Contract #: AR3090

STATE OF UTAH COOPERATIVE CONTRACT

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

Capgemini America, Inc.
Name

400 Broadacres Drive - 4th Floor; Suite 410
Street Address

Bloomfield	New Jersey	07003
City	*State*	*Zip*

Vendor # VC226513 Commodity Code #: 920-05 Legal Status of Contractor: For-Profit Corporation

Contact *Name:* Mark Stein *Phone Number:* 972-556-7606 *Email:* Mark.Stein@Capgemini.com

2. CONTRACT PORTFOLIO NAME: Cloud Solutions.

3. GENERAL PURPOSE OF CONTRACT: Provide Cloud Solutions under the service models awarded in Attachment B.

4. PROCUREMENT: This contract is entered into as a result of the procurement process on FY2018, Solicitation# SK18008

5. CONTRACT PERIOD: Effective Date: Wednesday, July 31, 2019. Termination Date: Tuesday, September 15, 2026 unless terminated early or extended in accordance with the terms and conditions of this contract.

6. Administrative Fee: Contractor shall pay to NASPO ValuePoint, or its assignee, a NASPO ValuePoint Administrative Fee of one-quarter of one percent (0.25% or 0.0025) of contract sales no later than 60 days following the end of each calendar quarter. The NASPO ValuePoint Administrative Fee shall be submitted quarterly and is based on sales of the Services.

7. ATTACHMENT A: NASPO ValuePoint Master Terms and Conditions, including the attached Exhibits
ATTACHMENT B: Scope of Services Awarded to Contractor
ATTACHMENT C: Pricing Discounts and Schedule
ATTACHMENT D: Contractor's Response to Solicitation # SK18008
ATTACHMENT E: Service Offering EULAs, SLAs

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

9. DOCUMENTS INCORPORATED INTO THIS CONTRACT BY REFERENCE BUT NOT ATTACHED:
   a. All other governmental laws, regulations, or actions applicable to the goods and/or services authorized by this contract.
   b. Utah Procurement Code, Procurement Rules, and Contractor's response to solicitation #SK18008.

10. Each signatory below represents that he or she has the requisite authority to enter into this contract.

IN WITNESS WHEREOF, the parties sign and cause this contract to be executed. Notwithstanding verbal or other representations by the parties, the "Effective Date" of this Contract shall be the date provided within Section 5 above.

CONTRACTOR

Mark Stein
Mark Stein (Jul 30, 2019)

Jul 30, 2019

Contractor's signature — Date

Mark Stein — July 30, 2019

Type or Print Name and Title

DIVISION OF PURCHASING

Jul 30, 2019

Director, Division of Purchasing — Date

Attachment A: NASPO ValuePoint Master Agreement Terms and Conditions

1. Master Agreement Order of Precedence

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

(1) A Participating Entity's Participating Addendum[1] ("PA");
(2) NASPO ValuePoint Master Agreement Terms & Conditions, including the applicable Exhibits[2] to the Master Agreement;
(3) The Solicitation;
(4) Contractor's response to the Solicitation, as revised (if permitted) and accepted by the Lead State; and
(5) A Service Level Agreement issued against the Participating Addendum.

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

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

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

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

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

Data means all information, whether in oral or written (including electronic) form, created by or in any way originating with a Participating Entity or Purchasing Entity, and all information that is the output of any computer processing, or other electronic manipulation, of any information that was created by or in any way originating with a Participating Entity or Purchasing Entity, in the course of using and configuring the Services provided under this Agreement.

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

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

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

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

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

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

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

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

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

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

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

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

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

Participating Addendum means a bilateral agreement executed by a Contractor and a Participating Entity incorporating this Master Agreement and any other additional Participating Entity specific language or other requirements, e.g. ordering procedures specific to the Participating Entity, other terms and conditions.

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

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

Upon execution of the Participating Addendum, a Participating State becomes a Participating Entity.

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

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

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

Protected Health Information (PHI) means individually identifiable health information transmitted by electronic media, maintained in electronic media, or transmitted or maintained in any other form or medium. PHI excludes education records covered by the Family Educational Rights and Privacy Act (FERPA), as amended, 20 U.S.C. 1232g, records described at 20 U.S.C. 1232g(a)(4)(B)(iv) and employment records held by a covered entity in its role as employer. PHI may also include information that is a subset of health information, including demographic information collected from an individual, and (1) is created or received by a health care provider, health plan, employer or health care clearinghouse; and (2) relates to the past, present or future physical or mental health or condition of an individual; the provision of health care to an individual; or the past, present or future payment for the provision of health care to an individual; and (a) that identifies the individual; or (b) with respect to which there is a reasonable basis to believe the information can be used to identify the individual.

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

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

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

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

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

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

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

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

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

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

The Lead State reserves the right to assign any rights or duties, including written assignment of contract administration duties to the NASPO Cooperative Purchasing

Organization LLC, doing business as NASPO ValuePoint.

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

7. Termination: Unless otherwise stated, this Master Agreement may be terminated by either party upon 120 days written notice prior to the effective date of the termination. Further, any Participating Entity may terminate its participation upon 90 days written notice, unless otherwise limited or stated in the Participating Addendum. Termination may be in whole or in part. Any termination under this provision shall not affect the rights and obligations attending orders outstanding at the time of termination, including any right of any Purchasing Entity to indemnification by the Contractor, rights of payment for Services delivered and accepted, data ownership, Contractor obligations regarding Purchasing Entity Data, rights attending default in performance an applicable Service Level of Agreement in association with any Order, Contractor obligations under Termination and Suspension of Service, and any responsibilities arising out of a Security Incident or Data Breach. Termination of the Master Agreement due to Contractor default may be immediate.

