact and urgency to the Customer Business Certain other Service Request designated by Customer as Emergency based on the impact and urgency to the Customer Business.
P2 Urgent Business Impacts	2 Hours	8 Hours	24 Hours	Non-Standard service request that the customer requires in order to complete day-to-day business activity Service Requests that are Urgent include: Non-emergency access revocation, Certain firewall changes designated by Customer as Urgent based on the impact and urgency to the Customer Business, and Certain other Service Request designated by Customer as Urgent based on the impact and urgency to the Customer Business.
P3 Low Non-Business Impacts	6 Hours	Upon Completion	Upon Completion	Minor service request with no urgency

Setup Services Warranty

REAN warrants that, for a period of 30 days from completion of on boarding, it has performed the Setup Services in substantial accordance with the SOW. Customer must notify REAN of any breach of this warranty no later than 30 days after completion of the Setup Services. Customer's exclusive remedy and REAN's sole obligation under this warranty will be for REAN to re-perform

any non-conforming portion of the Setup Services, or if REAN cannot remedy the breach within 30 days, then refund the portion of the fee attributable to such non-conforming portion of the Setup Services. This warranty will not apply to the extent Customer, its contractors or agents have modified any item or to the extent Customer's equipment does not meet the specifications provided by REAN.

Help Desk

All our clients have direct access to support for technical issues both minor and major and the ability to review existing tickets and request support based on the Service Level Agreement (SLA) supported from our 24/7 global helpdesk. REAN can be reached through an online ticketing system.

a. *Extended downtime.*

REAN Cloud: REAN will comply with 2 week advanced notice for scheduled down times per the requirements. They expect near zero downtime in the event of an outage at the primary facility.

b. *Suffers an unrecoverable loss of data.*

REAN Cloud: Amazon S3 provides a highly durable storage infrastructure designed for mission-critical and primary data storage. Amazon S3 redundantly stores data in multiple facilities and on multiple devices within each facility. To increase durability, Amazon S3 synchronously stores the Participating State or Entity's data across multiple facilities before confirming that the data has been successfully stored. In addition, Amazon S3 calculates checksums on all network traffic to detect corruption of data packets when storing or retrieving data. Unlike traditional systems, which can require laborious data verification and manual repair, Amazon S3 performs regular, systematic data integrity checks and is built to be automatically self-healing. Amazon S3's standard storage is:

- Backed with the Amazon S3 Service Level Agreement for availability
- Designed for 99.999999999% durability and 99.99% availability of objects over a year
- Designed to sustain the concurrent loss of data in two facilities

c. *Offeror experiences a system failure.*

REAN Cloud: Refer the Service Level Agreements in section 8.8.1.

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

REAN Cloud: Refer the Service Level Agreements in section 8.8.1.

e. *Describe your Recovery Point Objective (RPO) and Recovery Time Objective (RTO).*

REAN Standard RTO/ RPO

REAN uses a product called CloudEndure to test and enable disaster recovery, and to seamlessly migrate on-premises applications to the Cloud in the first place. They enable cloud workload mobility using continuous replication of the entire cloud application stack. A single click creates an exact replica of the entire workload, including its up-to-the-second consistent state at a target cloud location within minutes, complete with instances, attached volumes containing all the data, network topology, firewalls, and more.

While snapshot-based and backup solutions result in high RPO and degrade performance of the replicated machines, CloudEndure's real-time, continuous block-level data protection (CDP)

ensures maximum up time and minimal loss of data without consuming additional resources at the source application. REAN then creates a fully functioning, up-to-date copy of the application within minutes. The result is 1-click, fail-safe replication of the entire application to, across, and between multiple cloud locations. CloudEndure automatically discovers the network topology of the workload (IP addresses, subnets, load balancers, firewalls) and transforms it to the compatible format of the target cloud. This ensures that the functionality of the replica workload is identical to the source.

Microsoft: Azure Backup can now back up customer's on-premises application workloads, including Microsoft SQL Server, Hyper-V virtual machines, Microsoft SharePoint, and Microsoft Exchange. They can back up their applications to a local disk or to Azure, allowing them to eliminate local tape libraries and leverage the unlimited storage capability of Azure.

Participating State or Entities can also manage all their on-premises backups from a single user interface. Backup continues to support backups of their production IaaS virtual machines in Azure and to help protect their Windows client data and their shared files and folders.

System Center Data Protection Manager is an option for on-premises, Azure, or Cloud Only backup and recovery.

Azure Site Recovery: ASR's enhanced VMware to Azure scenario is now Generally Available.

This GA release, among other enhancements, is designed to help customers benefit from the following key functionality:

- Elimination of IaaS-based replication and orchestration components/appliance
- MSI-based unified setup of on-premises components, which significantly reduces the time and complexity to onboard to the scenario
- Non-disruptive disaster recovery testing with Test Failover
- ASR-integrated failback experience without vContinuum, with support for alternate location recovery, and original location recovery
- Disk-based replication from source machines, and driver installation without needing a source reboot
- Multi-VM Application and Crash-Consistent Replication for Windows and Linux
- Migration of protected machines from the in-market – Legacy – VMware to Azure scenario to the Enhanced VMware to Azure scenario
- Enterprise-grade enhancements such as support for FQDNs, custom ports, and installation paths
- Support for CentOS & RHEL 6.7, vCenter Server 6.0

8.8.2 Describe your methodologies for the following backup and restore services:

- a. Method of data backups**
- b. Method of server image backups**
- c. Digital location of backup storage (secondary storage, tape, etc.)**
- d. Alternate data center strategies for primary data centers within the continental United States.**

Insight Response: Insight has described how our CSP partners provide backup and restoration services.

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

- Computers are now virtual abstract resources instantiated via code rather than being hardware based.
- Capacity is available at incremental cost rather than up-front cost.
- Resource provisioning takes place in minutes, lending itself to real-time configuration.
- Server images are available on demand, can be maintained by an organization, and can be activated immediately.

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

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

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

A sample workflow is as follows:

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

The first figure depicts a traditional backup approach and the second figure depicts an Amazon EC2 backup approach.




Figure 21: Traditional AWS Backup Approach




Figure 22: AWS EC2 Backup Approach

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

Amazon Machine Image (AMI)

AMIs that customers register are automatically stored in their account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable. This means that the underlying storage mechanism for the AMIs is protected from multiple failure scenarios.

It is also possible to share AMIs between separate AWS accounts. Consequently, customers can create totally independent copies of the AMI by:

- Sharing the original AMI to another specified AWS account controlled by the customer.
- Starting a new instance based upon the shared AMI.
- Creating a new AMI from that running instance.

The new AMI is then stored in the second account and is an independent copy of the original AMI. Of course, customers can also create multiple copies of the AMI within the same account.

Configuration Files

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

Database and File Servers

Backing up data for database and file servers differs from the web and application layers. In general, database and file servers contain larger amounts of business data (tens of GB to multiple TB) that must be retained and protected at all times. In these cases, customers can leverage efficient data movement techniques such as snapshots to create backups that are fast, reliable, and space efficient.

For databases that are built upon RAID sets of Amazon EBS volumes (and have total storage less than 1 TB), an alternative backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. While the destination Amazon EBS volume will have slower performance, it is not being used for data access and can be easily snapshot to Amazon S3 using the Amazon EBS snapshot capability.

Disaster Recovery

The AWS cloud supports many popular DR architectures from “pilot light” environments that are ready to scale up at a moment’s notice to “hot standby” environments that enable rapid failover. With data centers in 12 regions around the world (4 in the United States), AWS provides a set of cloud-based DR services that enable rapid recovery of IT infrastructure and data.

AWS Capabilities for DR/ COOP/ Backup Solutions

With AWS, customers can eliminate the need for additional physical infrastructure, off-site data replication, and upkeep of spare capacity. AWS uses distinct and geographically diverse Availability Zones (AZs) that are engineered to be isolated from failures in other AZs. This

innovative and unique AWS feature enables customers to protect applications from the failure of a single location, resulting in significant cost savings and increased agility to change and optimize resources during a DR scenario.

AWS offers the following high-level DR capabilities:

- **Fast Performance:** Fast, disk-based storage and retrieval of files.
- **No Tape:** Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software.
- **Compliance:** Minimize downtime to avoid breaching Service Level Agreements (SLAs).
- **Elasticity:** Add any amount of data, quickly. Easily expire and delete without handling media.
- **Security:** Secure and durable cloud DR platform with industry-recognized certifications and audits.
- **Partners:** AWS solution providers and system integration partners to help with deployments.

Solution Use Cases

AWS can enable customers to cost-effectively operate multiple DR scenarios to include "backup & restore," "pilot light," "warm standby," and "multi-site". The classifications are arranged by how quickly a system can be available to users after a DR event.

Each DR option is discussed in more detail below:

- **Backup and Restore:** In most traditional environments, data is backed up to tape and sent off-site regularly. Recovery time will be the longest using this method, and lack of automation leads to increased costs. Using Amazon Simple Storage Service (Amazon S3) is ideal for backup data, as it is designed to provide 99.999999999% durability of objects over a given year. Transferring data to and from Amazon S3 is typically done via the network, and it is therefore accessible from any location. Also, with AWS Storage Gateway, customers can automatically back up on-premises data to Amazon S3.
- **Pilot Light for Simple Recovery into AWS Warm Standby Solution:** The idea of the pilot light is an analogy that comes from the gas heater. In a gas heater, a small idle flame that's always on can quickly ignite the entire furnace to heat up a house as needed. This scenario is analogous to a backup and restore scenario; however, customers must ensure that they have the most critical core elements of their system already configured and running in AWS (the pilot light). When the time comes for recovery, customers would rapidly provision a full-scale production environment around the critical core.
- **Warm Standby Solution in AWS:** The term "warm standby" is used to describe a DR scenario in which a scaled-down version of a fully functional environment is always running in the cloud. It further decreases recovery time because, in this case, some services are always running. By identifying business-critical systems, customers could fully duplicate these systems on AWS and have them always on.
- **Multi-Site Solution Deployed on AWS and On-Site:** A multi-site solution runs in AWS as well as on a customer's existing on-premise infrastructure in an active-active configuration. During a disaster situation, an organization can simply send all traffic to AWS servers, which can scale to handle their full production load.

Microsoft: For its Government Community Cloud Services (as defined in Microsoft's service terms and conditions), or "GCC," Customer Content is stored at rest in the United States. In the cases of the GCC versions of Exchange Online, SharePoint Online, Skype for Business, and Dynamics CRM Online, the Customer Content is stored in encrypted format, whereas in the GCC version of Azure Core Services, customers are given the option to encrypt non-public Customer Content.

For the non-GCC (public) versions of the equivalent services, as well as for Microsoft Intune Online Services, certain types of Customer Content are stored at rest in the United States, if set up by the users in the United States. The terms and conditions governing where Customer Data will be stored may be found in the Microsoft Online Services Terms. Finally, for the non-GCC version of Azure Core Services, customers are given the choice of which of Microsoft's worldwide data centers to store and/or process data in.

For purposes of the above, "Customer Content" means the subset of Customer Data created by users. For Office 365 Services, Customer Content shall at least include Exchange Online mailbox content (e-mail body, calendar entries, and the content of e-mail attachments), SharePoint Online site content and the files stored within that site, and Skype for Business Online archived conversations. For Microsoft Dynamics CRM Online Services, Customer Content shall be the entities of Customer Data managed by the Microsoft Dynamics CRM Online Services.

8.9 (E) DATA PROTECTION

8.9.1 Specify standard encryption technologies and options to protect sensitive data, depending on the particular service model that you intend to provide under this Master Agreement, while in transit or at rest.

Insight Response: As a Value-Added Reseller, Insight is not responsible for encrypting or protecting data. However, we have provided detail on how our CSP partners do protect data, while in transit or at rest.

REAN:

Encryption in Transit

Users will access cloud solutions using their web browsers, mobile devices, or desktop software. In all the cases, the data can be transferred on HTTPS connection to ensure encryption in transit.

In addition, users on the go requiring easy but secure remote access to their customer network can utilize VPN capability of S-VPC. S-VPC uses Sophos to provide a broad set of industry-standard VPN technologies including IPsec, SSL, Cisco VPN, iOS and native Windows VPN clients. The customer can set up the content to be accessible only when the mobile devices are on the VPN using the X.509 certificates deployed to their mobile devices using a Mobile Device Management (MDM) solution.

By enabling easy and convenient VPN only access to users, the customer will be able to ensure that only legitimate users with proper credentials are able to access the server. Then they will be required to produce a second factor authentication to access the content securely on HTTPS. This combination provides additional security of content and provides the network level access logs for compliance.

This also ensures that the servers are not exposed to the Internet, which greatly reduces the risk of loss or inadvertent use of data.

Encryption at Rest

The files in the cloud, Meta data in the database, and any other data can be stored encrypted at all times. The encryption keys used to encrypt data are stored in tamper resistant hardware security modules outside the AWS data centers in a Safenet DataSecure Appliance. Safenet DataSecure appliance key management and policy management is provided to ensure compliance, and maximize security.

ProtectV: With SafeNet ProtectV, the customer can encrypt and secure entire virtual machines, protecting these assets from theft or exposure. Further, ProtectV helps encrypt virtual storage, ensuring cloud data is isolated and secured— even in shared, multi-tenant cloud environments used for application hosting, data storage, or disaster recovery.

Key Management: With DataSecure, all cryptographic keys are kept in the centralized, hardened appliance to simplify administration while ensuring tight security for the broadest array of data types. Key versioning streamlines the time-consuming task of key rotation.

Policy Management: Administrators can set authentication and authorization policies that dictate which applications, databases, or file servers can be accessed by particular users in the clear. When combined with strong authentication, this policy-driven security provides a vital layer of protection. DataSecure also offers granular access controls to help customers comply with the separation of duties required in many security mandates.

An administrator can create a policy that prevents certain users from accessing sensitive data without interfering with their day-to-day system administration duties.

Logging, Auditing, and Reporting: When encrypting data within an enterprise, data, keys, and logs are often accessed, encrypted, managed, and generated on multiple devices, in multiple locations. To reduce the cost and complexity of security management, DataSecure provides a single, centralized interface for logging, auditing, and reporting access to data and keys. A centralized mechanism increases security and helps customers ensure compliance with industry mandates and government regulations.

Sophos: Sophos is offered as an annual license based on the number of users that the device sees. A user in the sense of Sophos UTM software licensing, are workstations, clients, servers, and other devices that have an IP address and are protected by or receive service from the UTM appliance. As soon as a user communicates with or through the UTM appliance, their IP address is added to the list of licensed devices in the appliance's local database. No distinction is made if the user communicates with the Internet or with a device in another LAN segment. DNS or DHCP queries to the UTM appliance are also counted. If several users communicate through a single device with only one IP address (e.g., mail server or web proxy), every user is counted as a separate user. The license mechanism only uses data from the last seven days. If an IP address has not been used in the last seven days, it is removed from the database. REAN Cloud solutions can utilize an unlimited user license, or other appropriate license, from Sophos.

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

Securing Data at Rest

There are several options for encrypting data at rest, ranging from completely automated AWS encryption solutions to manual, client-side options. Choosing the right 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.

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.

Microsoft: Outlined below is how data in transit and at rest is protected in the Microsoft environment.

Data at Rest: Performed by the customer by encrypting the virtual hard disk (VHD) files. Microsoft and third-party mechanisms are used.

Workloads (such as SQL Server) also support Transparent Data Encryption (TDE).

Technologies that assist with this are:

- Key Vault
- SQL Server Transparent Data Encryption
- Azure Disk Encryption

Third-party virtual machine volume encryption

Data in Transit: Performed by the customer by using transport encryption of traffic traversing exposed virtual machine network endpoints. Microsoft and third-party mechanisms are used.

Actions performed by Microsoft include disk encryption using BitLocker Drive Encryption for bulk import/export operations and encrypting traffic between Azure datacenters.

Technologies that assist with this are:

- HTTPS/REST API
- Azure endpoints
- Azure Import/Export service

Data Access: Performed by the customer by using native protections within the installed operating system to authenticate and authorize access to the virtual hard disk (VHD) data that is exposed through the operating system and published endpoints (for example, operating system file shares).

8.9.2 Describe whether or not it is willing to sign relevant and applicable Business Associate Agreement or any other agreement that may be necessary to protect data with a Purchasing Entity.

Microsoft: Microsoft's BAA is not negotiable, as they administer their HIPAA-based controls in a uniform manner. However, Microsoft believes that their BAA meets this requirement, such that no exception is required for this section 8.9.2.

AWS: Yes, Insight is willing to sign a BAA and Insight will work with the Participating State or Entity as needed.

8.9.3 Offeror must describe how it will only use data for purposes defined in the Master Agreement, participating addendum, or related service level agreement. Offeror shall not use the government data or government related data for any other purpose including but not limited to data mining. Offeror or its subcontractors shall not resell nor otherwise redistribute information gained from its access to the data received as a result of this RFP.

Insight Response: This requirement is not applicable to insight as we will not access the customer's data. Provided below is Microsoft's explanation of use of customer data.

Microsoft: Customer Data will be used only to provide a Purchasing Entity the Online Services including purposes compatible with providing those services. Offeror and Microsoft will not use Customer Data or derive information from it for any advertising or similar commercial purposes. As between the parties, the Purchasing Entity retains all right, title and interest in and to Customer Data. Neither Offeror nor Microsoft acquires any rights in Customer Data, other than the rights Customer grants to Offeror and its subcontractor, Microsoft, to provide the Online Services to Customer. This paragraph does not affect Microsoft's rights in software or Online Services Microsoft licenses to Purchasing Entity.

For DPT Services, Offeror and Microsoft use data mining solely for the purposes of providing those cloud services, subject to the above-mentioned restrictions. Microsoft will not use data mining in the DPT Services for unrelated commercial purposes, advertising or advertising-related purposes, or for any other purpose other than security or service delivery analysis that is not explicitly authorized.

8.10 (E) SERVICE LEVEL AGREEMENTS

8.10.1 Offeror must describe whether your sample Service Level Agreement is negotiable. If not describe how it benefits purchasing entity's not to negotiate your Service Level Agreement.

AWS: AWS' Service level Agreement for each of its product offerings is non-negotiable because of the rapidly evolving nature of AWS's product offerings.

AWS innovates extremely quickly, and released over 700 new features or Services in 2015. AWS has over a million 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. AWS needs the right to make changes across its customer base, and is not able to offer a Participating State or Entity a custom notice period.

Microsoft: The Service Level Agreement for Microsoft Online Services is not negotiable, as it pertains to standardized multitenant cloud services, uniformly delivered to many thousands of customers and millions of users, and relies upon automated processes and standard operating procedures. Purchasing Entities benefit to the extent that Microsoft's SLA is competitive and of industry standard quality, and also benefit from the cost savings which standardized multitenant cloud services provide over the types of customized outsourcing services that would allow for such a negotiation.

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

Microsoft: Enclosed within Insight's response is Microsoft's Service Level Agreement for Microsoft's Online Services. Insight is unable to negotiate Microsoft's SLA or submit to a Purchasing Entity's SLA requirements due to corporate policies surrounding the operational and security controls of its cloud service. For clarity, Microsoft's SLAs are administered in a consistent and in some cases automated way for all its customers, and may therefore not be customized. For any given cloud service, our SLA in effect as of the time a subscription order is first placed is locked and will not change during the term of a subscription order. Upon renewal of a Purchasing Entity's subscription order, Microsoft's then-current SLA will supersede the previous SLA. The Purchasing Entity's renewal of its subscription will constitute its written approval of the then-current (new) SLA. Microsoft's historical practice has been to improve its SLAs over time, and they have never before adversely changed any SLA terms.

AWS: AWS currently provides Service Level Agreements (SLAs) for several products. Due to the rapidly evolving nature of AWS's product offerings, SLAs are best reviewed directly on our website via the links below:

- Amazon EC2 SLA: [http:// aws.amazon.com/ ec2-sla/](http://aws.amazon.com/ec2-sla/)
- Amazon S3 SLA: [http:// aws.amazon.com/ s3-sla](http://aws.amazon.com/s3-sla)
- Amazon CloudFront SLA: [http:// aws.amazon.com/ cloudfront/ sla/](http://aws.amazon.com/cloudfront/sla/)
- Amazon Route 53 SLA: [http:// aws.amazon.com/ route53/ sla/](http://aws.amazon.com/route53/sla/)
- Amazon RDS SLA: [http:// aws.amazon.com/ rds-sla/](http://aws.amazon.com/rds-sla/)

SLAs must remain fluid for AWS because innovates extremely quickly. In 2015 alone, AWS released over 700 new features or Services.

AWS has over a million 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. AWS needs the right to make changes across its customer base, and is not able to offer a Participating State or Entity a custom notice period. Relevantly: AWS will provide 90 days prior notice before materially reducing benefits under a SLA.

Provided below is a sample of AWS's SLA for Amazon EC2.

Service Commitment

AWS will use commercially reasonable efforts to make Amazon EC2 and Amazon EBS each available with a Monthly Uptime Percentage (defined below) of at least 99.95%, in each case during any monthly billing cycle (the "Service Commitment"). In the event Amazon EC2 or Amazon EBS does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Monthly Uptime Percentage" is calculated by subtracting from 100% the percentage of minutes during the month in which Amazon EC2 or Amazon EBS, as applicable, was in the state of "Region Unavailable." Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon EC2 SLA Exclusion (defined below).
- "Region Unavailable" and "Region Unavailability" mean that more than one Availability Zone in which you are running an instance, within the same Region, is "Unavailable" to you.
- "Unavailable" and "Unavailability" mean:
 - For Amazon EC2, when all of your running instances have no external connectivity.
 - For Amazon EBS, when all of your attached volumes perform zero read write IO, with pending IO in the queue.
- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Commitments and Service Credits

Service Credits are calculated as a percentage of the total charges paid by you (excluding one-time payments such as upfront payments made for Reserved Instances) for either Amazon EC2 or Amazon EBS (whichever was Unavailable, or both if both were Unavailable) in the Region affected for the monthly billing cycle in which the Region Unavailability occurred in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0%	30%

We will apply any Service Credits only against future Amazon EC2 or Amazon EBS payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the Unavailability occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar (\$1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability, non-performance, or other failure by us to provide Amazon EC2 or Amazon EBS is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you must submit a claim by [opening a case in the AWS Support Center](#). To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

1. the words "SLA Credit Request" in the subject line;
2. the dates and times of each Unavailability incident that you are claiming;
3. the affected EC2 instance IDs or the affected EBS volume IDs; and
4. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the Monthly Uptime Percentage of such request is confirmed by us and is less than the Service Commitment, then we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon EC2 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon EC2 or Amazon EBS, or any other Amazon EC2 or Amazon EBS performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon EC2 or Amazon EBS; (iii) that result from any actions or inactions of you or any third party, including failure to acknowledge a recovery volume; (iv) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); (v) that result from failures of individual instances or volumes not attributable to Region Unavailability; (vi) that result from any maintenance as provided for pursuant to the AWS Agreement; or (vii) arising from our suspension and termination of your right to use Amazon EC2 or Amazon EBS in accordance with the AWS Agreement (collectively, the "Amazon EC2 SLA Exclusions"). If availability is impacted by factors other than those used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

Figure 23: Sample SLA for AWS EC2

8.11 (E) DATA DISPOSAL

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

Insight Response: Insight does not have access to data that would require us to have data disposal policies and procedures. However, we have described how our CSP partners address this requirement.

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

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

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

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.

REAN: In alignment with ISO 27001 standards, when a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M ("National Industrial Security Program Operating Manual ") or NIST 800-88 ("Guidelines for Media Sanitization") to destroy data as part of the decommissioning process.

If a hardware device is unable to be decommissioned using these procedures, the device will be degaussed or physically destroyed in accordance with industry-standard practices.

Microsoft: Microsoft Azure supports best practice procedures and a data removal solution which is NIST 800-88 compliant. Disk drives that can't be cleaned a destruction process that destroys it (i.e. shredding) and renders the recovery of information impossible (e.g., disintegrate, shred, pulverize, or incinerate). The appropriate means of disposal is determined by the asset type. Records of the types of media and other tracking information pertaining to the destruction are recorded. All Microsoft Azure services utilize approved media storage and disposal management services. Paper documents are destroyed by approved means at the pre-determined end-of-life cycle.

Microsoft Azure supports NIST 800-88 Guidelines on Media Sanitization, which address the major concern of ensuring that data is not released unexpectedly. Microsoft Azure guidelines encompass both physical and digital sanitization.

8.12 (E) PERFORMANCE MEASURES AND REPORTING

Insight Response: Insight has provided responses for question 8.12 as they pertain to our CSP partner solution.

8.12.1 Describe your ability to guarantee reliability and uptime greater than 99.5% . Additional points will be awarded for 99.9% or greater availability.

Amazon: Amazon S3's standard storage is based by Amazon S3's SLA for availability and is designed for 99.999999999% durability and 99.99% availability of objects over a year.

Microsoft: Microsoft Azure virtual machine (VM) can be made highly available by creating an Azure Availability Set. All data is replicated three times within each datacenter. Customers have the option to enable Geo-redundant storage for additional failover capabilities. Geo-redundant storage insures data is replicated three times in a local datacenter and another three copies in a datacenter a several hundred miles away.

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

Microsoft: Microsoft has commitments for uptime and connectivity for the individual Azure Services. These SLAs are outlined in depth in the Microsoft Online Services SLA document that are provided with the response.

AWS: AWS has commitments for uptime and connectivity for the individual services. These SLAs are outlined in depth by following the links provided below.

- Amazon EC2 SLA: <http://aws.amazon.com/ec2-sla/>
- Amazon S3 SLA: <http://aws.amazon.com/s3-sla>
- Amazon CloudFront SLA: <http://aws.amazon.com/cloudfront/sla/>
- Amazon Route 53 SLA: <http://aws.amazon.com/route53/sla/>
- Amazon RDS SLA: <http://aws.amazon.com/rds-sla/>

8.12.3 Specify and provide the process to be used for the participating entity to call/contact you for support, who will be providing the support, and describe the basis of availability.

Microsoft: Support is provided 24x7x365 with options for telephone and email. Microsoft will be providing technical support through Microsoft Premier Services.

AWS: AWS Support is a one-on-one, fast-response support channel that is staffed 24x7x365 with experienced and technical support engineers to help customers of all sizes and technical abilities successfully utilize the products and features provided by AWS.

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

Microsoft: Microsoft provides financially backed SLA's described in response 8.10. If Service Levels are not met, the Purchasing Entity will be entitled to a Service Credit. Service credit amounts vary by the Azure Service and are outlined in the Online Services SLA submitted with this response. Microsoft requires that customers submit an SLA breach claim to customer support by the end of the calendar month after the event has happened.

AWS: AWS does not offer incident response time SLAs at this point in time.

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

Microsoft: Microsoft provides financially backed SLA's described in response 8.10. Microsoft requires that customers submit an SLA breach claim to customer support by the end of the calendar month after the event has happened.

AWS: 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; Amazon recommends 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.

Routine, emergency, and configuration changes to existing AWS infrastructure are authorized, logged, tested, approved, and documented in accordance with industry norms for similar systems.

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

Microsoft: Azure sometimes restarts customer's VM as part of regular, planned maintenance updates in the Azure datacenters. Unplanned maintenance events can occur when Azure detects a serious hardware problem that affects their VM. For unplanned events, Azure automatically migrates the VM to a healthy host and restarts the VM.

Single VM, not part of an availability set, Azure notifies the subscription's Service Administrator by email at least one week before planned maintenance because the VMs could be restarted during the update. Applications running on the VMs could experience downtime.

Use Azure PowerShell to view the reboot logs when the reboot occurred due to planned maintenance. For details, see Viewing VM Reboot Logs.

AWS: The Shared Responsibility nature of the AWS solution dictates that the customer owns their architecture design for fault tolerance when using AWS. AWS Shared Responsibility provides 5 SLA's for disaster or problems with the Infrastructure.

Businesses are using the AWS cloud to enable faster disaster recovery of their critical IT systems without incurring the infrastructure expense of a second physical site. The AWS cloud supports many popular disaster recovery (DR) architectures from "pilot light" environments that are ready to scale up at a moment's notice to "hot standby" environments that enable rapid failover. With data centers in 12 regions around the world, AWS provides a set of cloud-based disaster recovery services that enable rapid recovery of Participating States and Entity's IT infrastructure and data.

Best Practices noted below:

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

AWS makes available to customers multiple resources to help organizations start using AWS for a DR/COOP and backup solution, including AWS produced whitepapers, industry reports such as Forrester, sample DR architecture drawings, and informational pages on the web.

8.12.7 Provide a sample of performance reports and specify if they are available over the Web and if they are real-time statistics or batch statistics.

AWS: AWS has a number of performance reports available for each of the components of the AWS product stack. A description and example of the reports for each is provided below.

AWS – EC2 Usage Reports

The usage reports provided by Amazon EC2 enable customers to analyze the usage of their instances in depth. The data in the usage reports is updated multiple times each day. Customers can filter the reports by AWS account, region, Availability Zone, operating system, instance type, purchasing option, tenancy, and tags. The report is available over the web via Amazon AWS Management Console. An example of this report is provided below.




Figure 24: AWS EC2 Usage Report Screenshot

AWS CloudFront Usage Report


The Amazon CloudFront console can display a graphical representation of the client's CloudFront usage that is based on a subset of the usage report data. They can display charts for a specified date range in the last 60 days, with data points every hour or every day. They can usually view

data about requests that CloudFront received as recently as four hours ago, but data can occasionally be delayed by as much as 24 hours. This report is available over the web via the Amazon AWS Management Console. An example of the usage reports is provided below.

CloudFront Usage Reports

Start Date	2016-02-23	Granularity	Daily (any period in previous 60)	Web Distribution	All Web Distributions (excludes)
End Date	2016-03-07	Billing Region	All Regions	Update	Download CSV

Number of Requests ([Millions](#) | [Thousands](#) | [Not Scaled](#)) [Show Details](#)



HTTP Requests: Total: 0 Average: 0 Minimum: 0 Maximum: 0

Data Transferred By Protocol ([Gigabytes](#) | [Megabytes](#) | [Kilobytes](#)) [Show Details](#)




Figure 25: CloudFront Usage Report Examples

AWS Billing Reports – (Available over the web via the Amazon AWS Management Console)

Billing reports provide information about customer's usage of AWS resources and estimated costs for that usage. Customers can have AWS generate billing reports that break down their estimated costs in different ways:

- By the hour, day, or month
- By each account in the Participating States and Entity's organization
- By product or product resource

- By tags that the Participating State or Entity defines itself

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

Microsoft: Azure provides a current dashboard of service health which is updated in 10 minutes intervals. As a subscriber to Azure services, users with access to the Azure administrative portal are provided real time performance statics of all services with the ability to drill down into each component and service. An example of this dashboard is provided below.




Figure 26: Microsoft Online Azure Status Dashboard

All raw data can be downloaded from the Azure portal and imported into a reporting tool, such as Excel. These reports can then be dissected to support any number of historical usage reports required. Additionally Power BI is an option to provide visualized data points and build KPIs across the customer's Azure subscription. This enables clients to build Business Intelligent reports and analytics on a real time or historical basis. Power BI is accessed through a portal and the reports can be downloaded locally.

AWS: Below are additional examples of AWS Billing Reports, EC2 Usage Reports and CloudFront Reports and Analytics. The examples display how the customer can access and print the reports.

Billing Reports:

Bills ?

Date: Download CSV Print

Summary	Amount
AWS Service Charges	\$0.02
Other Details	
Total	\$0.02

EC2 Usage Reports:

Instance Usage Reports

Granularity: Time range:

Filter: Group by:

Unit:

Reserved Instance Utilization Reports

Time range:

Filter:

CloudFront Reports & Analytics: The following metrics can be printed from CloudWatch statistics: Cache Statistics, Monitoring and Alarms, Top Referrers, Usage, Viewers

CloudFront Cache Statistics Reports

Start Date	2016-02-23	Granularity	Daily (any period in previous 60)	Web Distribution	All Web Distributions
End Date	2016-03-07	Viewer Location	All Locations	Update	Download CSV

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

Microsoft: On demand, scheduled automation deployment is offered 24x365

AWS: All customers receive Basic Support that is included with all AWS accounts. All plans, including Basic Support, provide 24x7 access to customer services, AWS documentation, whitepapers, and support forums.

If the Participating Entity should choose, higher level support plans are available – Developer, Business, and Enterprise. The higher the level, the more advanced support the customer will receive.

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

Microsoft: On demand, scheduled automation deployment is offered 24x365.

AWS: The Participating State or Entity will be responsible for architecting the scale up and scale down. However, this process is made easy with Amazon's auto scaling functionality.

Auto Scaling helps Participating States and Entities maintain application availability and allows the user to scale their Amazon EC2 capacity up or down automatically according to conditions they define. They can use Auto Scaling to help ensure that they are running their desired number of Amazon EC2 instances. Auto Scaling can also automatically increase the number of Amazon EC2 instances during demand spikes to maintain performance and decrease capacity during lulls to reduce costs. Auto Scaling is well suited both to applications that have stable demand patterns or that experience hourly, daily, or weekly variability in usage.

8.13 (E) CLOUD SECURITY ALLIANCE

Describe your level disclosure of compliance with CSA Star Registry for each Cloud solutions offered.

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

Insight Response: This requirement is not applicable to Insight because we do not manage the environment; however, we have provided disclosure statements for each of our CSP partners.

AWS:

a. AWS is compliant with Level 1 CSA STAR Registry Self-Assessment. Insight has enclosed AWS' self-assessment found within AWS' Risk and Compliance Whitepaper. The information requested can be located on pages 25-61. Please refer to AWS' self-assessment found within their Risk and Compliance Whitepaper, page 25-61. This is the latest CAIQ (v3) released by the CSA.

b. There is no response required for Exhibit B. Exhibit A questions refer to the Exhibit B for mapping references to common standards. Please refer to the completed AWS' self-assessment found within AWS' Risk and Compliance Whitepaper, page 25-61. This is the latest CAIQ (v3) released by the CSA.

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

- Level 2 Attestation is based on SOC2, which can be requested under NDA. The SOC 2 report audit attests that AWS has been validated by a third party auditor to confirm that AWS' control objectives are appropriately designed and operating effectively.
- Level 2 Certification is based on ISO 27001:2005 – the AWS ISO 27001:2005 certification has been submitted with the proposal response.

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

d. Per the CSA website, CSA Level 3 Continuous Monitoring is still under development. AWS has implemented and documented a Continuous Monitoring Plan which defines AWS' approach to conducting continuous monitoring with its authorizing officials within the FedRAMP Security Assessment Framework. It is based on the continuous monitoring process described in NIST SP 800-137, Information Security Continuous Monitoring for Federal Information Systems and Organization, and has been reviewed and validated by a third-party assessor as part of our annual FedRAMP Assessment. It is made available to customers within the AWS FedRAMP Package which can be obtained under NDA.

Microsoft:

Insight's subcontractor, Microsoft, is on the Board of Directors of the Cloud Security Alliance (CSA).

The responses below are intended to provide information on how Microsoft operates Azure services; customers have accountability to control and maintain their cloud environment once the service has been provisioned (for example, user access management with appropriate policies and procedures in accordance with regulatory requirements).

Adrienne Hall

General Manager for Issues & Crisis Management at Microsoft

Adrienne Hall is the General Manager for Issues & Crisis Management at Microsoft, overseeing communication regarding a wide range of topics. Hall works closely with colleagues to ensure accurate and timely information is delivered, providing the details for customers and partners to take action, if or when, needed. She represents Microsoft as a board member of the Cloud Security Alliance, is a recognized speaker on cloud computing, cybersecurity, and is a recurring author on the Microsoft Cyber Trust Blog. Hall also regularly meets with customers, law enforcement, and government agencies to share best practices for the protection of critical infrastructure and information technology systems. Hall is a recipient of awards and recognition from global organizations, including several from international law enforcement agencies for her collaboration on disrupting cybercrime. In 2013, Hall received the Women in Security Award from Professional Security Magazine UK for her contributions to the field of cybersecurity and in 2014, received a ten-year service award from the Executive Women's Forum for Security, Privacy, and Risk Management Professionals. Hall is a graduate of the University of British Columbia and the Michael G. Foster School of Business at the University of Washington.