8. Confidentiality, Non-Disclosure, and Injunctive Relief

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

b. Non-Disclosure. Each Party shall hold Confidential Information in confidence, using at least the industry standard of confidentiality, and shall not copy, reproduce, sell, assign, license, market, transfer or otherwise dispose of, give, or disclose Confidential Information to third parties or use Confidential Information for any purposes whatsoever other than what is necessary to the performance of Orders placed under this Master Agreement. Contractor shall advise each of its employees and agents of their obligations to keep Confidential Information confidential. Contractor shall use commercially reasonable efforts to assist Purchasing Entity in identifying and preventing any unauthorized use or disclosure of any Confidential Information. Without limiting the

generality of the foregoing, Contractor shall advise Purchasing Entity, applicable Participating Entity, and the Lead State immediately if Contractor learns or has reason to believe that any person who has had access to Confidential Information has violated or intends to violate the terms of this Master Agreement, and Contractor shall at its expense cooperate with Purchasing Entity in seeking injunctive or other equitable relief in the name of Purchasing Entity or Contractor against any such person. Except as directed by Purchasing Entity, Contractor will not at any time during or after the term of this Master Agreement disclose, directly or indirectly, any Confidential Information to any person, except in accordance with this Master Agreement, and that upon termination of this Master Agreement or at Purchasing Entity's request, Contractor shall turn over to Purchasing Entity all documents, papers, and other matter in Contractor's possession that embody Confidential Information. Notwithstanding the foregoing, Contractor may keep one copy of such Confidential Information necessary for quality assurance, audits and evidence of the performance of this Master Agreement.

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

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

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

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

10. Defaults and Remedies

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

(1) Nonperformance of contractual requirements; or

(2) A material breach of any term or condition of this Master Agreement; or

(3) Any certification, representation or warranty by Contractor in response to the solicitation or in this Master Agreement that proves to be untrue or materially misleading; or

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

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

b. Upon the occurrence of an event of default, the non-defaulting party shall issue a written notice of default, identifying the nature of the default, and providing a period of 30 calendar days in which the 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 liability for damages.

c. If Contractor is afforded an opportunity to cure and fails to cure the default within the period specified in the written notice of default, Contractor shall be in breach of its obligations under this Master Agreement and Lead State shall have the right to exercise any or all of the following remedies:

(1) Exercise any remedy provided by law; and

(2) Terminate this Master Agreement and any related Contracts or portions thereof; and

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

(4) Suspend Contractor's performance; and

(5)

d. Contractor may discontinue performance if: (a) the Purchasing Entity fails to pay amounts due; and (b) after a 60-day written notice, the Participating Entity has not cured a breach.

e. 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. Each party shall defend, indemnify and hold harmless the other (for purposes of NASPO, this includes: 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 from third parties including reasonable attorneys' fees and related costs for any death, injury, or damage to tangible property (excluding software or data) arising directly or indirectly the negligence or willful misconduct of the indemnitor, its employees or subcontractors or volunteers, at any tier, relating to the performance under the Master Agreement. Notwithstanding anything to the contrary contained herein, Contractor has no obligation for any indemnification claim arising out of or resulting from the acts or omissions of Participating Entities or Purchasing Entities or any of their respective employees, contractors or agents

b. Indemnification – Intellectual Property. The Contractor shall defend, indemnify and hold harmless NASPO, NASPO ValuePoint, the Lead State, Participating Entities, Purchasing Entities, along with their officers, agents, and employees as well as any person or entity for which they may be liable ("Indemnified Party"), from and against claims, damages or causes of action including reasonable attorneys' fees and related costs arising out of the claim that the Product or its use, infringes Intellectual Property rights ("Intellectual Property Claim") of another person or entity.

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

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

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

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

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

(2) The Indemnified Party shall notify the Contractor within a reasonable time after receiving notice of an Intellectual Property Claim. Even if the Indemnified Party fails to provide reasonable notice, the Contractor shall not be relieved from its obligations unless the Contractor can demonstrate that it was prejudiced in defending the Intellectual Property Claim resulting in increased expenses or loss to the Contractor and then only to the extent of the prejudice or expenses. If the Contractor promptly and reasonably investigates and defends any Intellectual Property Claim, it shall have control over the defense and settlement of it. However, the Indemnified Party must consent in writing for any money damages or obligations for which it may be responsible. The Indemnified Party shall furnish, at the Contractor's reasonable request and expense, information and assistance necessary for such defense. If the Contractor fails to vigorously pursue the defense or settlement of the Intellectual Property Claim, the Indemnified Party may assume the defense or settlement of it and the Contractor shall be liable for all costs and expenses, including reasonable attorneys' fees and related costs, incurred by the Indemnified Party in the pursuit of the Intellectual Property Claim. Unless otherwise agreed in writing, this section is not subject to any limitations of liability in this Master Agreement or in any other document executed in conjunction with this Master Agreement.