Submitted with Insight's proposal is Microsoft's Cloud Controls Matrix (CCM) and Consensus Assessments Initiative Questionnaire (CAIQ) for Microsoft Azure.

Azure's CCM responses are scoped to Azure services in alignment with their ISO 27001 and PCI DSS attestations, including Microsoft's physical datacenters:

- Compute (Virtual Machines, Cloud Services, RemoteApp)
- Web and Mobile (App Service, Mobile Apps, API Management)
- Data and Storage (SQL Database, Storage, StorSimple)
- Analytics (HDInsight, Data Factory)
- Networking (Virtual Networks)
- Hybrid Integration (BizTalk Services, Service Bus, Backup, Site Recovery)
- Identity and Access Management (Azure Active Directory, Multi-Factor Authentication)
- Developer Services (Visual Studio Online)
- Management (Preview Portal, Scheduler, Key Vault)

Azure validates services using third party penetration testing based upon the OWASP (Open Web Application Security Project) top ten and CREST-certified testers. The outputs of testing are tracked through the risk register, which is audited and reviewed on a regular basis to ensure compliance to Microsoft security practices.

8.14 (E) SERVICE PROVISIONING

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

Insight Response: In most cases, Insight is not involved in the implementation and ongoing management of the Cloud Solutions. Those are self-managed by the Purchasing Entity, with support from Microsoft. Microsoft does provide technical support through the cloud portals and administrative site, and generally responds to incidents in just a few hours. But unless Insight is engaged to provide technical or consulting services around the implementation of these Cloud Solutions, Insight is not involved. (If Insight is engaged to provide such services, emergency procedures and service level agreements will be included in Statements of Work that define the service provided.)

REAN:

Standard “Request” Service Levels

A Request may be submitted via the ticketing tool for changes or additions to the infrastructure that are not associated with resolving a Break/fix issue. Examples of Requests include: adding users, patching software and requests for information. Requesting this constitutes approval for REAN to conduct the work. Requests are assigned severity levels (e.g. P1, P2, P3) based on the urgency of the need to support the business.

Request Priorities	Initial Response /Case Assignment	Request Follow Up / Updates	Time to Fulfillment	Defined
P1 (Emergency) Service Request	CALL for Immediate Response (571) 252-9696	60 Minutes	60 Minutes	Emergency change to avoid or cure potential business impact Service Requests that are included as Emergency include: Emergency access revocation Certain firewall changes designated by Customer as Emergency based on the impact and urgency to the Customer Business Certain other Service Request designated by Customer as Emergency based on the impact and urgency to the Customer Business.
P2 Urgent Business Impacts	2 Hours	8 Hours	24 Hours	Non-Standard service request that the customer requires in order to complete day-to-day business activity Service Requests that are Urgent include: Non-emergency access revocation, Certain firewall changes designated by Customer as Urgent based on the impact and urgency to the Customer Business, and Certain other Service Request designated by Customer as Urgent based on the impact and urgency to the Customer Business.

P3 Low Non-Business Impacts	6 Hours	Upon Completion	Upon Completion	Minor service request with no urgency
---------------------------------------	---------	--------------------	--------------------	---------------------------------------

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

Insight Response: There is no standard lead-time for provisioning Cloud Solutions. The time needed is based on a number of factors, including the scope and complexity of such provisioning. For example, the time needed to stand up a small number of Exchange Online mailboxes is much less than the time needed to migrate a larger number of existing mailboxes to Exchange online, especially when the provisioning includes full redundancy, failover and advanced Exchange features.

However, Insight has a great deal of experience in provisioning Microsoft Cloud Solutions to a wide variety of organizations, and once the scope and nature of the work is assessed, we can provide a detailed and accurate forecast of the lead-time required for the provisioning. We can also make recommendations of options that will reduce the lead-time needed.

8.15 (E) BACK UP AND DISASTER PLAN

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

Insight Response: Retention policies for mailboxes are applied at the mailbox level, and can be done individually or in bulk. A default retention policy exists, which applies to every mailbox as it is created, and can be edited to suit agency needs.

Retention policies for SharePoint can be applied at the site level, and can be made to suit agency needs

A particular administrative role will be required when configuring or reporting on policy. Office 365 services are delineated by tenant; and a defined administrative role will span the tenant.

If a security functional boundary needs to be imposed restricting the scope of an administrator to a DNS domain, it must be done by separating the domain into its own tenant.

For Office 365, there is currently no boundary that can be set around a purchasing entity if more than one entity exists for a single DNS domain. However, roles can be separated by scope of function, even if that scope can't be focused on a single domain.

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

Insight Response: Disasters are mitigated to the extent of the customer's risk appetite, meaning Azure's native Disaster Recovery mechanisms take into account the ordinary failure rate of commodity hardware by using redundant systems for even the most basic consumer-oriented account. Additional redundancy within the same datacenter (or region) is available at additional cost for any account that provisions those options. Finally, all accounts can elect to configure Azure to be excessively redundant using multiple datacenters, multiple geographic regions and

multiple services (including Availability Sets, Recovery Groups, Site Recovery Vaults, Managed DNS, Traffic Management, and Scaling) to provide a near-zero RTO and RPO for any given service regardless if a failure is systemic or geographic.

8.15.3 Describe the infrastructure that supports multiple data centers within the United States, each of which supports redundancy, failover capability, and the ability to run large scale applications independently in case one data center is lost.

Insight Response: Azure is built in over 100 datacenters in 28 regions worldwide. Each datacenter provides redundant infrastructure and server platforms to customers, with the option to extend redundancy to an additional datacenter for failover, whether declared by Microsoft or the customer. Azure has the ability to scale to over 100,000 CPU cores for any given application, and that scaling is not limited to any single datacenter.

AWS: 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. AWS' availability and fault-tolerant design are outlined below.

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 Participating States and Entities 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.

Participating States and Entities should architect their AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure modes, including natural disasters or system failures.

Fault-Tolerant Design: Amazon's infrastructure has a high level of availability and provides Participating States and Entities 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 Participating States and Entities 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.

Participating States and Entities should architect their AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to remain resilient in the face of most failure scenarios, including natural disasters or system failures. However, they should be aware of location-dependent privacy and compliance requirements, such as the EU Data Privacy Directive. Data is not replicated between regions unless proactively done so by the customer, thus allowing customers with these types of data placement and privacy requirements the ability to establish compliant environments. It should be noted that all communications between regions is across public Internet infrastructure; therefore, appropriate encryption methods should be used to protect sensitive data.

As of this writing, there are twelve regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), AWS GovCloud (US), EU (Ireland), EU (Frankfurt), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), South America (Sao Paulo), and China (Beijing).

AWS GovCloud (US) is an isolated AWS Region designed to allow US government agencies and customers to move workloads into the cloud by helping them meet certain regulatory and compliance requirements. The AWS GovCloud (US) framework allows US government agencies and their contractors to comply with U.S. International Traffic in Arms Regulations (ITAR) regulations as well as the Federal Risk and Authorization Management Program (FedRAMP) requirements. AWS GovCloud (US) has received an Agency Authorization to Operate (ATO) from the US Department of Health and Human Services (HHS) utilizing a FedRAMP accredited Third Party Assessment Organization (3PAO) for several AWS services.

The AWS GovCloud (US) Region provides the same fault-tolerant design as other regions, with two Availability Zones. In addition, the AWS GovCloud (US) region is a mandatory AWS Virtual Private Cloud (VPC) service by default to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private (RFC 1918) addresses.

8.16 (E) SOLUTION ADMINISTRATION

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

Insight Response: For Microsoft solutions, the purchasing entity is in full control of the setup and management of identity and user accounts, through their own Active Directory as well as administrative portals that are included in the Cloud Solutions. Administrators who are authorized by the Purchasing Entity can also use these tools to assign licenses to users, to create and enforce group policies, and to provide or restrict users from accessing company data and other assets.

REAN: REAN S-VPC environment provides various convenient options to the end users to access the environment and initiate their VPN connections. These include:

- ◆ HTML5 based remote access VPN that they can initiate from any HTML5 compatible browser with requiring any plug-in.
- ◆ SSL remote access VPN that provides additional security by a double authentication using X.509 certificates and username/password.
- ◆ IPSec based VPN using native Windows or Mac VPN clients
- ◆ Mobile VPN using native iPhone VPN client to securely connect to VPC

System administrator access control is provided through the integration of GU identity and access management solution. This suite supplements the AWS Management Console by vaulting administrator's credentials, enforcing separation of duties and recording all accesses and actions.

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

Insight Response: This requirement does not apply to Insight as a Value Added Reseller. Provided below is how our CSP partners and service provider addresses this requirement.

REAN: REAN designs, develops and deploys a packaged secure virtual private cloud (S-VPC) framework that facilitates assurance of information protection. The Amazon S-VPC lets customers provision a private, isolated section of the AWS Cloud. We integrated a front-end user layer utilizing a unified threat management suite. The suite provides firewall services, intrusion protection/detection services, secure Virtual Private Network (VPN) connectivity, packet filtering and web application firewall protection not available via AWS standard offerings.

This front-end protects against denial-of-service attacks, worms, and hacker exploits; secures email from spam and viruses; filters web browsing; and provides wireless network protection. An application layer leverages the pay-per-use approach and elastic AWS services infrastructure in the S-VPC to deliver a scalable and highly available solution.

The keys for data encryption can be within the customer corporate data center with oversight provided by Participating States and Entity's system administrators. System administrator access control is provided through the integration of the customer's identity and access management solution. This suite supplements the AWS Management Console by vaulting administrator's credentials, enforcing separation of duties and recording all accesses and actions. Finally, they provide a management layer that provides continuous real-time forensics to monitor for patterns of malicious activity across the S-VPC framework.

This framework can enable the customer to provision almost any application in the application layer while benefitting from the security and scalability of the S-VPC framework. As part of REAN Cloud engineering service offerings they can support the customer in making their applications cloud ready and integrated into this pre-defined, proven framework.

The security of the framework above has been validated by independent third-party auditors and meets the Federal Information Security Management Act (FISMA) moderate security level, the Payment Card Industry (PCI) security standard, the Service Organization Control (SOC) 1 standard, and is fully compliant with the Health Insurance Portability and Accountability Act (HIPAA) standard.

The following figure demonstrates sample application architecture built on this S-VPC framework, the REAN Cloud Secure Mobile Collaboration Solution.




Figure 27: REAN Cloud Example Secure Architecture: Mobile Collaboration Solution

Microsoft: The Microsoft Cloud Solutions include anti-virus and anti-malware protections for data at rest and data in motion between the customer and the Microsoft Cloud, as well as between points within the Microsoft Cloud.

For Office 365 Services, Microsoft Azure Core Services, Microsoft Dynamics CRM Online Services, and Microsoft Intune Online Services (as each is defined in the Microsoft Online Services Terms), Microsoft will implement and maintain all appropriate administrative, physical, technical and procedural safeguards in accordance with the terms and conditions of the Microsoft Online Services Terms, at all times during the term of the Master Agreement, to secure Customer Data from Security Incident, protect Customer Data and the applicable Online Services from hacks, introduction of viruses, disabling devices, malware and other forms of malicious or inadvertent acts that can disrupt a Purchasing Entity's access to its Customer Data.

Additionally, see the section of Microsoft's Service Level Agreement pertaining to anti-virus.

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

Insight Response: With the Microsoft Cloud Solutions, the Purchasing Entity is in full control of their data, and can migrate some or all of it to another Cloud provider, another hosting company, or the Purchasing Entity's own datacenter at any time. Insight Services can be engaged to manage the migration of this data if necessary.

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

Insight Response: The administration of the Microsoft Cloud Solutions is done by the Purchasing Entity themselves, but the solution does allow different administrators to administer different parts of the solution. It is possible and even common for different participating entities

and sub-entities to separately manage for themselves the services included in the Microsoft Cloud solution.

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

Insight Response: Microsoft Cloud Solutions use a given set of administration tools that are used to administer the solutions, but there is some flexibility in the use of these tools. Without knowing the defined administration policies themselves, it is difficult to confirm that Microsoft's tools and the Purchasing Entity's policies will work together seamlessly. But we can confirm that these tools are similar in nature to administration tools used to administer other Microsoft technology, so there is no reason to expect that if such policies are supported by current on premise Microsoft administration tools, the move to a Cloud solution would change that.

8.17 (E) HOSTING AND PROVISIONING

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

Insight Response: Insight has provided responses for both AWS and Microsoft cloud solutions.

AWS: Amazon has many years of experience in designing, constructing, and operating large-scale datacenters. This experience has been applied to the AWS platform and infrastructure. AWS datacenters are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access datacenter floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff.

AWS only provides datacenter access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or AWS. All physical access to datacenters by AWS employees is logged and audited routinely.

Microsoft: Provisioning is done by default using a web-based Management Portal, or Azure Resource Manager, or PowerShell commandlets, or an API using one of the provided SDKs (in JSON, REST, Node.js, PHP, Python, or Java). Additional provisioning stacks include Chef/Puppet (available as a native add-in), and select third-party tools available in the Azure Marketplace.

8.17.2 Provide tool sets at minimum for:

1. Deploying new servers (determining configuration for both stand alone or part of an existing server farm, etc.)

Microsoft: By default this is accomplished using a web-based Management Portal, or Azure Resource Manager, or PowerShell commandlets, or an API using one of the provided SDKs (in JSON, REST, Node.js, PHP, Python, or Java).

2. Creating and storing server images for future multiple deployments

Microsoft: By default this is accomplished using a web-based Management Portal, or Azure Resource Manager, or PowerShell commandlets, or an API using one of the provided SDKs (in JSON, REST, Node.js, PHP, Python, or Java).

3. Securing additional storage space

Microsoft: By default this is accomplished using a web-based Management Portal, or Azure Resource Manager, or PowerShell commandlets, or an API using one of the provided SDKs (in JSON, REST, Node.js, PHP, Python, or Java).

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).

Microsoft: Monitoring provides the granular usage statistics for every service in Azure. Billing and financial consumption is limited to the Account Owner. Resource and Service usage is available either in the Management Portal or via an API. Alerting and Response Management are provided by various services within Azure.

AWS: Provided below are details on AWS' tools as they pertain to the requirements of the RFP.

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

Use Existing Management Tools

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

- AWS Management Portal for vCenter enables customers to manage their AWS resources using VMware vCenter. The portal installs as a vCenter plug-in within the existing vCenter environment. Once installed, it enables customers to migrate VMware VMs to Amazon EC2 and manage AWS resources from within vCenter. The AWS resources that customers create using the portal can be located in their AWS account, even though those resources have been created using vCenter. For experienced VMware administrators, AWS Management Portal for vCenter provides a familiar look and feel that can make it easy to start using AWS. AWS Management Portal for vCenter is available at no additional charge.
- The Amazon EC2 VM Import Connector extends the capabilities of VMware vCenter to provide a familiar graphical user interface customers can use to import their preexisting Virtual Machines (VMs) to Amazon EC2. Using the connector, importing a VM is as simple as selecting a VM from the vSphere infrastructure, and specifying the AWS region, Availability Zone, operating system, instance size, security group, and Amazon Virtual

Private Cloud (Amazon VPC) details (if desired) into which the VM should be imported. Once the VM has been imported, customers can launch it as an instance from the AWS Management Console and immediately take advantage of all the features of Amazon EC2.

- AWS Management Pack for Microsoft System Center enables customers to view and monitor their AWS resources directly in the Operations Manager console. This way, customers can use a single, familiar console to monitor all of their resources, whether they are on-premises or in the AWS cloud. Participating States and Entities get a consolidated view of all AWS resources across regions and Availability Zones. It also has built-in integration with Amazon CloudWatch so that the metrics and alarms defined in Amazon CloudWatch surface as performance counters and alerts in the Operations Manager console.

8.18 (E) TRIAL AND TESTING PERIODS (PRE- AND POST- PURCHASE)

8.18.1 Describe your testing and training periods that your offer for your service offerings.

Insight Response: When delivered by Insight Services, our recommended testing is done in three phases.

Phase 1 Device Under test (DUT) is used to test performance and proficiency

- Virtual lab
- Performance testing

Phase 2 Systems Integration Testing (SIT) this is the high-level testing process in which testers verify that all related systems maintain data integrity and can operate in coordination with other systems in the same environment.

- Interoperability testing
- Coexistence testing
- Functionality testing

Phase 3 User Acceptance testing (UAT) where actual software users test the software to make sure it can handle required tasks in real-world scenarios, according to specifications.

- Pre-production pilot
- End user services

For all of our services, we provide in-depth knowledge transfer and admin training with all appropriate architecture documentation to support the design of the solution.

REAN Cloud: REAN typically tailors its knowledge transfer and delivery to each customer's needs and requirements. A key deliverable, though, in the early stages of an engagement is the communications plan which details roles and responsibilities, accountability, methods and modalities of communication, document management, training media and other key components of knowledge transfer. In practical terms, they can do internet-based or on site briefings accompanied by practical hands on demonstrations and guidance.

Again, this is largely driven by requirements of the customer. REAN typically emails a deliverable to the customer and schedules a walkthrough of the deliverable if required. Sometimes acceptance testing is required as well. Customer is usually limited to a certain period of time it has to review deliverable and confirm acceptance or rejection of that deliverable.

Microsoft: Offeror's subcontractor, Microsoft, currently, as of the date of the Proposal, has a mechanism by which 30-day Trial subscriptions may be ordered for some, but not all, of the cloud services offered hereunder. Microsoft will provide additional information about this upon request of Lead State, Participating States, or any Purchasing Entity.

It is possible for any Purchasing Entity to purchase a separate subscription for the purpose of establishing a second environment for test and/or staging purposes. Such separate Subscription would be at an additional cost, and additional contract paperwork may be required.

8.18.2 Describe how you intend to provide a test and/ or proof of concept environment for evaluation that verifies your ability to meet mandatory requirements.

Insight Response: For all Insight services offered around Microsoft cloud technologies, Insight has the ability to stand up a test environment that is identical to the production environment and built on the same platform as the production environment but totally separated. This gives us the ability to do a proof of concept that will totally reflect the production environment to provide the most optimal testing possible.

8.18.3 Offeror must describe what training and support it provides at no additional cost.

Insight Response: Insight offers full admin training and pre-sales support and limited post-sales support for both Azure and Office 365. This training and support of offered at no cost to the customer.

8.19 (E) INTEGRATION AND CUSTOMIZATION

8.19.1 Describe how the Solutions you provide can be integrated to other complementary applications, and if you offer standard-based interface to enable additional integrations.

Insight Response: Insight and REAN will help customers carefully consider and choose the right services and help with the integration of those services into customer's IT environment, and applicable laws and regulations. We will help customer enhance security and/or meet more stringent compliance requirements by leveraging technology such as host-based firewalls, host-based intrusion detection/prevention and encryption.

Microsoft Cloud Solutions are specifically designed to support third-party technologies, and were built with standards-based interfaces to enable integrations with these non-Microsoft tools. Many of these third-party technologies were built to run on Microsoft Windows platforms in the first place, and the Windows in the Azure Cloud and in Office 365 is the same as the Windows that these technologies were designed for. Microsoft even makes previous versions of Windows and SQL Server available in Azure, so that older applications that were designed for these previous versions can run in Azure as well.

Microsoft: Microsoft Cloud Solutions are specifically designed to support third-party technologies, and were built with standards-based interfaces to enable integrations with these non-Microsoft tools. Many of these third-party technologies were built to run on Microsoft Windows platforms in the first place, and the Windows in the Azure Cloud and in Office 365 is the same as the Windows that these technologies were designed for. Microsoft even makes previous versions of Windows and SQL Server available in Azure, so that older applications that were designed for these previous versions can run in Azure as well.

8.19.2 Describe the ways to customize and personalize the Solutions you provide to meet the needs of specific Purchasing Entities.

Insight Response: From proof of concept to fully, secured and operational cloud environments, REAN Cloud provided pragmatic solutions based on best practices for performance, security, compliance, and cost in a phased approach customized to the customer's specific needs.

REAN Cloud provides customized operations support through its Managed Services (MGS) offering, which provides the customer with 24x7x365 enterprise technical support.

Microsoft technology has always been developed with the end user in mind, so Microsoft Cloud Solutions are highly customizable and easily adaptable to the needs of the end user or their organization. Whether the Purchasing Entities want a SaaS, PaaS or IaaS solution, whether they want all or part of their business productivity tools in the Cloud, and whether they want to access it using Windows PCs or iOS devices, Microsoft Cloud Solutions are available. Purchasing Entities have a great deal of choice in the size, speed and performance level of all Cloud services they deploy, and have several options for business productivity tools as well. And Microsoft is adding new options, new services and new features every month.

Microsoft: Microsoft technology has always been developed with the end user in mind, so Microsoft Cloud Solutions are highly customizable and easily adaptable to the needs of the end user or their organization. Whether the Purchasing Entities want a SaaS, PaaS or IaaS solution, whether they want all or part of their business productivity tools in the Cloud, and whether they want to access it using Windows PCs or iOS devices, Microsoft Cloud Solutions are available. Purchasing Entities have a great deal of choice in the size, speed and performance level of all

Cloud services they deploy, and have several options for business productivity tools as well. And Microsoft is adding new options, new services and new features every month.

8.20 (E) MARKETING PLAN

Describe your how you intend to market your Services to NASPO ValuePoint and Participating Entities.

Insight Response: Insight will partner with the NASPO ValuePoint business development office to help increase business transacted through the contract, increase number of sales leads, expand the number of customer relationships beyond main point-of-contact, increase digital presence, acquire new customers by highlighting Insight and NASPO ValuePoint's value proposition, and drive current Insight SLED clients toward the contract. Our national, established working relationship across many SLED organizations, and our deep knowledge of NASPO ValuePoint contract rules and regulations, requirements, and initiatives combined with our experienced sales and services teams give us the differentiators to help the State of Utah and NASPO ValuePoint make the Cloud Solutions contract a success.

We strongly believe that our capabilities differentiate Insight from others, and it is these distinguishing factors that will contribute to growth of the Cloud Solutions contract. Insight will prepare a marketing outreach plan that will include information on training that is currently available through NASPO ValuePoint to our customers and other referrals. In addition, through a digital presence and links to the NASPO ValuePoint website, Insight will develop a trifold that is co-branded with Insight's and NASPO ValuePoint's logo to promote the contract and provide customers with information on the contract vehicle and services that NASPO ValuePoint offers.

Insight Public Sector's strategy for marketing and selling the cloud solutions that are being responded to in this RFP to eligible NASPO ValuePoint Participating Entities includes both dedicated local resources to this contract, the Insight Public Sector dedicated public sector sales and delivery organization, and our centralized cloud software/product specialists organization.

Dedicated Public Sector Sales Team - Insight Public Sector has maintained a dedicated public sector sales team for over eighteen years. Many of the Insight Public Sector account executives across the nation have multiple years of seasoned experience in utilizing NASPO ValuePoint contracts for their customers. Sales team members maintain specific accounts with public sector entities, consistently maintaining over a 95% customer satisfaction rating.

Insight Public Sector's current sales force is comprised of three areas: Inside Sales Executives, Field Based Sales Executives and Sales Support Representatives. The Public Sector Sales Organization has been dedicated to K-12 Education, Universities, and State/Local government agencies across the nation for over eighteen (18) years.

Educating Insight Public Sector Staff:

Upon award of contract, Insight Public Sector will educate all Insight Public Sector staff of the capabilities and requirements set forth in the Cloud Solutions contract. Insight Public Sector will assign one (1) dedicated Point of Contact (POC) to be the Contract Manager for all requests made under this contract. The Contract Manager will be the primary communicator to the Insight Public Sector Sales Organization, and this person will monitor this contract closely and add additional support as needed.

Educating Customers - Upon award of contract, Insight Public Sector will implement a marketing plan to educate current and potential customers on the capabilities and requirements set forth in the NASPO ValuePoint Cloud Solutions contract. In conjunction with our partners, Insight Public Sector proposes to present the following campaign:

Marketing Collateral – Insight Public Sector has specific marketing collateral in place for all Insight Public Sector representatives to use in marketing cloud services to Participating Entities - state agencies, cities, counties, K-12 and Higher Education.

Cloud Software/ Product Sales Specialist(s) – Due to the complex nature of cloud solution sales, Insight Public Sector will staff specific consulting experts to assist sales representatives and pre-sales engineering staff with necessary information and pricing structures that will ensure the Participating Entity is receiving clear, concise proposals that are tailored to each entities specific need.

Sales Calls: Insight Public Sector Account Executives will inform all eligible customers of the advantages of utilizing the Cloud Solutions contract and will reinforce the results of the phone campaign by making Face-to-Face sales calls.

8.21 (E) RELATED VALUE-ADDED SERVICES TO CLOUD SOLUTIONS

Describe the valued-added services that you can provide as part of an awarded contract, e.g. consulting services pre- and post- implementation. Offerors may detail professional services in the RFP limited to assisting offering activities with initial setup, training and access to the services.

Insight Response: As Participating States and Entity's partner, Insight offers Courtesy services for both Azure and Office 365. These services include full admin training, initial tenant access and setup, user creation, virtual machine creation for a limited number of VMs, storage creation, and creation of the required virtual networks. These services are to help our customer to quickly understand and start using these products.

8.22 (E) SUPPORTING INFRASTRUCTURE

8.22.1 Describe what infrastructure is required by the Purchasing Entity to support your Solutions or deployment models.

Insight Response: Microsoft Cloud Solutions will work with little or no infrastructure, since it provides its own infrastructure within the Microsoft Cloud. End users need to have reliable access to the Internet and compatible devices (generally wireless-ready PCs, iOS and Android devices). But other infrastructure elements exist in Microsoft datacenters around the world. If the Purchasing Entity isn't comfortable with sending and receiving information over the public Internet, Microsoft offers a private, dedicated connection service called Azure ExpressRoute.

8.22.2 If required, who will be responsible for installation of new infrastructure and who will incur those costs?

Insight Response: To the extent that any infrastructure is required, the Purchasing Entity will be responsible for providing and maintaining it. But Microsoft Cloud Solutions are designed to alleviate the need for an on-premise infrastructure, so Purchasing Entities should expect to see their IT infrastructure costs go down as they move to these solutions.

8.23 (E) ALIGNMENT OF CLOUD COMPUTING REFERENCE ARCHITECTURE

Clarify how their architecture compares to the NIST Cloud Computing Reference Architecture, in particular, to describe how they align with the three domains e.g. Infrastructure as a Service (IaaS), Software as a Service (SaaS), and Platform as a Service (PaaS).

Insight Response: Azure is primarily a PaaS offering, with a large variety of SaaS services available. While some components of Azure fall into NIST's definition of IaaS, the chief discriminator is Physical Assets, of which Azure provides none. Functionality and redundancy of physical hardware (including plant, environment, power, network and access) is provided, which meets the fiduciary definition of IaaS. Customers often have a choice between PaaS and SaaS offerings to meet their business objectives.

Office 365 is a SaaS service for a group of software plus services subscriptions that provides productivity software and related services to its subscribers. For consumers, the service allows the use of Microsoft Office apps on Windows and OS X, provides storage space on Microsoft's cloud storage service OneDrive, and grants 60 Skype minutes per month. For business and enterprise users, Office 365 offers plans including e-mail and social networking services through hosted versions of Exchange Server, Skype for Business Server, SharePoint and Office Online, integration with Yammer, as well as access to the Office software.

7. Confidential, Protected, or Proprietary Information

Insight Response: None.

8. Exceptions and/or Additions to the Standard Terms and Conditions

Insight Response: Per the requirements of the RFP, Insight has included proposed exceptions and/or additions to the master Agreement Terms and Conditions, including the exhibits, in response to this section. We have included this as a separate attachment to the main response document.


Attachment E


E1 AWS BAA


E2 MS BAA

E3 MS SLA


AWS BUSINESS ASSOCIATE ADDENDUM











HIPAA Business Associate Agreement

If Customer is a Covered Entity or a Business Associate and includes Protected Health Information in Customer Data (as such terms are defined below), execution of a license agreement that includes the Online Services Terms ("Agreement") will incorporate the terms of this HIPAA Business Associate Agreement ("BAA") into that Agreement. If there is any conflict between a provision in this BAA and a provision in the Agreement, this BAA will control.

1. Definitions.

Except as otherwise defined in this BAA, capitalized terms shall have the definitions set forth in HIPAA and Customer's Agreement.

"Breach Notification Rule" means the Breach Notification for Unsecured Protected Health Information Final Rule.

"Business Associate" shall have the same meaning as the term "business associate" in 45 CFR § 160.103 of HIPAA.

"Covered Entity" shall have the same meaning as the term "covered entity" in 45 CFR § 160.103 of HIPAA.

"Dynamics CRM Online Services" means Dynamics CRM Online services made available through volume licensing or the Microsoft online services portal, excluding Dynamics CRM for supported devices, which includes but is not limited to Dynamics CRM Online services for tablets and/or smartphones and any separately branded service made available with or connected to Dynamics CRM Online such as Microsoft Social Engagement, Parature, from Microsoft, and Microsoft Dynamics Marketing.

"HIPAA" collectively means the administrative simplification provision of the Health Insurance Portability and Accountability Act enacted by the United States Congress, and its implementing regulations, including the Privacy Rule, the Breach Notification Rule, and the Security Rule, as amended from time to time, including by the Health Information Technology for Economic and Clinical Health ("HITECH") Act and by the Modifications to the HIPAA Privacy, Security, Enforcement, and Breach Notification Rules under the Health Information Technology for Economic and Clinical Health Act and the Genetic Information Nondiscrimination Act; Other Modifications to the HIPAA Rules; Final Rule.

"Microsoft Azure Core Services" means the following features of Microsoft Azure Services: Cloud Services (web and worker roles), Virtual Machines (including with SQL Server), Storage (Blobs, Tables, Queues), Virtual Network, Traffic Manager, Batch, Web Sites, BizTalk Services, Media Services, Mobile Services, Service Bus, Notification Hub, Workflow Manager, Express Route, Scheduler, Multi-Factor Authentication, Active Directory, Rights Management Service, SQL Database, HDInsight and any other features identified as included on the Microsoft Azure Trust Center.

"Microsoft Intune Online Services" means the cloud service portion of Microsoft Intune such as the Microsoft Intune Add-on Product or a management service provided by Microsoft Intune such as Mobile Device Management for Office 365. It does not include any on-premises software made available with a Microsoft Intune subscription.

"Microsoft Online Services," for this BAA only, means Microsoft Dynamics CRM Online Services, Office 365 Services, Microsoft Azure Core Services, and/or Microsoft Intune.

“Office 365 Services” means the following services, each as a standalone service or as included in an Office 365-branded plan or suite: Exchange Online, Exchange Online Archiving, Exchange Online Protection, Advanced Threat Protection, SharePoint Online, OneDrive for Business, Project Online, Skype for Business Online, Sway, Office Online, and Yammer Enterprise. Office 365 Services do not include Office 365 ProPlus, any portion of PSTN Services that operate outside of Microsoft’s control, any client software, or any separately branded service made available with an Office 365-branded plan or suite, such as a Bing or a service branded “for Office 365.”

“Privacy Rule” means the Standards for Privacy of Individually Identifiable Health Information.

“Protected Health Information” shall have the same meaning as the term “protected health information” in 45 CFR § 160.103 of HIPAA, provided that it is limited to such protected health information that is received by Microsoft from, or created, received, maintained, or transmitted by Microsoft on behalf of, Customer.

“Security Rule” means the Security Standards for the Protection of Electronic Protected Health Information.

2. *Permitted Uses and Disclosures of Protected Health Information.*

- a. Performance of the Agreement for Microsoft Online Services.** Except as otherwise limited in this BAA, Microsoft may Use and Disclose Protected Health Information for, or on behalf of, Customer as specified in the Agreement; provided that any such Use or Disclosure would not violate HIPAA if done by Customer, unless expressly permitted under paragraph b of this Section.
- b. Management, Administration, and Legal Responsibilities.** Except as otherwise limited in this BAA, Microsoft may Use and Disclose Protected Health Information for the proper management and administration of Microsoft and/or to carry out the legal responsibilities of Microsoft, provided that any Disclosure may occur only if: (1) Required by Law; or (2) Microsoft obtains written reasonable assurances from the person to whom the Protected Health Information is Disclosed that it will be held confidentially and Used or further Disclosed only as Required by Law or for the purpose for which it was Disclosed to the person, and the person notifies Microsoft of any instances of which it becomes aware in which the confidentiality of the Protected Health Information has been breached.

3. *Responsibilities of the Parties with Respect to Protected Health Information.*

- a. Microsoft’s Responsibilities.** To the extent Microsoft is acting as a Business Associate, Microsoft agrees to the following:
 - (i) Limitations on Use and Disclosure.** Microsoft shall not Use and/or Disclose the Protected Health Information other than as permitted or required by the Agreement and/or this BAA or as otherwise Required by Law; Microsoft shall not disclose, capture, maintain, scan, index, transmit, share or Use Protected Health Information for any activity not authorized under the Agreement and/or this BAA. Microsoft Online Services shall not use Protected Health Information for any advertising, Marketing or other commercial purpose of Microsoft or any third party. Microsoft shall not violate the HIPAA prohibition on the sale of Protected Health Information. Microsoft shall make reasonable efforts to Use, Disclose, and/or request the

minimum necessary Protected Health Information to accomplish the intended purpose of such Use, Disclosure, or request.

- (ii) **Safeguards.** Microsoft shall: (1) use reasonable and appropriate safeguards to prevent inappropriate Use and Disclosure of Protected Health Information other than as provided for in this BAA; and (2) comply with the applicable requirements of 45 CFR Part 164 Subpart C of the Security Rule.
- (iii) **Reporting.** Microsoft shall report to Customer: (1) any Use and/or Disclosure of Protected Health Information that is not permitted or required by this BAA of which Microsoft becomes aware; (2) any Security Incident of which it becomes aware, provided that notice is hereby deemed given for Unsuccessful Security Incidents and no further notice of such Unsuccessful Security Incidents shall be given; and/or (3) any Breach of Customer's Unsecured Protected Health Information that Microsoft may discover (in accordance with 45 CFR § 164.410 of the Breach Notification Rule). Notification of a Breach will be made without unreasonable delay, but in no event more than thirty (30) calendar days after discovery of a Breach. Taking into account the level of risk reasonably likely to be presented by the Use, Disclosure, Security Incident, or Breach, the timing of other reporting will be made consistent with Microsoft's and Customer's legal obligations.

For purposes of this Section, "Unsuccessful Security Incidents" mean, without limitation, pings and other broadcast attacks on Microsoft's firewall, port scans, unsuccessful log-on attempts, denial of service attacks, and any combination of the above, as long as no such incident results in unauthorized access, acquisition, Use, or Disclosure of Protected Health Information. Notification(s) under this Section, if any, will be delivered to contacts identified by Customer pursuant to Section 3b(ii) (Contact Information for Notices) of this BAA by any means Microsoft selects, including through e-mail. Microsoft's obligation to report under this Section is not and will not be construed as an acknowledgement by Microsoft of any fault or liability with respect to any Use, Disclosure, Security Incident, or Breach.