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

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

16. Insurance

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

Participating Addendum.

b. The minimum acceptable limits shall be as indicated below, with Contractor responsible for any deductible for each of the following categories:

(1) Commercial General Liability covering premises operations, independent contractors, products and completed operations, blanket contractual liability, personal injury (including death), advertising liability, and property damage, with a limit of not less than $1 million per occurrence/$3 million general aggregate; The limit requirement may be satisfied through a combination of policies including Umbrella/Excess Liability**.**

(2) CLOUD MINIMUM INSURANCE COVERAGE:

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

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

(4) Professional Liability. As applicable, Professional Liability Insurance Policy in the minimum amount of $10,000,000 per claim and $10,000,000 in the aggregate, and addressing network security and privacy including categories in subsection 2 above.

c. Contractor shall pay premiums on all insurance policies. Such policies shall also reference this Master Agreement Contractor agrees to provide thirty (30) calendar days written notice of intended revocation thereof shall have been given to Purchasing Entity and Participating Entity.

d. Prior to commencement of performance, Contractor shall provide to the Lead State a written endorsement to the Contractor's general liability insurance policy or other documentary evidence acceptable to the Lead State that (1) includes the Participating States identified in the Request for Proposal as additional insureds which may be met through the use of a "blanket" insured endorsement, and (2) provides that the Contractor's liability insurance policy shall be primary, with any liability insurance of any Participating State as secondary and noncontributory. Contractor also agrees to provide thirty (30) days written notice of any material alteration, cancellation, non-renewal, or expiration of the coverage contained in this Agreement. Unless otherwise agreed in any Participating Addendum, the Participating Entity's rights and Contractor's obligations are the same as those specified in the first sentence of this subsection. Before performance of any Purchase Order issued after execution of a Participating

Addendum authorizing it, the Contractor shall provide to a Purchasing Entity or Participating Entity who requests it the same information described in this subsection.

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

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

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

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

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

19. Ordering

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

b. This Master Agreement permits Purchasing Entities to define project-specific requirements and informally compete the requirement among other firms having a Master Agreement on an "as needed" basis. This procedure may also be used when requirements are aggregated or other firm commitments may be made to achieve

reductions in pricing. This procedure may be modified in Participating Addenda and adapted to Purchasing Entity rules and policies. The Purchasing Entity may in its sole discretion determine which firms should be solicited for a quote. The Purchasing Entity may select the quote that it considers most advantageous, cost and other factors considered.

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

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

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

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

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

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

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

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

termination of this Master Agreement, notwithstanding the term of any such indefinite delivery order agreement.

20. Participants and Scope

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

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

c. Unless otherwise stipulated in a Participating Entity's Participating Addendum, specific services accessed through the NASPO ValuePoint cooperative Master Agreements for Cloud Services by state executive branch agencies, as required by a Participating Entity's statutes, are subject to the authority and approval of the Participating Entity's Chief Information Officer's Office[3].

d. Obligations under this Master Agreement are limited to those Participating Entities who have signed a Participating Addendum and Purchasing Entities within the scope of those Participating Addenda. States or other entities permitted to participate may use an informal competitive process to determine which Master Agreements to participate in through execution of a Participating Addendum. Financial obligations of Participating States are limited to the orders placed by the departments or other state agencies and institutions having available funds. Participating States incur no financial obligations on behalf of political subdivisions.

e. NASPO ValuePoint is not a party to the Master Agreement. It is a nonprofit cooperative purchasing organization assisting states in administering the NASPO

[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.

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.

21. Payment: Orders under this Master Agreement are fixed-price or fixed-rate orders, not cost reimbursement contracts. Unless otherwise stipulated in the Participating Addendum, Payment is normally made within 30 days following the date of a correct invoice is received. Purchasing Entities reserve the right to withhold disputed amounts of an invoice up to an amount equal to two (2) months' invoices. After 45 days the Contractor may assess overdue account charges up to a maximum rate of one percent per month on the outstanding balance. Payments will be remitted by mail. Payments may be made via a State or political subdivision "Purchasing Card" with no additional charge.

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

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

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

23. Operations Management: Contractor shall maintain the administrative, physical, technical, and procedural infrastructure associated with the provision of the Product in a manner that is, at all times during the term of this Master Agreement, at a level equal to or more stringent than those specified in the Solicitation.

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

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

Contractor shall not use any information collected in connection with this Master Agreement, including Purchasing Entity Data, for any purpose other than fulfilling its obligations under this Master Agreement.

26. Records Administration and Audit.

a. The Contractor shall maintain books, records, documents, and other evidence pertaining to this Master Agreement and orders placed by Purchasing Entities under it to the extent and in such detail as shall adequately reflect performance and administration of payments and fees. Contractor shall permit the Lead State, a Participating Entity, a Purchasing Entity, the federal government (including its grant awarding entities and the U.S. Comptroller General), and any other duly authorized agent of a governmental agency, to audit, inspect, examine, copy and/or transcribe Contractor's books, documents, papers and records directly pertinent to this Master Agreement or orders placed by a Purchasing Entity under it for the purpose of making audits, examinations, excerpts, and transcriptions. This right shall survive for a period of six (6) years following termination of this Agreement or final payment for any order placed by a Purchasing Entity against this Agreement, whichever is later, to assure compliance with the terms hereof or to evaluate performance hereunder.

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

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

d. The Contractor shall allow the Purchasing Entity to audit conformance to the Master Agreement and applicable Participating Addendum terms. The purchasing entity may perform this audit or contract with a third party (who is not a competitor of the Contractor) 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 commercially reasonable efforts to restore or assist in restoring the system to operational capacity.

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

30. Data Privacy: The Contractor must comply with all applicable laws related to data privacy and security, including IRS Pub 1075 to the extent applicable to the operation of its business in the provision of the Services. Prior to entering into a SLA with a Purchasing Entity, the Contractor and Purchasing Entity must cooperate and hold a meeting to determine the Data Categorization to determine whether the Contractor will hold, store, or process High Risk Data, Moderate Risk Data and Low Risk Data. The Contractor must document the Data Categorization in the SLA or Statement of Work.

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

a. Contractor has acquired any and all rights, grants, assignments, conveyances,

licenses, permissions, and authorization for the Contractor to provide the Services described in this Master Agreement.

b. Contractor will perform materially as described in this Master Agreement, SLA, Statement of Work, including any performance representations contained in the Contractor's response to the Solicitation by the Lead State.

c. Contractor represents and warrants that the representations contained in its response to the Solicitation by the Lead State.

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

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

f. The Contractor warrants that it will use commercially reasonable efforts to cause Products it provides under this Master Agreement to be free of malware. The Contractor must use industry-leading technology to detect and remove worms, Trojans, rootkits, rogues, dialers, spyware, etc.

g. Purchasing Entity represents, warrants and covenants that it shall perform its obligations and responsibilities under the Participating Addendum in a manner that does not infringe or misappropriate, or constitute an infringement or misappropriation of, any patent, copyright, trademark, trade secret or other intellectual property, proprietary or privacy rights of Contractor or any third party

h. Contractor's sole liability with respect to warranty claims related to material performance of its obligations shall be the reperformance of the deficient services or, if unable to reperform, replace or to void the invoiced amount for such services. Further, Contractor is not responsible for warranty claims resulting from third-party hardware or software failures, or third-party actions.

i. Purchasing Entity acknowledges and agrees that all software or hardware installed by Contractor is subject to the respective manufacturer's warranty.

EXCEPT AS OTHERWISE EXPRESSLY PROVIDED IN THE AGREEMENT, THE PARTIES MAKE NO REPRESENTATIONS, WARRANTIES OR CONDITIONS (STATUTORY OR OTHERWISE), EXPRESS OR IMPLIED, REGARDING ANY MATTER, INCLUDING THE MERCHANTABILITY, SUITABILITY, FITNESS FOR A PARTICULAR USE OR PURPOSE, OR RESULTS TO BE DERIVED FROM THE USE OF ANY SERVICE, SOFTWARE, HARDWARE, DELIVERABLES OR OTHER MATERIALS PROVIDED UNDER THE AGREEMENT.

32. Transition Assistance:

a. The Contractor shall reasonably cooperate with other parties in connection with all Services to be delivered under this Master Agreement, including without limitation any successor service provider to whom a Purchasing Entity's Data is transferred in connection with the termination or expiration of this Master Agreement. The Contractor shall assist a Purchasing Entity in exporting and extracting a Purchasing Entity's Data, in a format usable without the use of the Services and as agreed by a Purchasing Entity, at no additional cost to the Purchasing Entity. Any transition services requested by a Purchasing Entity involving additional knowledge transfer and support may be subject to a separate transition Statement of Work.

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

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

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

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

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

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

37. Governing Law and Venue

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

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

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

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

38. No Guarantee of Service Volumes: The Contractor acknowledges and agrees that the Lead State and NASPO ValuePoint makes no representation, warranty or condition

as to the nature, timing, quality, quantity or volume of business for the Services or any other products and services that the Contractor may realize from this Master Agreement, or the compensation that may be earned by the Contractor by offering the Services. The Contractor acknowledges and agrees that it has conducted its own due diligence prior to entering into this Master Agreement as to all the foregoing matters.

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

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

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

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

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

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

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

tool).

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

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

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

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

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

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

a. Contractor agrees to work cooperatively with NASPO ValuePoint personnel. Contractor agrees to present plans to NASPO ValuePoint for the education of Contractor's contract administrator(s) and sales/marketing workforce regarding the Master Agreement contract, including the competitive nature of NASPO ValuePoint procurements, the Master agreement and participating addendum process, and the manner in which qualifying entities can participate in the Master Agreement.

b. Contractor agrees, as Participating Addendums become executed, if requested by ValuePoint personnel to provide plans to launch the program within the participating state. Plans will include time frames to launch the agreement and confirmation that the Contractor's website has been updated to properly reflect the contract offer as available in the participating state.

c. Contractor agrees, absent anything to the contrary outlined in a Participating Addendum, to consider customer proposed terms and conditions, as deemed important to the customer, for possible inclusion into the customer agreement. Contractor will ensure that their sales force is aware of this contracting option.

d. Contractor agrees to participate in an annual contract performance review at a location selected by the Lead State and NASPO ValuePoint, which may include a discussion of marketing action plans, target strategies, marketing materials, as well as Contractor reporting and timeliness of payment of administration fees.

e. Contractor acknowledges that the NASPO ValuePoint logos may not be used by Contractor in sales and marketing until a logo use agreement is executed with NASPO ValuePoint.

f. The Lead State expects to evaluate the utilization of the Master Agreement at the annual performance review. Lead State may, in its discretion, terminate the Master Agreement pursuant to section 6 when Contractor utilization does not warrant further administration of the Master Agreement. The Lead State may exercise its right to not renew the Master Agreement if vendor fails to record or report revenue for three consecutive quarters, upon 60-calendar day written notice to the Contractor. This subsection does not limit the discretionary right of either the Lead State or Contractor to terminate the Master Agreement pursuant to section 7.

g. Contractor agrees, within 30 days of their effective date, to notify the Lead State and NASPO ValuePoint of any contractual most-favored-customer provisions in third-part contracts or agreements that may affect the promotion of this Master Agreements or whose terms provide for adjustments to future rates or pricing based on rates, pricing in, or Orders from this master agreement. Upon request of the Lead State or NASPO ValuePoint, Contractor shall provide a copy of any such provisions.