- (iv) **Subcontractors.** In accordance with 45 CFR §§ 164.502(e)(1)(ii) and 164.308(b)(2) of HIPAA, Microsoft shall require its Subcontractors who create, receive, maintain, or transmit Protected Health Information on behalf of Microsoft to agree in writing to: (1) the same or more stringent restrictions and conditions that apply to Microsoft with respect to such Protected Health Information; (2) appropriately safeguard the Protected Health Information; and (3) comply with the applicable requirements of 45 CFR Part 164 Subpart C of the Security Rule. Microsoft remains responsible for its subcontractors' compliance with obligations in this BAA.
- (v) **Disclosure to the Secretary.** Microsoft shall make available its internal practices, records, and books relating to the Use and/or Disclosure of Protected Health Information received from Customer to the Secretary of the Department of Health and Human Services for purposes of determining Customer's compliance with HIPAA, subject to attorney-client and other applicable legal privileges. Microsoft shall respond to any such request from the Secretary in accordance with the Section titled "Disclosure of Customer Data" in the Agreement.
- (vi) **Access.** If Microsoft maintains Protected Health Information in a Designated Record Set for Customer, then Microsoft, at the request of Customer, shall within

fifteen (15) days make access to such Protected Health Information available to Customer in accordance with 45 CFR § 164.524 of the Privacy Rule.

(vii) Amendment. If Microsoft maintains Protected Health Information in a Designated Record Set for Customer, then Microsoft, at the request of Customer, shall make available such Protected Health Information to Customer for amendment and incorporate any reasonably requested amendment in the Protected Health Information in accordance with 45 CFR § 164.526 of the Privacy Rule.

(viii) Accounting of Disclosure. Microsoft, at the request of Customer, shall within fifteen (15) days make available to Customer such information relating to Disclosures made by Microsoft as required for Customer to make any requested accounting of Disclosures in accordance with 45 CFR § 164.528 of the Privacy Rule.

(ix) Performance of a Covered Entity's Obligations. To the extent Microsoft is to carry out a Covered Entity obligation under the Privacy Rule, Microsoft shall comply with the requirements of the Privacy Rule that apply to Customer in the performance of such obligation.

b. Customer Responsibilities.

(i) No Impermissible Requests. Customer shall not request Microsoft to Use or Disclose Protected Health Information in any manner that would not be permissible under HIPAA if done by a Covered Entity (unless permitted by HIPAA for a Business Associate).

(ii) Contact Information for Notices. Customer hereby agrees that any reports, notification, or other notice by Microsoft pursuant to this BAA may be made electronically. Customer shall provide contact information to MSO-HIPAA@microsoft.com or such other location or method of updating contact information as Microsoft may specify from time to time and shall ensure that Customer's contact information remains up to date during the term of this BAA. Contact information must include name of individual(s) to be contacted, title of individual(s) to be contacted, e-mail address of individual(s) to be contacted, name of Customer organization, and, if available, either contract number or subscriber identification number.

(iii) Safeguards and Appropriate Use of Protected Health Information. Customer is responsible for implementing appropriate privacy and security safeguards to protect its Protected Health Information in compliance with HIPAA. Without limitation, it is Customer's obligation to:

- 1) Not include Protected Health Information in: (1) information Customer submits to technical support personnel through a technical support request or to community support forums; and (2) Customer's address book or directory information. In addition, Microsoft does not act as, or have the obligations of, a Business Associate under HIPAA with respect to Customer Data once it is sent to or from Customer outside Microsoft Online Services over the public Internet.
- 2) Implement privacy and security safeguards in the systems, applications, and software Customer controls, configures, and uploads into the Microsoft Online Services.

4. *Applicability of BAA.*

This BAA is applicable to Microsoft Online Services. Microsoft may, from time to time, update the definition of Microsoft Online Services in this BAA to include additional Microsoft online services. Any such updated definitions will apply to Customer without additional action by Customer. It is Customer's obligation to not store or process Protected Health Information in a Microsoft online service until this BAA is effective as to the applicable service.

5. *Term and Termination.*

- a. **Term.** This BAA shall continue in effect until the earlier of (1) termination by a Party for breach as set forth in Section 5b, below, or (2) expiration of Customer's Agreement.
- b. **Termination for Breach.** Upon written notice, either Party immediately may terminate the Agreement and this BAA if the other Party is in material breach or default of any obligation in this BAA. Either party may provide the other a thirty (30) calendar day period to cure a material breach or default within such written notice.
- c. **Return, Destruction, or Retention of Protected Health Information Upon Termination.** Upon expiration or termination of this BAA, Microsoft shall return or destroy all Protected Health Information in its possession, if it is feasible to do so, and as set forth in the applicable termination provisions of the Agreement. If it is not feasible to return or destroy any portions of the Protected Health Information upon termination of this BAA, then Microsoft shall extend the protections of this BAA, without limitation, to such Protected Health Information and limit any further Use or Disclosure of the Protected Health Information to those purposes that make the return or destruction infeasible for the duration of the retention of the Protected Health Information.

6. *Miscellaneous.*

- a. **Interpretation.** The Parties intend that this BAA be interpreted consistently with their intent to comply with HIPAA and other applicable federal and state law. Except where this BAA conflicts with the Agreement, all other terms and conditions of the Agreement remain unchanged. Any captions or headings in this BAA are for the convenience of the Parties and shall not affect the interpretation of this BAA.
- b. **BAAs; Waiver.** This BAA may not be modified or amended except in a writing duly signed by authorized representatives of the Parties. A waiver with respect to one event shall not be construed as continuing, as a bar to, or as a waiver of any right or remedy as to subsequent events.
- c. **No Third Party Beneficiaries.** Nothing express or implied in this BAA is intended to confer, nor shall anything in this BAA confer, upon any person other than the Parties, and the respective successors or assigns of the Parties, any rights, remedies, obligations, or liabilities whatsoever.
- d. **Severability.** In the event that any provision of this BAA is found to be invalid or unenforceable, the remainder of this BAA shall not be affected thereby, but rather the remainder of this BAA shall be enforced to the greatest extent permitted by law.
- e. **No Agency Relationship.** It is not intended that an agency relationship (as defined under the Federal common law of agency) be established hereby expressly or by implication between Customer and Microsoft under HIPAA or the Privacy Rule,

Security Rule, or Breach Notification Rule. No terms or conditions contained in this BAA shall be construed to make or render Microsoft an agent of Customer.

Volume
Licensing

Service Level Agreement for Microsoft Online Services February 1, 2016

Table of Contents

TABLE OF CONTENTS	2	HDINSIGHT	21
INTRODUCTION	3	HOCKEYAPP	22
ABOUT THIS DOCUMENT	3	KEY VAULT	22
PRIOR VERSIONS OF THIS DOCUMENT	3	MACHINE LEARNING – BATCH EXECUTION SERVICE (BES) AND MANAGEMENT APIS SERVICE	23
CLARIFICATIONS AND SUMMARY OF CHANGES TO THIS DOCUMENT	3	MACHINE LEARNING – REQUEST RESPONSE SERVICE (RRS)	23
GENERAL TERMS	4	MEDIA SERVICES – CONTENT PROTECTION SERVICE	23
DEFINITIONS	4	MEDIA SERVICES – ENCODING SERVICE	24
TERMS	4	MEDIA SERVICES – INDEXER SERVICE	24
SERVICE SPECIFIC TERMS	6	MEDIA SERVICES – LIVE CHANNELS	25
MICROSOFT DYNAMICS	6	MEDIA SERVICES – STREAMING SERVICE	25
MICROSOFT DYNAMICS AX	6	MOBILE ENGAGEMENT	26
MICROSOFT DYNAMICS CRM	6	MOBILE SERVICES	26
OFFICE 365 SERVICES	7	MULTI-FACTOR AUTHENTICATION SERVICE	26
DUET ENTERPRISE ONLINE	7	OPERATIONAL INSIGHTS	27
EXCHANGE ONLINE	7	REMOTEAPP	27
EXCHANGE ONLINE ARCHIVING	8	SCHEDULER	28
EXCHANGE ONLINE PROTECTION	8	SEARCH	28
OFFICE 365 BUSINESS	8	SERVICE-BUS SERVICE – EVENT HUBS	29
OFFICE 365 CUSTOMER LOCKBOX	9	SERVICE-BUS SERVICE – NOTIFICATION HUBS	29
OFFICE 365 PROPLUS	9	SERVICE-BUS SERVICE – QUEUES AND TOPICS	30
OFFICE ONLINE	9	SERVICE-BUS SERVICE – RELAYS	30
OFFICE 365 VIDEO	10	SITE RECOVERY SERVICE – ON-PREMISES-TO-AZURE	30
ONEDRIVE FOR BUSINESS	10	SITE RECOVERY SERVICE – ON-PREMISES-TO-ON-PREMISES	31
PROJECT ONLINE	10	SQL DATABASE SERVICE (BASIC, STANDARD AND PREMIUM TIERS)	31
SHAREPOINT ONLINE	11	SQL DATABASE SERVICE (WEB AND BUSINESS TIERS)	32
SKYPE FOR BUSINESS ONLINE	11	STORAGE SERVICE	32
SKYPE FOR BUSINESS ONLINE – PSTN CALLING AND PSTN CONFERENCING	12	STORSIMPLE SERVICE	33
SKYPE FOR BUSINESS ONLINE – VOICE QUALITY	12	STREAM ANALYTICS – API CALLS	34
YAMMER ENTERPRISE	12	STREAM ANALYTICS – JOBS	34
ENTERPRISE MOBILITY SERVICES	13	TRAFFIC MANAGER SERVICE	35
AZURE ACTIVE DIRECTORY BASIC	13	VIRTUAL MACHINES	35
AZURE ACTIVE DIRECTORY PREMIUM	13	VPN GATEWAY	36
AZURE RIGHTS MANAGEMENT	14	VISUAL STUDIO ONLINE – BUILD SERVICE	36
MICROSOFT INTUNE	14	VISUAL STUDIO ONLINE – LOAD TESTING SERVICE	37
MICROSOFT AZURE SERVICES	14	VISUAL STUDIO ONLINE – USER PLANS SERVICE	37
API MANAGEMENT SERVICES	14	OTHER ONLINE SERVICES	38
APP SERVICE	15	BING MAPS ENTERPRISE PLATFORM	38
APPLICATION GATEWAY	16	BING MAPS MOBILE ASSET MANAGEMENT	38
AUTOMATION SERVICE	16	POWER BI PRO	39
BACKUP SERVICE	16	TRANSLATOR API	39
BATCH SERVICE	17		
BIZTALK SERVICES	17	APPENDIX A – SERVICE LEVEL COMMITMENT FOR VIRUS DETECTION AND BLOCKING, SPAM EFFECTIVENESS, OR FALSE POSITIVE	40
CACHE SERVICES	18		
CDN SERVICE	19	APPENDIX B - SERVICE LEVEL COMMITMENT FOR UPTIME AND EMAIL DELIVERY	41
CLOUD SERVICES	19		
DATA FACTORY – ACTIVITY RUNS	20		
DATA FACTORY – API CALLS	20		
DOCUMENTDB	20		
EXPRESSROUTE	21		

Introduction

About this Document

This Service Level Agreement for Microsoft Online Services (this “SLA”) is a part of your Microsoft volume licensing agreement (the “Agreement”). Capitalized terms used but not defined in this SLA will have the meaning assigned to them in the Agreement. This SLA applies to the Microsoft Online Services listed herein (a “Service” or the “Services”), but does not apply to separately branded services made available with or connected to the Services or to any on-premise software that is part of any Service.

If we do not achieve and maintain the Service Levels for each Service as described in this SLA, then you may be eligible for a credit towards a portion of your monthly service fees. We will not modify the terms of your SLA during the initial term of your subscription; however, if you renew your subscription, the version of this SLA that is current at the time of renewal will apply throughout your renewal term. We will provide at least 90 days’ notice for adverse material changes to this SLA. You can review the most current version of this SLA at any time by visiting <http://www.microsoftvolumelicensing.com/SLA>.

Prior Versions of this Document

This SLA provides information on Services currently available. Earlier versions of this document are available at <http://www.microsoftvolumelicensing.com>. To find the needed version, a customer may contact its reseller or Microsoft Account Manager.

Clarifications and Summary of Changes to this Document

Below are recent additions, deletions and other changes to this SLA. Also listed below, are clarifications of Microsoft policy in response to common customer questions.

Additions	Deletions
Microsoft Dynamics AX	Skype for Business Online – Cloud PBX
Skype for Business Online – Voice Quality	Skype for Business Online – PSTN Conferencing
HockeyApp	

Service Specific Terms

[Skype for Business Online – PSTN Calling](#): The Skype for Business Online – PSTN Calling entry and Skype for Business Online – PSTN Conferencing entry were combined into a single entry.

[Table of Contents / Definitions](#)

[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

General Terms

Definitions

“Applicable Monthly Period” means, for a calendar month in which a Service Credit is owed, the number of days that you are a subscriber for a Service.

“Applicable Monthly Service Fees” means the total fees actually paid by you for a Service that are applied to the month in which a Service Credit is owed.

“Downtime” is defined for each Service in the Services Specific Terms below. Except for Microsoft Azure Services, Downtime does not include Scheduled Downtime. Downtime does not include unavailability of a Service due to limitations described below and in the Services Specific Terms.

“Error Code” means an indication that an operation has failed, such as an HTTP status code in the 5xx range.

“External Connectivity” is bi-directional network traffic over supported protocols such as HTTP and HTTPS that can be sent and received from a public IP address.

“Incident” means (i) any single event, or (ii) any set of events, that result in Downtime.

“Management Portal” means the web interface, provided by Microsoft, through which customers may manage the Service.

“Scheduled Downtime” means periods of Downtime related to network, hardware, or Service maintenance or upgrades. We will publish notice or notify you at least five (5) days prior to the commencement of such Downtime.

“Service Credit” is the percentage of the Applicable Monthly Service Fees credited to you following Microsoft’s claim approval.

“Service Level” means the performance metric(s) set forth in this SLA that Microsoft agrees to meet in the delivery of the Services.

“Service Resource” means an individual resource available for use within a Service.

“Success Code” means an indication that an operation has succeeded, such as an HTTP status code in the 2xx range.

“Support Window” refers to the period of time during which a Service feature or compatibility with a separate product or service is supported.

“User Minutes” means the total number of minutes in a month, less all Scheduled Downtime, multiplied by the total number of users.

Terms

Claims

In order for Microsoft to consider a claim, you must submit the claim to customer support at Microsoft Corporation including all information necessary for Microsoft to validate the claim, including but not limited to: (i) a detailed description of the Incident; (ii) information regarding the time and duration of the Downtime; (iii) the number and location(s) of affected users (if applicable); and (iv) descriptions of your attempts to resolve the Incident at the time of occurrence.

For a claim related to Microsoft Azure, we must receive the claim within two months of the end of the billing month in which the Incident that is the subject of the claim occurred. For claims related to all other Services, we must receive the claim by the end of the calendar month following the month in which the Incident occurred. For example, if the Incident occurred on February 15th, we must receive the claim and all required information by March 31st.

We will evaluate all information reasonably available to us and make a good faith determination of whether a Service Credit is owed. We will use commercially reasonable efforts to process claims during the subsequent month and within forty-five (45) days of receipt. You must be in compliance with the Agreement in order to be eligible for a Service Credit. If we determine that a Service Credit is owed to you, we will apply the Service Credit to your Applicable Monthly Service Fees.

If you purchased more than one Service (not as a suite), then you may submit claims pursuant to the process described above as if each Service were covered by an individual SLA. For example, if you purchased both Exchange Online and SharePoint Online (not as part of a suite), and during the term of the subscription an Incident caused Downtime for both Services, then you could be eligible for two separate Service Credits (one for each Service), by submitting two claims under this SLA. In the event that more than one Service Level for a particular Service is not met because of the same Incident, you must choose only one Service Level under which to make a claim based on the Incident.

Service Credits

Service Credits are your sole and exclusive remedy for any performance or availability issues for any Service under the Agreement and this SLA. You may not unilaterally offset your Applicable Monthly Service Fees for any performance or availability issues.

Service Credits apply only to fees paid for the particular Service, Service Resource, or Service tier for which a Service Level has not been met. In cases where Service Levels apply to individual Service Resources or to separate Service tiers, Service Credits apply only to fees paid for the affected

Service Resource or Service tier, as applicable. The Service Credits awarded in any billing month for a particular Service or Service Resource will not, under any circumstance, exceed your monthly service fees for that Service or Service Resource, as applicable, in the billing month. If you purchased Services as part of a suite or other single offer, the Applicable Monthly Service Fees and Service Credit for each Service will be pro-rated.

If you purchased a Service from a reseller, you will receive a service credit directly from your reseller and the reseller will receive a Service Credit directly from us. The Service Credit will be based on the estimated retail price for the applicable Service, as determined by us in our reasonable discretion.

Limitations

This SLA and any applicable Service Levels do not apply to any performance or availability issues:

1. Due to factors outside our reasonable control (for example, natural disaster, war, acts of terrorism, riots, government action, or a network or device failure external to our data centers, including at your site or between your site and our data center);
2. That result from the use of services, hardware, or software not provided by us, including, but not limited to, issues resulting from inadequate bandwidth or related to third-party software or services;
3. Caused by your use of a Service after we advised you to modify your use of the Service, if you did not modify your use as advised;
4. During or with respect to preview, pre-release, beta or trial versions of a Service, feature or software (as determined by us) or to purchases made using Microsoft subscription credits;
5. That result from your unauthorized action or lack of action when required, or from your employees, agents, contractors, or vendors, or anyone gaining access to our network by means of your passwords or equipment, or otherwise resulting from your failure to follow appropriate security practices;
6. That result from your failure to adhere to any required configurations, use supported platforms, follow any policies for acceptable use, or your use of the Service in a manner inconsistent with the features and functionality of the Service (for example, attempts to perform operations that are not supported) or inconsistent with our published guidance;
7. That result from faulty input, instructions, or arguments (for example, requests to access files that do not exist);
8. That result from your attempts to perform operations that exceed prescribed quotas or that resulted from our throttling of suspected abusive behavior;
9. Due to your use of Service features that are outside of associated Support Windows; or
10. For licenses reserved, but not paid for, at the time of the Incident.

Services purchased through Open, Open Value, and Open Value Subscription volume licensing agreements, and Services in an Office 365 Small Business Premium suite purchased in the form of a product key are not eligible for Service Credits based on service fees. For these Services, any Service Credit that you may be eligible for will be credited in the form of service time (i.e., days) as opposed to service fees, and any references to "Applicable Monthly Service Fees" is deleted and replaced by "Applicable Monthly Period."