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

46. Savings Clause. Contractor shall not be liable for any delay to or failure or non-performance of its obligations under the Agreement to the extent caused by any act or omission of [state, purchasing state, etc. include any third parties that may be working on the state's behalf] or an act or omission of any auditors or agents of State, or any act or omission of any other third party acting on behalf of State which causes the Contractor not to be able to perform its obligations under the Agreement and which Contractor could not reasonably have avoided without incurring material additional expenditure; provided that Contractor has used all commercially reasonable efforts to avoid the impact of the relevant failure or default in a timely manner. Contractor shall provide the State with notice of such non-performance promptly after becoming aware of the occurrence of such an event and shall attempt performance notwithstanding the State's failure to perform, with the State reimbursing Contractor for its additional and reasonable out-of-pocket expenses incurred in undertaking such efforts to the extent these have been agreed in advance by the State.

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

Contractor's liability for any claim, loss or liability arising out of, or connected with the Services provided, and whether based upon default o other liability such as breach of contract, warranty Negligence, misrepresentation or otherwise, shall in no case exceed direct damages in: (i) an amount not to exceed a total of twelve (12) months charges or (ii) five million dollars ($5, 000,000), whichever is greater.

Notwithstanding the above, neither the Contractor nor the Purchasing Entity shall be liable for any consequential, indirect or special damages of any kind which may result directly or indirectly from such performance, including, without limitation, damages resulting from loss of use or loss of profit by the Purchasing Entity, the Contractor, or by others. The limitation of liability in this Section will not apply to claims for bodily injury or death.

48. Entire Agreement: This Master Agreement, along with any attachment, contains the entire understanding of the parties hereto with respect to the Master Agreement unless a term is modified in a Participating Addendum with a Participating Entity. No click-through, or other end user terms and conditions or agreements required by the Contractor ("Additional Terms") provided with any Services hereunder shall be binding on Participating Entities or Purchasing Entities, even if use of such Services requires an affirmative "acceptance" of those Additional Terms before access is permitted.

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

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

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

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

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

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

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

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

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

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

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

4. Security Incident or Data Breach Notification:

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

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

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

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

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

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

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

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

7. Termination and Suspension of Service:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Attachment B – Scope of Services Awarded to Contractor

1.1 Awarded Service Model(s).

Contractor is awarded the following Service Model:

- Software as a Service (SaaS)

1.2 Risk Categorization.*

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

Service Model	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered
SaaS	Provided, using DocFinity® SaaS solution	Provided, using DocFinity® SaaS solution	Provided, using DocFinity® SaaS solution	Private Cloud (two options) 1. Hosted in our secure, HIPAA- compliant Companion Data Services Data Center in Columbia, South Carolina (USA) 2. Hosted in Amazon Web Services government-rated cloud

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

2.1 Deployment Models.

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

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

together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds)

Attachment C - Pricing Discounts and Schedule

Contractor: Capgemini America, Inc.

Pricing Notes

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

Cloud Service Model: Infrastructure as a Service (IaaS)

Description	Minimum Discount % Off
OEM: Amazon Web Services (AWS)	2.00%
OEM: Microsoft Corporation - Azure	7.25%
OEM: Virtustream	5.75%
Average IaaS OEM Discount Off	**5.00%**

Cloud Service Model: Platform as a Service (PaaS)

Description	Discount
OEM: Amazon Web Services (AWS)	2.00%
OEM: Microsoft Corporation - Azure	7.25%
OEM: Capgemini iPaaS	9.75%
Average PaaS OEM Discount Off	**6.33%**

Cloud Service Model: Software as a Service (SaaS)

Description	Discount
OEM: Amazon Web Services (AWS)	2.00%
OEM: Microsoft Corporation - Azure	9.75%
OEM: BMC - Remedy	20.25%
OEM: ServiceNow	5.75%
Average SaaS OEM Discount Off	**9.44%**

Additional Value Added Services

	Onsite Hourly Rate**		Remote Hourly Rate	
Item Description	**NVP Price**	**Catalog Price**	**NVP Price**	**Catalog Price**
Maintenance Services	n/a	n/a	n/a	n/a
Professional Services	n/a	n/a	n/a	n/a
Deployment Services	n/a	n/a	n/a	n/a
Integration Services)	n/a	n/a	n/a	n/a
Consulting/Advisory Services	n/a	n/a	n/a	n/a
Architectural Design Services	n/a	n/a	n/a	n/a
Statement of Work Services	n/a	n/a	n/a	n/a
Partner Services	n/a	n/a	n/a	n/a
Training Deployment Services	n/a	n/a	n/a	n/a
Service Delivery Manager	$288.42	$339.32	$288.42	$339.32
Program Manager	$247.00	$290.59	$247.00	$290.59
Project Manager	$202.74	$238.52	$202.74	$238.52
Transition Manager	$202.74	$238.52	$202.74	$238.52
Project Manager	$202.74	$238.52	$202.74	$238.52
Project Analyst	$143.83	$169.21	$143.83	$169.21
Team Manager/Leader	$101.63	$119.57	$101.63	$119.57
Operations Manager	$102.66	$120.78	$102.66	$120.78
Senior Delivery Architect	$288.42	$339.32	$288.42	$339.32
Infra Architect	$202.74	$238.52	$202.74	$238.52
Cloud Lead Architect	$247.00	$290.59	$247.00	$290.59
Cloud Network Architect	$202.74	$238.52	$202.74	$238.52
Cloud Infrastructure Lead	$171.10	$201.29	$171.10	$201.29

Attachment C - Pricing Discounts and Schedule

Contractor: Capgemini America, Inc.

Cloud Administrator	$171.10	$201.29	$171.10	$201.29
Cloud Infrastructure Engineer	$143.83	$169.21	$143.83	$169.21
Cloud Network Engineer	$143.83	$169.21	$143.83	$169.21
SIAM Architect	$247.00	$290.59	$247.00	$290.59
SIAM Workstream Lead	$202.74	$238.52	$202.74	$238.52
SIAM Infrastructure Lead	$171.10	$201.29	$171.10	$201.29
SIAM Administrator	$171.10	$201.29	$171.10	$201.29
SIAM Process Analyst	$143.83	$169.21	$143.83	$169.21
Database Architect	$247.00	$290.59	$247.00	$290.59
Database Administrator	$171.10	$201.29	$171.10	$201.29
CyberSecurity Architect	$274.59	$323.05	$274.59	$323.05
Senior CyberSecurity Consultant	$225.39	$265.17	$225.39	$265.17
CyberSecurity Consultant	$190.21	$223.78	$190.21	$223.78
CyberSecurity Analyst	$159.89	$188.11	$159.89	$188.11
iPaaS Architect	$375.65	$441.94	$375.65	$441.94
IPaaS Delivery Manager	$245.67	$289.02	$245.67	$289.02
IPaaS Platform Engineer	$141.04	$165.93	$141.04	$165.93
Traning Coordinator	$247.00	$290.59	$247.00	$290.59
Training Analyst	$171.10	$201.29	$171.10	$201.29
Training Lead	$143.83	$169.21	$143.83	$169.21
Help Desk 1st Line IT Service Desk	$47.35	$55.70	$47.35	$55.70
Help Desk Resolver	$67.67	$79.61	$67.67	$79.61
Help Desk Reporting Anaylst	$84.91	$99.89	$84.91	$99.89
Help Desk Knowledge Manager	$101.63	$119.57	$101.63	$119.57
Help Desk/OCM Manager	$102.66	$120.78	$102.66	$120.78

**Onsite Hourly rates stated are a fully burdened rate that include labor, overhead, and any other costs related to the service. These rates do not include any per diem and/or travel costs. Capgemini will discuss with individual Purchasing Entities, and determine any necessary per Diem and/or travel costs.

Capgemini | NASPO ValuePoint

Technical Requirements
Capgemini Response

July 6, 2018

Table of Contents

8. Technical Requirements

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

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

8.1 (M)(E) TECHNICAL REQUIREMENTS

8.1.1 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 characteristics defined in NIST Special Publication 800-145.

Capgemini has extensive experience in the services required in this contract and will demonstrate in this response the breadth of our capabilities and our experience delivering these services in the market. We deliver services that meet the essential characteristics of the NIST Publication 800-145 in all three areas of IaaS, PaaS, and SaaS. We bring value-added professional services defined in detail throughout this technical requirements response. These services will assist the State of Utah and other Purchasing Entities in designing and implementing a cloud strategy that eases the transformation from the private data center, to hybrid cloud implementations, to full cloud-first strategies around the elastic consumption of public clouds.

Below is a table that describes the NIST compliant offerings for NASPO Purchasing Entities on which Capgemini proposed solutions will be built:

NIST Compliant Offerings	IaaS	PaaS	SaaS
Public Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft
Community Cloud	Amazon Microsoft	Amazon Microsoft	Amazon Microsoft
Hybrid Cloud	Virtustream	Capgemini (iPaaS)	ServiceNow BMC Remedy on Demand
Private Cloud	Virtustream	Capgemini (iPaaS)	ServiceNow BMC Remedy on Demand

NIST Essential Characteristics:

With the breadth and depth of these offerings, Capgemini meets all the essential characteristics of the NIST publication 800-145 by providing services that are self-service, have broad network access across private and public networks, have rapid provisioning and decommissioning characteristics and are measured and billed in a cloud consumption model. In **Section 8.1.3** we will elaborate on the individual product offerings.

NIST Service Models:

Capgemini is meeting all **three cloud service models** (IaaS, PaaS, SaaS) by partnering with Amazon Web Services (AWS), Microsoft Azure services, along with Virtustream for Private Hosting. Capgemini

is also partnering with ServiceNow and BMC to offer Software as Service capabilities that naturally extend onto Cloud environments to assist in the ITIL management of services via customer portals/dashboards, instrumented processes, automation, and measured billing capabilities. Capgemini also offers our unique API service iPaaS, which is a Platform as a Service model to extended application programming interfaces to interconnect multiple cloud environments seamlessly.