[Table of Contents](#) / [Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Service Specific Terms

Microsoft Dynamics

Microsoft Dynamics AX

Additional Definitions:

"Active Tenant" means a tenant with an active high availability production topology in the Management Portal that (A) has been deployed to a Partner Application Service; and (B) has an active database that users can log into.

"Partner Application Service" means a partner application built on top of and combined with the Platform that (A) is used for processing your organization's actual business transactions; and (B) has reserve compute and storage resources equal to or greater than one of the Scale Units your partner selected for the applicable partner application.

"Maximum Available Minutes" means the total accumulated minutes during a billing month in which an Active Tenant was deployed in a Partner Application Service using an active high availability production topology.

"Platform" means the Service's client forms, SQL server reports, batched operations, and API endpoints, or the Service's retail APIs that are used for commerce or retail purposes only.

"Scale Unit" means the increments by which compute and storage resources are added to or removed from a Partner Application Service.

"Service Infrastructure" means the authentication, computing, and storage resources that Microsoft provides in connection with the Service.

Downtime: Any period of time when end users are unable to login to their Active Tenant, due to a failure in the unexpired Platform or the Service Infrastructure as Microsoft determines from automated health monitoring and system logs. Downtime does not include Scheduled Downtime, the unavailability of Service add-on features, the inability to access the Service due to your modifications of the Service, or periods where the Scale Unit capacity is exceeded.

Monthly Uptime Percentage: The Monthly Uptime Percentage for a given Active Tenant in a calendar month is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.5%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Microsoft Dynamics CRM

Downtime: Any period of time when end users are unable to read or write any Service data for which they have appropriate permission but this does not include non-availability of Service add-on features.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

[Table of Contents / Definitions](#)

Office 365 Services

Duet Enterprise Online

Downtime: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply when the inability to read or write any portion of a SharePoint Online site is caused by any failure of third party software, equipment, or services that are not controlled by Microsoft, or Microsoft software that is not being run by Microsoft itself as part of the Service.

Additional Terms: You will be eligible for a Service Credit for Duet Enterprise Online only when you are eligible for a Service Credit for the SharePoint Online Plan 2 User SLs that you have purchased as a prerequisite for your Duet Enterprise Online User SLs.

[Table of Contents / Definitions](#)

Exchange Online

Downtime: Any period of time when users are unable to send or receive email with Outlook Web Access.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Additional Terms: See Appendix 1 – Service Level Commitment for Virus Detection and Blocking, Spam Effectiveness, or False Positive.

[Table of Contents / Definitions](#)[Table of Contents](#)[Introduction](#)[General Terms](#)[Service Specific Terms](#)[Appendices](#)

Exchange Online Archiving

Downtime: Any period of time when users are unable to access the email messages stored in their archive.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to the Enterprise CAL suite purchased through Open Value and Open Value Subscription volume licensing agreements.

[Table of Contents / Definitions](#)

Exchange Online Protection

Downtime: Any period of time when the network is not able to receive and process email messages.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to the Enterprise CAL suite purchased through Open Value and Open Value Subscription volume licensing agreements.

Additional Terms: See (i) Appendix 1 – Service Level Commitment for Virus Detection and Blocking, Spam Effectiveness, or False Positive and (ii) Appendix 2 – Service Level Commitment for Uptime and Email Delivery.

[Table of Contents / Definitions](#)

Office 365 Business

Downtime: Any period of time when Office applications are put into reduced functionality mode due to an issue with Office 365 activation.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

[Table of Contents](#)

→

[Introduction](#)

→

[General Terms](#)

→

[Service Specific Terms](#)

→

[Appendices](#)

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Office 365 Customer Lockbox

Downtime: Any period of time when Customer Lockbox is put into reduced functionality mode due to an issue with Office 365.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Office 365 ProPlus

Downtime: Any period of time when Office applications are put into reduced functionality mode due to an issue with Office 365 activation.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Office Online

Downtime: Any period of time when users are unable to use the Web Applications to view and edit any Office document stored on a SharePoint Online site for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Office 365 Video

Downtime: Any period of time when users are unable to upload, view or edit videos in the video portal when they have appropriate permissions and valid content.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User Minutes - Downtime}{User Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Level Commitment:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

OneDrive for Business

Downtime: Any period of time when users are unable to view or edit files stored on their personal OneDrive for Business storage.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User Minutes - Downtime}{User Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Project Online

Downtime: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection with Project Web App for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

[Table of Contents](#)

→

[Introduction](#)

→

[General Terms](#)

→

[Service Specific Terms](#)

→

[Appendices](#)

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

SharePoint Online

Downtime: Any period of time when users are unable to read or write any portion of a SharePoint Online site collection for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Skype for Business Online

Downtime: Any period of time when end users are unable to see presence status, conduct instant messaging conversations, or initiate online meetings.¹

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes - Downtime}{User\ Minutes} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

¹Online meeting functionality applicable only to Skype for Business Online Plan 2 Service.

[Table of Contents](#) / [Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Skype for Business Online – PSTN Calling and PSTN Conferencing

Downtime: Any period of time when end users are unable to initiate a PSTN call or unable to dial into a PSTN conference.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

Where Downtime is measured in user-minutes; that is, for each month Downtime is the sum of the length (in minutes) of each incident that occurs during that month multiplied by the number of users impacted by that incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Skype for Business Online – Voice Quality

Additional Definitions:

“Eligible Call” is a Skype for Business placed call (within a subscription) that meets both conditions below:

- The call was placed from a Skype for Business Certified IP Desk phones on wired Ethernet
- Packet Loss, Jitter and Latency issues on the call were due to networks managed by Microsoft.

“Total Calls” is the total number of Eligible Calls

“Poor Quality Calls” is the total number of Eligible Calls that are classified as poor because of Packet Loss, Jitter and Latency issues in the networks managed by Microsoft. (For details on the measurements and thresholds refer <http://aka.ms/callquality>)

Monthly Good Call Rate: The Monthly Good Call Rate is calculated using the following formula:

$$\frac{\text{Total Calls} - \text{Poor Quality Calls}}{\text{Total Calls}} \times 100$$

Service Credit:

Monthly Good Call Rate	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Yammer Enterprise

Downtime: Any period of time greater than ten minutes when more than five percent of end users are unable to post or read messages on any portion of the Yammer network for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Enterprise Mobility Services

Azure Active Directory Basic

Downtime: Any period of time when users are not able to log in to the service, log in to the Access Panel, access applications on the Access Panel and reset passwords; or any period of time IT administrators are not able to create, read, write and delete entries in the directory and/or provision/de-provision users to applications in the directory.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)

Azure Active Directory Premium

Downtime: Any period of time when users are not able to log in to the service, log in to the Access Panel, access applications on the Access Panel and reset passwords; or any period of time IT administrators are not able to create, read, write and delete entries in the directory and/or provision/de-provision users to applications in the directory.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents](#) / [Definitions](#)
[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Azure Rights Management

Downtime: Any period of time when end users cannot create or consume IRM documents and email.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Microsoft Intune

Downtime: Any period of time when the Customer's IT administrator or users authorized by Customer are unable to log on with proper credentials. Scheduled Downtime will not exceed 10 hours per calendar year.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{User Minutes} - \text{Downtime}}{\text{User Minutes}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This Service Level does not apply to any: (i) On-premises software licensed as part of the Service subscription, or (ii) Internet-based services (excluding Microsoft Intune Service) that provide updates to any on-premise software licensed as part of the Service subscription.

[Table of Contents / Definitions](#)

Microsoft Azure Services

API Management Services

Additional Definitions:

"Deployment Minutes" is the total number of minutes that a given API Management instance has been deployed in Microsoft Azure during a billing month.

"Maximum Available Minutes" is the sum of all Deployment Minutes across all API Management instances deployed by you in a given Microsoft Azure subscription during a billing month.

"Proxy" is the component of the API Management Service responsible for receiving API requests and forwarding them to the configured dependent API.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Downtime: The total accumulated Deployment Minutes, across all API Management instances deployed by you in a given Microsoft Azure subscription, during which the API Management Service is unavailable. A minute is considered unavailable for a given API Management instance if all continuous attempts to perform operations through the Proxy throughout the minute result in either an Error Code or do not return a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit for Standard Tier:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Credit for Premium Tier deployments scaled across two or more regions:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents / Definitions](#)

App Service

Additional Definitions:

“App” is a Web App or Mobile App deployed by Customer within the App Service, excluding web apps in the Free and Shared tiers.

“Deployment Minutes” is the total number of minutes that a given App has been set to running in Microsoft Azure during a billing month. Deployment Minutes is measured from when the App was created or the Customer initiated an action that would result in running the App to the time the Customer initiated an action that would result in stopping or deleting the Web App.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Apps deployed by Customer in a given Microsoft Azure subscription during a billing month

Downtime: is the total accumulated Deployment Minutes, across all Apps deployed by Customer in a given Microsoft Azure subscription, during which the App is unavailable. A minute is considered unavailable for a given App when there is no connectivity between the App and Microsoft’s Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Additional Terms: Service Credits are applicable only to fees attributable to your use of Web Apps or Mobile Apps and not to fees attributable to other types of apps available through the App Service, which are not covered by this SLA.

[Table of Contents / Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Application Gateway

Additional Definitions:

“Application Gateway Cloud Service” refers to a collection of one or more Application Gateway instances configured to perform HTTP load balancing services.

“Maximum Available Minutes” is the total accumulated minutes during a billing month during which an Application Gateway Cloud Service comprising two or more medium or larger Application Gateway instances has been deployed in a Microsoft Azure subscription.

Downtime: is the total accumulated Maximum Available Minutes during a billing month for a given Application Gateway Cloud Service during which the Application Gateway Cloud Service is unavailable. A given minute is considered unavailable if all attempts to connect to the Application Gateway Cloud Service throughout the minute are unsuccessful.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Automation Service

Additional Definitions:

“Delayed Jobs” is the total number of Jobs, for a given Microsoft Azure subscription, that fail to start within thirty (30) minutes of their Planned Start Times.

“Job” means the execution of a Runbook.

“Planned Start Time” is a time at which a Job is scheduled to begin executing.

“Runbook” means a set of actions specified by you to execute within Microsoft Azure.

“Total Jobs” is the total number of Jobs scheduled for execution during a given billing month, for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Jobs} - \text{Delayed Jobs}}{\text{Total Jobs}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Backup Service

Additional Definitions:

“Backup” or **“Back Up”** is the process of copying computer data from a registered server to a Backup Vault.

“Backup Agent” refers to the software installed on a registered server that enables the registered server to Back Up or Restore one or more Protected Items.

“Backup Vault” refers to a container in which you may register one or more Protected Items for Backup.

“Deployment Minutes” is the total number of minutes during which a Protected Item has been scheduled for Backup to a Backup Vault.

“Failure” means that either the Backup Agent or the Service fails to fully complete a properly configured Backup or Recovery operation due to unavailability of the Backup Service.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Protected Items for a given Microsoft Azure subscription during a billing month.

“Protected Item” refers to a collection of data, such as a volume, database, or virtual machine that has been scheduled for Backup to the Backup Service such that it is enumerated as a Protected Item in the Protected Items tab in the Recovery Services section of the Management Portal.

“Recovery” or **“Restore”** is the process of restoring computer data from a Backup Vault to a registered server.

Downtime: The total accumulated Deployment Minutes across all Protected Items scheduled for Backup by you in a given Microsoft Azure subscription during which the Backup Service is unavailable for the Protected Item. The Backup Service is considered unavailable for a given Protected Item from the first Failure to Back Up or Restore the Protected Item until the initiation of a successful Backup or Recovery of a Protected Item, provided that retries are continually attempted no less frequently than once every thirty minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Batch Service

Additional Definitions:

“Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“Error Rate” is the total number of Failed Requests divided by Total Requests during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

“Excluded Requests” are requests within Total Requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

“Failed Requests” is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or fail to return a Success Code within 5 seconds.

“Total Requests” is the total number of authenticated REST API requests, other than Excluded Requests, to perform operations against Batch accounts attempted within a one-hour interval within a given Azure subscription during a billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

BizTalk Services

Additional Definitions:

“BizTalk Service Environment” refers to a deployment of the BizTalk Services created by you, as represented in the Management Portal, to which you may send runtime message requests.

“Deployment Minutes” is the total number of minutes that a given BizTalk Service Environment has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all BizTalk Service Environments deployed by you in a given Microsoft Azure subscription during a billing month.

“Monitoring Storage Account” refers to the Azure Storage account used by the BizTalk Services to store monitoring information related to the execution of the BizTalk Services.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Downtime: The total accumulated Deployment Minutes, across all BizTalk Service Environments deployed by you in a given Microsoft Azure subscription, during which the BizTalk Service Environment is unavailable. A minute is considered unavailable for a given BizTalk Service Environment when there is no connectivity between your BizTalk Service Environment and Microsoft's Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Basic, Standard, and Premium tiers of the BizTalk Services. The Developer tier of the Microsoft Azure BizTalk Services is not covered by this SLA.

Additional Terms: When submitting a claim, you must ensure that complete monitoring data is maintained within the Monitoring Storage Account and is made available to Microsoft.

[Table of Contents / Definitions](#)

Cache Services

Additional Definitions:

"Cache" refers to a deployment of the Cache Service created by you, such that its Cache Endpoints are enumerated in the Cache tab in the Management Portal.

"Cache Endpoints" refers to endpoints through which a Cache may be accessed.

"Deployment Minutes" is the total number of minutes that a given Cache has been deployed in Microsoft Azure during a billing month.

"Maximum Available Minutes" is the sum of all Deployment Minutes across all Caches deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes, across all Caches deployed by you in a given Microsoft Azure subscription, during which the Cache is unavailable. A minute is considered unavailable for a given Cache when there is no connectivity throughout the minute between one or more Cache Endpoints associated with the Cache and Microsoft's Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Cache Service, which includes the Azure Managed Cache Service or the Standard tier of the Azure Redis Cache Service. The Basic tier of the Azure Redis Cache Service is not covered by this SLA.

[Table of Contents / Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

CDN Service

Downtime: To assess Downtime, Microsoft will review data from any commercially reasonable independent measurement system used by you.

You must select a set of agents from the measurement system's list of standard agents that are generally available and represent at least five geographically diverse locations in major worldwide metropolitan areas (excluding PR of China).

Measurement System tests (frequency of at least one test per hour per agent) will be configured to perform one HTTP GET operation according to the model below:

1. A test file will be placed on your origin (e.g., Azure Storage account).
2. The GET operation will retrieve the file through the CDN Service, by requesting the object from the appropriate Microsoft Azure domain name hostname.
3. The test file will meet the following criteria:
 - i. The test object will allow caching by including explicit "Cache-control: public" headers, or lack of "Cache-Control: private" header.
 - ii. The test object will be a file at least 50KB in size and no larger than 1MB.
 - iii. Raw data will be trimmed to eliminate any measurements that came from an agent experiencing technical problems during the measurement period.

Monthly Uptime Percentage: The percentage of HTTP transactions in which the CDN responds to client requests and delivers the requested content without error. Monthly Uptime Percentage of the CDN Service is calculated as the number of times the object was delivered successfully divided by the total number of requests (after removing erroneous data).

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99.5%	25%

[Table of Contents](#) / [Definitions](#)

Cloud Services

Additional Definitions:

"Cloud Services" refers to a set of compute resources utilized for Web and Worker Roles.

"Maximum Available Minutes" is the total accumulated minutes during a billing month for all Internet facing roles that have two or more instances deployed in different Update Domains. Maximum Available Minutes is measured from when the Tenant has been deployed and its associated roles have been started resultant from action initiated by you to the time you have initiated an action that would result in stopping or deleting the Tenant.

"Tenant" represents one or more roles each consisting of one or more role instances that are deployed in a single package.

"Update Domain" refers to a set of Microsoft Azure instances to which platform updates are concurrently applied.

"Web Role" is a Cloud Services component run in the Azure execution environment that is customized for web application programming as supported by IIS and ASP.NET.

"Worker Role" is a Cloud Services component run in the Azure execution environment that is useful for generalized development, and may perform background processing for a Web Role.

Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Data Factory – Activity Runs

Additional Definitions:

“**Activity Run**” means the execution or attempted execution of an activity

“**Delayed Activity Runs**” is the total number of attempted Activity Runs in which an activity fails to begin executing within four (4) minutes after the time at which it is scheduled for execution and all dependencies that are prerequisite to execution have been satisfied.

“**Total Activity Runs**” is the total number of Activity Runs attempted during in a billing month for a given Microsoft Azure Subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Activity Runs} - \text{Delayed Activity Runs}}{\text{Total Activity Runs}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Data Factory – API Calls

Additional Definitions:

“**Excluded Requests**” is the set of requests within Total Requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

“**Failed Requests**” is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or otherwise fail to return a Success Code within two minutes.

“**Resources**” means pipelines, data sets, and linked services created within a Data Factory.

“**Total Requests**” is the set of all requests, other than Excluded Requests, to perform operations against Resources within active pipelines during a billing month for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Requests} - \text{Failed Requests}}{\text{Total Requests}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

DocumentDB

Additional Definitions:

“**Average Error Rate**” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“**Database Account**” is a DocumentDB account containing one or more databases.

“**Error Rate**” is the total number of Failed Requests divided by Total Requests, across all Resources in a given Azure subscription, during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

“**Excluded Requests**” are requests within Total Requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

“**Failed Requests**” is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or fail to return a Success Code within 5 seconds.

“**Resource**” is a set of URI addressable entities associated with a Database Account.

“**Total Request**” is the set of all requests, other than Excluded Requests, to perform operations issued against Resources attempted within a one-hour interval within a given Azure subscription during a billing month.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents / Definitions](#)

ExpressRoute

Additional Definitions:

“Dedicated Circuit” means a logical representation of connectivity offered through the ExpressRoute Service between your premises and Microsoft Azure through an exchange provider or a network service provider, where such connectivity does not traverse the public Internet.

“Maximum Available Minutes” is the total number of minutes that a given Dedicated Circuit is linked to one or more Virtual Networks in Microsoft Azure during a billing month in a given Microsoft Azure subscription.