NIST Deployment Models:

Capgemini is meeting all four **cloud deployment methods**, Public, Community, Hybrid, and Private. Capgemini will provide Low and Moderate risk data categorizations of computing and storage resources, such as Amazon EC2 and Microsoft Azure virtual environments, which are suitable public and community deployment methods. Capgemini will partner with Virtustream to provide High-risk data environments where Hybrid and Private environments are needed for such sensitive information or private information is a key characteristic in that application environment. In **Section 8.1.3** we will break out the four deployment methods in greater detail and label the data risk categorizations.

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

Capgemini **certifies** we are willing to **comply** with the requirements of **Attachment C - NIST Service Models** IaaS, PaaS, and SaaS models and **Attachment D - Scope of Services**. Capgemini accomplishes this by providing such service models for Low, Moderate, and High-risk data categorizations, and the five essential characteristics, which include on-demand self-service, broad network access, resource pooling, rapid elasticity, and a measured service with Service Level Agreements. In **Section 8.1.3** we will describe our Capgemini PaaS capabilities, our partner IaaS, PaaS, and SaaS capabilities and overall deployment models and services to meet all NIST compliance checkmarks and summarize our Product Offerings Catalog that fulfills all cloud service and deployment methods.

8.1.3 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**.

Our response to this section is organized using the following outline. In it, we describe how our offerings adhere to the services, definitions, and deployment models identified in Attachment D. Capgemini structured this section as a direct response from Attachment D and address all key areas, broken into the partners and direct services that we offer.

- Cloud Based Services Providers
- Categorization of Risk
- Services Characteristics
  - BMC & ServiceNow
  - Amazon, Microsoft, Virtustream
  - Capgemini
- Service Models
  - BMC & ServiceNow
  - Amazon, Microsoft, Virtustream
  - Capgemini
- Deployment Methods
  - BMC & ServiceNow
  - Amazon, Microsoft, Virtustream
  - Capgemini

- Cloud Services Providers Offering Descriptions
  - BMC Remedy on Demand and ServiceNow
  - Amazon AWS
  - Microsoft Azure
  - Virtustream
  - Capgemini Enterprise iPaaS

Cloud Based Services Providers

Capgemini is a Cloud Based Service Provider meeting the NIST definition from Attachment D, by reselling IaaS, PaaS, and SaaS services, coordinating access to cloud solutions, and fulfilling the offerings through a channel of authorized partners. In addition, Capgemini has developed and offers iPaaS Cloud Service as a direct API integration platform to connect various Cloud Service providers together.

If a section, such as 8.2 Subcontractors as an example, applies to all Cloud Service Providers Capgemini did not break out each partner in its own subsection. For sections where different criteria or solution requirements differ we subdivided out each partner in its own section header.

Through this document, Capgemini will refer to Cloud Service Providers and or Partners in the following context.

- Amazon Web Services
- Microsoft Azure Services
- Virtustream for Private Hosting
- ServiceNow
- BMC Remedy on Demand
- Capgemini iPaaS

Categorization of Risk

The following table from Attachment G outlines the Capgemini offerings according to the NIST data risk profiles, addressing the ability to store and secure all data profiles.

Service Model	Low Risk Data	Moderate Risk Data	High Risk Data	Deployment Models Offered:
SaaS	Amazon Microsoft ServiceNow BMC	Amazon Microsoft ServiceNow BMC		Private, Community, Hybrid, Public
IaaS	Amazon Microsoft Virtustream	Amazon Microsoft Virtustream	Virtustream	Private, Community, Hybrid, Public
PaaS	Amazon Microsoft	Amazon Microsoft		Private, Community, Hybrid, Public

Service Characteristics

BMC & ServiceNow (SaaS)

Capgemini is offering two Service Management tools as SaaS offerings; BMC Remedy on Demand and ServiceNow. In keeping with the NIST essential characteristics, both solutions are offered on a SaaS delivery model deployed on Private Cloud infrastructure. The following table applies the NIST characteristics to BMC and ServiceNow solutions.

Service Characteristics	BMC Remedy on Demand	ServiceNow
On-demand self-service	Automated provisioning of instances post-commercial transaction processing	
Broad network access	All functionality available for all levels of access through a web browser, network access, and appropriately negotiated authorization protocols via internet or intranet access.	
Resource pooling	The specific instances of the deployed application are not multi-tenant by design. As a SaaS service model is used, specific hardware allocations are immaterial and irrelevant.	
Rapid elasticity	Allocation of resources to meet demand is done on a rapid basis to meet the agreed performance service levels. The number of computing resources consumed as a result is not a concern as the service is measured.	
Measured service	Basic access to software functionality is generally measured through User Subscriptions that are based on a concurrent or named basis. Other services may be measured differently.	Basic access to software is measured through User Subscriptions that are based on a named basis. Other add-on services may be measured differently.

Amazon, Microsoft, and Virtustream (IaaS, PaaS, SaaS)

Capgemini will resell the Amazon AWS, Microsoft Azure, and Virtustream cloud services to the Purchasing Entities. In the table below, we explain how the three cloud services adhere to provide the five NIST service characteristics.