“Virtual Network” refers to a virtual private network that includes a collection of user-defined IP addresses and subnets that form a network boundary within Microsoft Azure.

“VPN Gateway” refers to a gateway that facilitates cross-premises connectivity between a Virtual Network and a customer on-premises network.

Downtime: The total accumulated minutes during a billing month for a given Microsoft Azure subscription during which the Dedicated Circuit is unavailable. A minute is considered unavailable for a given Dedicated Circuit if all attempts by you within the minute to establish IP-level connectivity to the VPN Gateway associated with the Virtual Network fail for longer than thirty seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Additional Terms: Monthly Uptime Percentage and Service Credits are calculated for each Dedicated Circuit used by you.

[Table of Contents / Definitions](#)

HDInsight

Additional Definitions:

“Cluster Internet Gateway” means a set of virtual machines within an HDInsight Cluster that proxy all connectivity requests to the Cluster.

“Deployment Minutes” is the total number of minutes that a given HDInsight Cluster has been deployed in Microsoft Azure.

“HDInsight Cluster” or **“Cluster”** means a collection of virtual machines running a single instance of the HDInsight Service.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Clusters deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes when the HDInsight Service is unavailable. A minute is considered unavailable for a given Cluster if all continual attempts within the minute to establish a connection to the Cluster Internet Gateway fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

[Table of Contents](#)

→

[Introduction](#)

→

[General Terms](#)

→

[Service Specific Terms](#)

→

[Appendices](#)

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

HockeyApp

Additional Definitions:

"HockeyApp Dashboard" means the web interface provided to developers to view and manage applications using the HockeyApp Service.

"Maximum Available Minutes" is the total number of minutes in a billing month.

Downtime: is the total accumulated minutes in a billing month during which the HockeyApp Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the HockeyApp Dashboard or to the HockeyApp API throughout the minute either result in an Error Code or do not return a response within one minute. For purposes of the HockeyApp API, HTTP response codes 408, 429, 500, 503, and 511 are not considered Error Codes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Key Vault

Additional Definitions:

"Deployment Minutes" is the total number of minutes that a given key vault has been deployed in Microsoft Azure during a billing month.

"Excluded Transactions" are transactions for creating, updating, or deleting key vaults, keys, or secrets.

"Maximum Available Minutes" is the sum of all Deployment Minutes across all Key Vaults deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment Minutes, across all key vaults deployed by Customer in a given Microsoft Azure subscription, during which the key vault is unavailable. A minute is considered unavailable for a given key vault if all continuous attempts to perform transactions, other than Excluded Transactions, on the key vault throughout the minute either return an Error Code or do not result in a Success Code within 5 seconds from Microsoft's receipt of the request.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)
[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Machine Learning – Batch Execution Service (BES) and Management APIs Service

Additional Definitions:

“**Failed Transactions**” is the set of all requests within Total Transaction Attempts that return an Error Code.

“**Total Transaction Attempts**” is the total number of authenticated REST BES and Management API requests by you during a billing month for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: Service Levels and Service Credits are applicable to your use of the Machine Learning BES and Management API Service. The Free Machine Learning tier is not covered by this SLA.

[Table of Contents / Definitions](#)

Machine Learning – Request Response Service (RRS)

Additional Definitions:

“**Failed Transactions**” is the set of all requests within Total Transaction Attempts that return an Error Code.

“**Total Transaction Attempts**” is the total number of authenticated REST RRS and Management API requests by you during a billing month for a given Microsoft Azure subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Service Level Exceptions: Service Levels and Service Credits are applicable to your use of the Machine Learning RRS and Management API Service. The Free Machine Learning tier is not covered by this SLA.

[Table of Contents / Definitions](#)

Media Services – Content Protection Service

Additional Definitions:

“**Failed Transactions**” are all Valid Key Requests included in Total Transaction Attempts that result in an Error Code or otherwise do not return a Success Code within 30 seconds after receipt by the Content Protection Service.

“**Total Transaction Attempts**” are all Valid Key Requests made by you during a billing month for a given Azure subscription.

“**Valid Key Requests**” are all requests made to the Content Protection Service for existing content keys in a Customer's Media Service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Media Services – Encoding Service

Additional Definitions:

“Encoding” means the processing of media files per subscription as configured in the Media Services Tasks.

“Failed Transactions” is the set of all requests within Total Transaction Attempts that do not return a Success Code within 30 seconds from Microsoft’s receipt of the request.

“Media Service” means an Azure Media Services account, created in the Management Portal, associated with your Microsoft Azure subscription. Each Microsoft Azure subscription may have more than one associated Media Service.

“Media Services Task” means an individual operation of media processing work as configured by you. Media processing operations involve encoding and converting media files.

“Total Transaction Attempts” is the total number of authenticated REST API requests with respect to a Media Service made by you during a billing month for a subscription. Total Transaction Attempts does not include REST API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Media Services – Indexer Service

Additional Definitions:

“Encoding Reserved Unit” means encoding reserved units purchased by the customer in an Azure Media Services account

“Failed Transactions” is the set of Indexer Tasks within Total Transaction Attempts that either, a) do not complete within a time period that is 3 times the duration of the input file, or b) do not start processing within 5 minutes of the time that an Encoding Reserved Unit becomes available for use by the Indexer Task.

“Indexer Task” means a Media Services Task that is configured to index an MP3 input file with a minimum five-minute duration.

“Total Transaction Attempts” is the total number of Indexer Tasks attempted to be executed using an available Encoding Reserved Unit by Customer during a billing month for a subscription.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)
[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Media Services – Live Channels

Additional Definitions:

“**Channel**” means an end point within a Media Service that is configured to receive media data.

“**Deployment Minutes**” is the total number of minutes that a given Channel has been purchased and allocated to a Media Service and is in a running state during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Channels purchased and allocated to a Media Service during a billing month.

“**Media Service**” means an Azure Media Services account, created in the Management Portal, associated with your Microsoft Azure subscription. Each Microsoft Azure subscription may have more than one associated Media Service.

Downtime: The total accumulated Deployment Minutes when the Live Channels Service is unavailable. A minute is considered unavailable for a given Channel if the Channel has no External Connectivity during the minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Media Services – Streaming Service

Additional Definitions:

“**Deployment Minutes**” is the total number of minutes that a given Streaming Unit has been purchased and allocated to a Media Service during a billing month.

“**Maximum Available Minutes**” is the sum of all Deployment Minutes across all Streaming Units purchased and allocated to a Media Service during a billing month.

“**Media Service**” means an Azure Media Services account, created in the Management Portal, associated with your Microsoft Azure subscription. Each Microsoft Azure subscription may have more than one associated Media Service.

“**Media Service Request**” means a request issued to your Media Service.

“**Streaming Unit**” means a unit of reserved egress capacity purchased by you for a Media Service.

“**Valid Media Services Requests**” are all qualifying Media Service Requests for existing media content in a customer’s Azure Storage account associated with its Media Service when at least one Streaming Unit has been purchased and allocated to that Media Service. Valid Media Services Requests do not include Media Service Requests for which total throughput exceeds 80% of the Allocated Bandwidth.

Downtime: The total accumulated Deployment Minutes when the Streaming Service is unavailable. A minute is considered unavailable for a given Streaming Unit if all continuous Valid Media Service Requests made to the Streaming Unit throughout the minute result in an Error Code.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)
[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Mobile Engagement

Additional Definitions:

"Average Error Rate" for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

"Error Rate" is the total number of Failed Requests divided by Total Requests during a given one-hour interval. If the Total Requests in a given one-hour interval is zero, the Error Rate for that interval is 0%.

"Excluded Requests" is the set of REST API requests that result in an HTTP 4xx status code, other than an HTTP 408 status code.

"Failed Requests" is the set of all requests within Total Requests that either return an Error Code or an HTTP 408 status code or fail to return a Success Code within 30 seconds.

"Mobile Engagement Application" is an Azure Mobile Engagement service instance.

"Total Requests" is the total number of authenticated REST API requests, other than Excluded Requests, made to Mobile Engagement Applications within a given Azure subscription during a billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

The Free Mobile Engagement tier is not covered by this SLA.

[Table of Contents / Definitions](#)

Mobile Services

Additional Definitions:

"Failed Transactions" include any API calls included in Total Transaction Attempts that result in either an Error Code or do not return a Success Code.

"Total Transaction Attempts" are the total accumulated API calls made to the Azure Mobile Services during a billing month for a given Microsoft Azure subscription for which the Azure Mobile Services are running.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Standard and Premium Mobile Services tiers. The Free Mobile Services tier is not covered by this SLA.

[Table of Contents / Definitions](#)

Multi-Factor Authentication Service

Additional Definitions:

"Deployment Minutes" is the total number of minutes that a given Multi-Factor Authentication provider has been deployed in Microsoft Azure during a billing month.

"Maximum Available Minutes" is the sum of all Deployment Minutes across all Multi-Factor Authentication providers deployed by you in a given Microsoft Azure subscription during a billing month.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Downtime: The total accumulated Deployment Minutes, across all Multi-Factor Authentication providers deployed by you in a given Microsoft Azure subscription, during which the Multi-Factor Authentication Service is unable to receive or process authentication requests for the Multi-Factor Authentication provider.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Operational Insights

Additional Definitions:

“Batch” means a group of Log Data entries that are either uploaded to the Operational Insights Service or read from storage by the Operational Insights Service within a given period of time. Batches queued for indexing are displayed in the usage section of the Management Portal.

“Log Data” refers to information regarding a supported event, such as IIS and Windows events, that is logged by a computer and for which the Operational Insights Service has been configured to be processed by the Service index.

“Delayed Batches” is the total number of Batches within Total Queued Batches that fail to complete indexing within six hours of the Batch being queued.

“Total Queued Batches” is the total number of Batches queued for indexing by the Operational Insights Service during a given billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total Queued Batches} - \text{Delayed Batches}}{\text{Total Queued Batches}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

RemoteApp

Additional Definitions:

“Application” means a software application that is configured for streaming to a device using the RemoteApp Service.

“Maximum Available Minutes” is the sum of all User Application Minutes across all Users granted access to one or more Applications in a given Azure subscription during a billing month.

“User” means a specific user account that is able to stream an Application using the RemoteApp Service, as enumerated in the Management Portal.

“User Application Minutes” is the total number of minutes in a billing month during which you have granted a User access to an Application.

Downtime: The total accumulated User Minutes during which the RemoteApp Service is unavailable. A minute is considered unavailable for a given User when the User is unable to establish connectivity to an Application.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the RemoteApp Service. The RemoteApp free trial is not covered by this SLA.

[Table of Contents / Definitions](#)

Scheduler

Additional Definitions:

“Maximum Available Minutes” is the total number of minutes in a billing month.

“Planned Execution Time” is a time at which a Scheduled Job is scheduled to begin executing.

“Scheduled Job” means an action specified by you to execute within Microsoft Azure according to a specified schedule.

Downtime: The total accumulated minutes in a billing month during which one or more of your Scheduled Jobs is in a state of delayed execution. A given Scheduled Job is in a state of delayed execution if it has not begun executing after a Planned Execution Time, provided that such delayed execution time shall not be considered Downtime if the Scheduled Job begins executing within thirty (30) minutes after a Planned Execution Time.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Search

Additional Definitions:

“Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“Error Rate” is the total number of Failed Requests divided by Total Requests, across all Search Service Instances in a given Azure subscription, during a given one-hour interval. If the Total Requests in a one-hour interval is zero, the Error Rate for that interval is 0%.

“Excluded Requests” are all requests that are throttled due to exhaustion of resources allocated for a Search Service Instance, as indicated by an HTTP 503 status code and a response header indicating the request was throttled.

“Failed Requests” is the set of all requests within Total Requests that fail to return either a Success Code or HTTP 4xx response.

“Replica” is a copy of a search index within a Search Service Instance.

“Search Service Instance” is an Azure Search service instance containing one or more search indexes.

“Total Requests” is the set of (i) all requests to update a Search Service Instance having three or more Replicas, plus (ii) all requests to query a Search Service Instance having two or more Replicas, other than Excluded Requests, within a one-hour interval within a given Azure subscription during a billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Free Search tier is not covered by this SLA.

[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Service-Bus Service – Event Hubs

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Event Hub has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Event Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Event Hubs tiers during a billing month.

“Message” refers to any user-defined content sent or received through Service Bus Relays, Queues, Topics, or Notification Hubs, using any protocol supported by Service Bus.

Downtime: The total accumulated Deployment Minutes, across all Event Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Event Hubs tiers, during which the Event Hub is unavailable. A minute is considered unavailable for a given Event Hub if all continuous attempts to send or receive Messages or perform other operations on the Event Hub throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Basic and Standard Event Hubs tiers. The Free Event Hubs tier is not covered by this SLA.

Service-Bus Service – Notification Hubs

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Notification Hub has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Notification Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Notification Hubs tiers during a billing month.

Downtime: The total accumulated Deployment Minutes, across all Notification Hubs deployed by you in a given Microsoft Azure subscription under the Basic or Standard Notification Hubs tiers, during which the Notification Hub is unavailable. A minute is considered unavailable for a given Notification Hub if all continuous attempts to send notifications or perform registration management operations with respect to the Notification Hub throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: The Service Levels and Service Credits are applicable to your use of the Basic and Standard Notification Hubs tiers. The Free Notification Hubs tier is not covered by this SLA.

Service-Bus Service – Queues and Topics

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Queue or Topic has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Queues and Topics deployed by you in a given Microsoft Azure subscription during a billing month.

“Message” refers to any user-defined content sent or received through Service Bus Relays, Queues, Topics, or Notification Hubs, using any protocol supported by Service Bus.

Downtime: The total accumulated Deployment Minutes, across all Queues and Topics deployed by you in a given Microsoft Azure subscription, during which the Queue or Topic is unavailable. A minute is considered unavailable for a given Queue or Topic if all continuous attempts to send or receive Messages or perform other operations on the Queue or Topic throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Service-Bus Service – Relays

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Relay has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Relays deployed by you in a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes, across all Relays deployed by you in a given Microsoft Azure subscription, during which the Relay is unavailable. A minute is considered unavailable for a given Relay if all continuous attempts to establish a connection to the Relay throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Site Recovery Service – On-Premises-to-Azure

Additional Definitions:

“Failover” is the process of transferring control, either simulated or actual, of a Protected Instance from a primary site to a secondary site.

“On-Premises-to-Azure Failover” is the Failover of a Protected Instance from a non-Azure primary site to an Azure secondary site. You may designate a particular Azure datacenter as a secondary site, provided that if Failover to the designated datacenter is not possible, Microsoft may replicate to a different datacenter in the same region.

“Protected Instance” refers to a virtual or physical machine configured for replication by the Site Recovery Service from a primary site to a secondary site. Protected Instances are enumerated in the Protected Items tab in the Recovery Services section of the Management Portal.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

“Recovery Time Objective (RTO)” means the period of time beginning when you initiate a Failover of a Protected Instance experiencing either a planned or unplanned outage for On-Premises-to-Azure replication to the time when the Protected Instance is running as a virtual machine in Microsoft Azure, excluding any time associated with manual action or the execution of your scripts.

Monthly Recovery Time Objective: The Monthly Recovery Time Objective for a specific Protected Instance configured for On-Premises-to-Azure replication in a given billing month is four hours for an unencrypted Protected Instance and six hours for an encrypted Protected Instance. One hour will be added to the monthly Recovery Time Objective for each additional 25GB over the initial 100GB Protected Instance size.

Service Credit (Assuming Protected Instance of 100GB, or less):

Protected Instance	Monthly Recovery Time Objective	Service Credit
Unencrypted	> 4 hours	100%
Encrypted	> 6 hours	100%

Additional Terms: Monthly Recovery Time Objective and Service Credits are calculated for each Protected Instance used by you.

[Table of Contents](#) / [Definitions](#)

Site Recovery Service – On-Premises-to-On-Premises

Additional Definitions:

“Failover” is the process of transferring control, either simulated or actual, of a Protected Instance from a primary site to a secondary site.

“Failover Minutes” is the total number of minutes in a billing month during which a Failover of a Protected Instance configured for On-Premises-to-On-Premises replication has been attempted but not completed.

“Maximum Available Minutes” is the total number of minutes that a given Protected Instance has been configured for On-Premises-to-On-Premises replication by the Site Recovery Service during a billing month.

“On-Premises-to-On-Premises Failover” is the Failover of a Protected Instance from a non-Azure primary site to a non-Azure secondary site.

“Protected Instance” refers to a virtual or physical machine configured for replication by the Site Recovery Service from a primary site to a secondary site. Protected Instances are enumerated in the Protected Items tab in the Recovery Services section of the Management Portal.

Downtime: The total accumulated Failover Minutes in which the Failover of a Protected Instance is unsuccessful due to unavailability of the Site Recovery Service, provided that retries are continually attempted no less frequently than once every thirty minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Additional Terms: Monthly Recovery Time Objective and Service Credits are calculated for each Protected Instance used by you.

[Table of Contents](#) / [Definitions](#)

SQL Database Service (Basic, Standard and Premium Tiers)

Additional Definitions:

“Database” means any Basic, Standard, or Premium Microsoft Azure SQL Database.

“Deployment Minutes” is the total number of minutes that a given Basic, Standard, or Premium Database has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Basic, Standard, and Premium Databases for a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes across all Basic, Standard, and Premium Databases deployed by you in a given Microsoft Azure subscription during which the Database is unavailable. A minute is considered unavailable for a given Database if all continuous attempts by you to establish a connection to the Database within the minute fail.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

SQL Database Service (Web and Business Tiers)

Additional Definitions:

“Database” means any Web or Business Microsoft Azure SQL Database.

“Deployment Minutes” is the total number of minutes that a given Web or Business Database has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Web and Business Databases for a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated Deployment Minutes across all Web and Business Databases deployed by you in a given Microsoft Azure subscription during which the Database is unavailable. A minute is considered unavailable for a given Database if all continuous attempts by you to establish a connection to the Database within the minute fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents](#) / [Definitions](#)

Storage Service

Additional Definitions:

“Average Error Rate” for a billing month is the sum of Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

“Excluded Transactions” are storage transactions that do not count toward either Total Storage Transactions or Failed Storage Transactions. Excluded Transactions include pre-authentication failures; authentication failures; attempted transactions for storage accounts over their prescribed quotas; creation or deletion of containers, tables, or queues; clearing of queues; and copying blobs between storage accounts.

“Error Rate” is the total number of Failed Storage Transactions divided by the Total Storage Transactions during a set time interval (currently set at one hour). If the Total Storage Transactions in a given one-hour interval is zero, the error rate for that interval is 0%.

“Failed Storage Transactions” is the set of all storage transactions within Total Storage Transactions that are not completed within the Maximum Processing Time associated with their respective transaction type, as specified in the table below. Maximum Processing Time includes only the time spent processing a transaction request within the Storage Service and does not include any time spent transferring the request to or from the Storage Service.

Request Types	Maximum Processing Time
PutBlob and GetBlob (includes blocks and pages) Get Valid Page Blob Ranges	Two (2) seconds multiplied by the number of MBs transferred in the course of processing the request
Copy Blob	Ninety (90) seconds (where the source and destination blobs are within the same storage account)
PutBlockList	Sixty (60) seconds

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Request Types	Maximum Processing Time
GetBlockList	
Table Query List Operations	Ten (10) seconds (to complete processing or return a continuation)
Batch Table Operations	Thirty (30) seconds
All Single Entity Table Operations All other Blob and Message Operations	Two (2) seconds

These figures represent maximum processing times. Actual and average times are expected to be much lower.

Failed Storage Transactions do not include:

1. Transaction requests that are throttled by the Storage Service due to a failure to obey appropriate back-off principles.
2. Transaction requests having timeouts set lower than the respective Maximum Processing Times specified above.
3. Read transactions requests to RA-GRS Accounts for which you did not attempt to execute the request against Secondary Region associated with the storage account if the request to the Primary Region was not successful.
4. Read transaction requests to RA-GRS Accounts that fail due to Geo-Replication Lag.

“Geo Replication Lag” for GRS and RA-GRS Accounts is the time it takes for data stored in the Primary Region of the storage account to replicate to the Secondary Region of the storage account. Because GRS and RA-GRS Accounts are replicated asynchronously to the Secondary Region, data written to the Primary Region of the storage account will not be immediately available in the Secondary Region. You can query the Geo Replication Lag for a storage account, but Microsoft does not provide any guarantees as to the length of any Geo Replication Lag under this SLA.

“Geographically Redundant Storage (GRS) Account” is a storage account for which data is replicated synchronously within a Primary Region and then replicated asynchronously to a Secondary Region. You cannot directly read data from or write data to the Secondary Region associated with GRS Accounts.

“Locally Redundant Storage (LRS) Account” is a storage account for which data is replicated synchronously only within a Primary Region.

“Primary Region” is a geographical region in which data within a storage account is located, as selected by you when creating the storage account. You may execute write requests only against data stored within the Primary Region associated with storage accounts.

“Read Access Geographically Redundant Storage (RA-GRS) Account” is a storage account for which data is replicated synchronously within a Primary Region and then replicated asynchronously to a Secondary Region. You can directly read data from, but cannot write data to, the Secondary Region associated with RA-GRS Accounts.

“Secondary Region” is a geographical region in which data within a GRS or RA-GRS Account is replicated and stored, as assigned by Microsoft Azure based on the Primary Region associated with the storage account. You cannot specify the Secondary Region associated with storage accounts.

“Total Storage Transactions” is the set of all storage transactions, other than Excluded Transactions, attempted within a one-hour interval across all storage accounts in the Storage Service in a given subscription.

“Zone Redundant Storage (ZRS) Account” is a storage account for which data is replicated across multiple facilities. These facilities may be within the same geographical region or across two geographical regions.

Monthly Uptime Percentage: Monthly Uptime Percentage is calculated using the following formula:

$$100\% - \text{Average Error Rate}$$

Service Credit – LRS, ZRS, GRS and RA-GRS (write requests) Accounts:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Credit – RA-GRS (read requests) Accounts:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents / Definitions](#)

StorSimple Service

Additional Definitions:

“Backup” is the process of backing up data stored on a registered StorSimple device to one or more associated cloud storage accounts within Microsoft Azure.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

“Cloud Tiering” is the process of transferring data from a registered StorSimple device to one or more associated cloud storage accounts within Microsoft Azure.

“Deployment Minutes” is the total number of minutes during which a Managed Item has been configured for Backup or Cloud Tiering to a StorSimple storage account in Microsoft Azure.

“Failure” means the inability to fully complete a properly configured Backup, Tiering, or Restoring operation due to unavailability of the StorSimple Service.

“Managed Item” refers to a volume that has been configured to Backup to the cloud storage accounts using the StorSimple Service.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Managed Items for a given Microsoft Azure subscription during a billing month.

“Restoring” is the process of copying data to a registered StorSimple device from its associated cloud storage account(s).

Downtime: The total accumulated Deployment Minutes across all Managed Items configured for Backup or Cloud Tiering by you in a given Microsoft Azure subscription during which the StorSimple Service is unavailable for the Managed Item. The StorSimple Service is considered unavailable for a given Managed Item from the first Failure of a Backup, Cloud Tiering, or Restoring operation with respect to the Managed Item until the initiation of a successful Backup, Cloud Tiering, or Restoring operation of the Managed Item, provided that retries are continually attempted no less frequently than once every thirty minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes} - \text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Stream Analytics – API Calls

Additional Definitions:

“Total Transaction Attempts” is the total number of authenticated REST API requests to manage a streaming job within the Stream Analytics Service by Customer during a billing month for a given Microsoft Azure subscription.

“Failed Transactions” is the set of all requests within Total Transaction Attempts that return an Error Code or otherwise do not return a Success Code within five minutes from Microsoft’s receipt of the request.

“Monthly Uptime Percentage” for API calls within the Stream Analytics Service is represented by the following formula:

$$\text{Monthly Uptime \%} = \frac{\text{Total Transaction Attempts} - \text{Failed Transactions}}{\text{Total Transaction Attempts}}$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Stream Analytics – Jobs

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given job has been deployed within the Stream Analytics Service during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all jobs deployed by Customer in a given Microsoft Azure subscription during a billing month.

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Downtime is the total accumulated Deployment Minutes, across all jobs deployed by Customer in a given Microsoft Azure subscription, during which the job is unavailable. A minute is considered unavailable for a deployed job if the job is neither processing data nor available to process data throughout the minute.

Monthly Uptime Percentage for jobs within the Stream Analytics Service is represented by the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Traffic Manager Service

Additional Definitions:

“Deployment Minutes” is the total number of minutes that a given Traffic Manager Profile has been deployed in Microsoft Azure during a billing month.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all Traffic Manager Profiles deployed by you in a given Microsoft Azure subscription during a billing month.

“Traffic Manager Profile” or **“Profile”** refers to a deployment of the Traffic Manager Service created by you containing a domain name, endpoints, and other configuration settings, as represented in the Management Portal.

“Valid DNS Response” means a DNS response, received from at least one of the Traffic Manager Service name server clusters, to a DNS request for the domain name specified for a given Traffic Manager Profile.

Downtime: The total accumulated Deployment Minutes, across all Profiles deployed by you in a given Microsoft Azure subscription, during which the Profile is unavailable. A minute is considered unavailable for a given Profile if all continual DNS queries for the DNS name specified in the Profile that are made throughout the minute do not result in a Valid DNS Response within two seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Virtual Machines

Additional Definitions:

“Availability Set” refers to two or more Virtual Machines deployed across different Fault Domains to avoid a single point of failure.

“Fault Domain” is a collection of servers that share common resources such as power and network connectivity.

“Maximum Available Minutes” is the total accumulated minutes during a billing month for all Internet facing Virtual Machines that have two or more instances deployed in the same Availability Set. Maximum Available Minutes is measured from when at least two Virtual Machines in the same Availability Set have both been started resultant from action initiated by you to the time you have initiated an action that would result in stopping or deleting the Virtual Machines.

“Virtual Machine” refers to persistent instance types that can be deployed individually or as part of an Availability Set.

Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

[Table of Contents](#)

→

[Introduction](#)

→

[General Terms](#)

→

[Service Specific Terms](#)

→

[Appendices](#)

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

[Table of Contents / Definitions](#)

VPN Gateway

Additional Definitions:

“Maximum Available Minutes” is the total accumulated minutes during a billing month which a given VPN Gateway has been deployed in a Microsoft Azure subscription.

“Virtual Network” refers to a virtual private network that includes a collection of user-defined IP addresses and subnets that form a network boundary within Microsoft Azure.

“VPN Gateway” refers to a gateway that facilitates cross-premises connectivity between a Virtual Network and a customer on-premises network.

Downtime: Is the total accumulated VPN Gateway Maximum Available Minutes during which a VPN Gateway is unavailable. A minute is considered unavailable if all attempts to connect to the VPN Gateway within a thirty-second window within the minute are unsuccessful.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Visual Studio Online – Build Service

Additional Definitions:

“Build Service” is a feature that allows customers to build their applications in Visual Studio Online.

“Maximum Available Minutes” is the total number of minutes for which the paid Build Service has been enabled for a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated minutes for a given Microsoft Azure subscription during which the Build Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Build Service to perform operations initiated by you throughout the minute either result in an Error Code or do not return a response.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes-Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)
[Table of Contents](#)

[Introduction](#)

[General Terms](#)

[Service Specific Terms](#)

[Appendices](#)

Visual Studio Online – Load Testing Service

Additional Definitions:

“Load Testing Service” is a feature that allows customers to generate automated tasks to test the performance and scalability of applications.

“Maximum Available Minutes” is the total number of minutes for which the paid Load Testing Service has been enabled for a given Microsoft Azure subscription during a billing month.

Downtime: The total accumulated minutes for a given Microsoft Azure subscription during which the Load Testing Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Load Testing Service to perform operations initiated by you throughout the minute either result in an Error Code or do not return a response.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

Visual Studio Online – User Plans Service

Additional Definitions:

“Build Service” is a feature that allows customers to build their applications in Visual Studio Online.

“Deployment Minutes” is the total number of minutes for which a User Plan has been purchased during a billing month.

“Load Testing Service” is a feature that allows customers to generate automated tasks to test the performance and scalability of applications.

“Maximum Available Minutes” is the sum of all Deployment Minutes across all User Plans for a given Microsoft Azure subscription during a billing month.

“User Plan” refers to the set of features and capabilities selected for a user within a Visual Studio Online account in a Customer subscription. User Plan options and the features and capabilities per User Plan are described on the <http://www.visualstudio.com> website.

Downtime: The total accumulated Deployment Minutes, across all User Plans for a given Microsoft Azure subscription, during which the User Plan is unavailable. A minute is considered unavailable for a given User Plan if all continuous HTTP requests to perform operations, other than operations pertaining to the Build Service or the Load Testing Service, throughout the minute either result in an Error Code or do not return a response.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Maximum Available Minutes}-\text{Downtime}}{\text{Maximum Available Minutes}} \times 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

[Table of Contents / Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Other Online Services

Bing Maps Enterprise Platform

Downtime: Any period of time when the Service is not available as measured in Microsoft's data centers, provided that you access the Service using the methods of access, authentication and tracking methods documented in the Bing Maps Platform SDKs.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured as the total number of minutes during the month when the aspects of the Service set forth above are unavailable.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to Bing Maps Enterprise Platform purchased through Open Value and Open Value Subscription volume licensing agreements.

Service Credits will not apply if: (i) you fail to implement any Services updates within the time specified in the Bing Maps Platform API's Terms of Use; and (ii) you do not provide Microsoft with at least ninety (90) days' advance notice of any known significant usage volume increase, with significant usage volume increase defined as 50% or more of the previous month's usage.

[Table of Contents / Definitions](#)

Bing Maps Mobile Asset Management

Downtime: Any period of time when the Service is not available as measured in Microsoft's data centers, provided that you access the Service using the methods of access, authentication and tracking methods documented in the Bing Maps Platform SDKs.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured as the total number of minutes during the month when the aspects of the Service set forth above are unavailable.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Service Level Exceptions: This SLA does not apply to Bing Maps Enterprise Platform purchased through Open Value and Open Value Subscription volume licensing agreements.

Service Credits will not apply if: (i) you fail to implement any Services updates within the time specified in the Bing Maps Platform API's Terms of Use; and (ii) you do not provide Microsoft with at least ninety (90) days' advance notice of any known significant usage volume increase, with significant usage volume increase defined as 50% or more of the previous month's usage.

[Table of Contents / Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Power BI Pro

Downtime: Any period of time when users are unable to read or write any portion of Power BI data to which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured in user-minutes; that is, for each month, Downtime is the sum of the length (in minutes) of each Incident that occurs during that month multiplied by the number of users impacted by that Incident.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

Translator API

Downtime: Any period of time when users are not able to perform translations.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{\text{Total number of minutes in a month} - \text{Downtime}}{\text{Total number of minutes in a month}} \times 100$$

where Downtime is measured as the total number of minutes during the month when the aspects of the Service set forth above are unavailable.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

[Table of Contents / Definitions](#)

[Table of Contents](#)



[Introduction](#)



[General Terms](#)



[Service Specific Terms](#)



[Appendices](#)

Appendix A – Service Level Commitment for Virus Detection and Blocking, Spam Effectiveness, or False Positive

With respect to Exchange Online and EOP licensed as a standalone Service or via ECAL suite, or Exchange Enterprise CAL with Services, you may be eligible for Service Credits if we do not meet the Service Level described below for: (1) Virus Detection and Blocking, (2) Spam Effectiveness, or (3) False Positive. If any one of these individual Service Levels is not met, you may submit a claim for a Service Credit. If one Incident causes us to fail more than one SLA metric for Exchange Online or EOP, you may only make one Service Credit claim for that incident per Service.

1. Virus Detection and Blocking Service Level

- a. "Virus Detection and Blocking" is defined as the detection and blocking of Viruses by the filters to prevent infection. "Viruses" is broadly defined as known malware, which includes viruses, worms, and Trojan horses.
- b. A Virus is considered known when widely used commercial virus scanning engines can detect the virus and the detection capability is available throughout the EOP network.
- c. Must result from a non-purposeful infection.
- d. The Virus must have been scanned by the EOP virus filter.
- e. If EOP delivers an email that is infected with a known virus to you, EOP will notify you and work with you to identify and remove it. If this results in the prevention of an infection, you won't be eligible for a Service Credit under the Virus Detection and Blocking Service Level.
- f. The Virus Detection and Blocking Service Level shall not apply to:
 - i. Forms of email abuse not classified as malware, such as spam, phishing and other scams, adware, and forms of spyware, which due to its targeted nature or limited use is not known to the anti-virus community and thus not tracked by anti-virus products as a virus.
 - ii. Corrupt, defective, truncated, or inactive viruses contained in NDRs, notifications, or bounced emails.
- g. The Service Credit available for the Virus Detection and Blocking Service is: 25% Service Credit of Applicable Monthly Service Fee if an infection occurs in a calendar month, with a maximum of one claim allowed per calendar month.

2. Spam Effectiveness Service Level

- a. "Spam Effectiveness" is defined as the percentage of inbound spam detected by the filtering system, measured on a daily basis.
- b. Spam effectiveness estimates exclude false negatives to invalid mailboxes.
- c. The spam message must be processed by our service and not be corrupt, malformed, or truncated.
- d. The Spam Effectiveness Service Level does not apply to email containing a majority of non-English content.
- e. You acknowledge that classification of spam is subjective and accept that we will make a good faith estimation of the spam capture rate based on evidence timely supplied by you.
- f. The Service Credit available for the Spam Effectiveness Service is:

% of Calendar Month that Spam Effectiveness is below 99%	Service Credit
>25%	25%
> 50%	50%
100%	100%

3. False Positive Service Level

- a. "False Positive" is defined as the ratio of legitimate business email incorrectly identified as spam by the filtering system to all email processed by the service in a calendar month.
- b. Complete, original messages, including all headers, must be reported to the abuse team.
- c. Applies to email sent to valid mailboxes only.
- d. You acknowledge that classification of false positives is subjective and understand that we will make a good faith estimation of the false positive ratio based on evidence timely supplied by you.
- e. This False Positive Service Level shall not apply to:
 - i. bulk, personal, or pornographic email
 - ii. email containing a majority of non-English content
 - iii. email blocked by a policy rule, reputation filtering, or SMTP connection filtering
 - iv. email delivered to the junk folder
- f. The Service Credit available for the False Positive Service is:

False Positive Ratio in a Calendar Month	Service Credit
> 1:250,000	25%
> 1:10,000	50%
> 1:100	100%

Appendix B - Service Level Commitment for Uptime and Email Delivery

With respect to EOP licensed as a standalone Service, ECAL suite, or Exchange Enterprise CAL with Services, you may be eligible for Service Credits if we do not meet the Service Level described below for (1) Uptime and (2) Email Delivery.

1. **Monthly Uptime Percentage:**

If the Monthly Uptime Percentage for EOP falls below 99.999% for any given month, you may be eligible for the following Service Credit:

Monthly Uptime Percentage	Service Credit
<99.999%	25%
<99.0%	50%
<98.0%	100%

2. **Email Delivery Service Level:**

- a. "Email Delivery Time" is defined as the average of email delivery times, measured in minutes over a calendar month, where email delivery is defined as the elapsed time from when a business email enters the EOP network to when the first delivery attempt is made.
- b. Email Delivery Time is measured and recorded every 5 minutes, then sorted by elapsed time. The fastest 95% of measurements are used to create the average for the calendar month.
- c. We use simulated or test emails to measure delivery time.
- d. The Email Delivery Service Level applies only to legitimate business email (non-bulk email) delivered to valid email accounts.
- e. This Email Delivery Service Level does not apply to:
 1. Delivery of email to quarantine or archive
 2. Email in deferral queues
 3. Denial of service attacks (DoS)
 4. Email loops
- f. The Service Credit available for the Email Delivery Service is:

Average Email Delivery Time (as defined above)	Service Credit
> 1	25%
> 4	50%
> 10	100%