Service Characteristics	Amazon AWS	Microsoft Azure	Virtustream
On-Demand Self-service	AWS enables self-service through the AWS Management client portal. Through this portal, end users can provision computing power, storage, networks and software in a simple and flexible way. On-demand self-service allows end users to request resources quickly, enabling a pay as you go model for resources that have been provisioned.	Microsoft Azure provides a Management portal that provides the self-service capabilities to quickly and easily provision compute, networking, storage and software services. The on-demand self-service capability allows for a pay as you go model for consumption of resources provisioned.	Virtustream's xStream Cloud Management Platform provides a unified, cloud-agnostic control plane that integrates infrastructure orchestration, enterprise application automation and a suite of business intelligence and service management tools. This powerful combination allows for the delivery of enterprise IT as a service, with true consumption economics, application service-level

Service Characteristics	Amazon AWS	Microsoft Azure	Virtustream
			agreements, and best-in-class security and compliance features.
Broad network access	AWS provides direct network connections via customer supplied Internet or direct MPLS, VPN.	Azure provides direct network connections via customer supplied Internet or direct MPLS, VPN.	Virtustream provides direct network connections via customer supplied Internet or direct MPLS, VPN.
Resource Pooling	AWS provides resource pooling capabilities by leveraging EC2 for computing, S3, and Elastic Block Storage volumes for storage pools and networking configurations through virtual private cloud isolation and pooling of networking services, subnets, network groups and security controls.	Microsoft Azure computing fabric allows group resources to generate resource pools for computing resources in particular virtual machines. Blob, Table, Queue, and databases are also grouped to generate Storage pool configurations. Networking resources can also be allocated in by provisioning virtual networks for isolation and pooling of networking services, defining network security groups and subnets establish to secure access to pooled resources.	Virstustream xStream is the first true enterprise hybrid cloud solution—secure and compliant whether private or public. Enterprises need to run a complex IT environment that supports multiple applications, with a broad mix of hardware and software. xStream provides a hybrid cloud solution—private, public or mixed—which enables the enterprise to securely benefit from the cloud with multi-tenant efficiency, dynamic scale, and on-demand delivery. Enterprises can now operate secure, compliant private clouds on-site and combine them with virtual private/public clouds in a hybrid solution for both legacy and web-scale applications.
Rapid Elasticity	Amazon AWS provides EC2 for elastic computing services. EC2 provides services to the provision as many or as few virtual servers as needed, configuring security and networking, and storage as necessary. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes on workloads. Elastic Beanstalk automatically handles the web application deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring.	Microsoft developed the Azure Fabric to provide elastic computing services that allow the ability to quickly expand or decrease computer processing, memory, and storage resources to meet changing workload and web application demands. Azure Service Fabric allows the automated provisioning and management of hundreds and thousands	Virtustream supports a consumption-based resource model which can dynamically allocate the correct number of virtual servers, memory, storage and another solution component to fulfill your current and future computing needs. µVM, pronounced “microVM,” this is Virtustream’s fine-grained unit of measurement designed

Service Characteristics	Amazon AWS	Microsoft Azure	Virtustream
		of computing resources, containers, and applications to meet the workload demands for cloud-native applications, big data, and machine learning.	to accurately measure the actual consumption of cloud resources.
Measured Service	AWS provides the capability to generate usage reports of services consumed through the AWS Management portal. AWS measures consumption of cloud services provides monitoring of the cloud services provisioned, including billing, use of resource and capacity planning AWS provides a pay as you go model for the usage of their cloud services. AWS provides customers with an account to consume cloud services, also an invoice of what has been consumed throughout the month is generated. Customers are billed by AWS in a monthly a billing cycle for cloud resource consumption.	Azure provides the capability to monitor consumption of cloud resources with plans to forecast spending for cloud resources. Azure provides usage reports of services consumed through the month accessible through the Azure Management portal. Azure provides billing, use of resource and capacity planning for cloud resources provisioned. Azure provides a pay as you go model for the usage of their cloud services. Azure provides customers with a subscription model and account to allow them to provision cloud services. Azure generates an invoice of what has been consumed throughout the month. Azure bills customers on a monthly a billing cycle for cloud resource consumption.	A µVM is a unit of computing resources, comprised of CPU, memory, storage IOPS, and associated local network bandwidth. One µVM = 200MHz of CPU, 768MB of memory, 40 storage disk input/output operations per second (IOPS), and 2 Megabytes (MBps) of local network bandwidth. The usage of each µVM resource component (CPU, memory, IOPS, and local network bandwidth) is measured at five-minute intervals—one unit each for 200MHz of CPU, 768MB of memory, 40 storage disk input/output operations per second (IOPS), and 2MBps of local network bandwidth. The highest of the four components are averaged per hour, and the hour values are averaged across the month to determine the overall µVM usage for the month

Capgemini (PaaS)

Capgemini Enterprise iPaaS (integration Platform as a Service) is our key offering to provide cloud agnostic API and hybrid integration platform services that support agile business processes and data and application integration.

Capgemini Enterprise iPaaS adheres to the NIST service characteristics as identified in the following table.

Service Characteristics	Capgemini Enterprise iPaaS
On-demand self-service	The Capgemini Enterprise iPaaS includes Operational Management tooling to allow the purchaser to deploy integration and microservice payloads onto the platform and to provision additional compute capacity.

Service Characteristics	Capgemini Enterprise iPaaS
Broad network access	The Capgemini Enterprise iPaaS provides access capability from any device, while enforcing authentication and communication security policies.
Resource pooling	The Capgemini Enterprise iPaaS is a single-tenanted solution by design. This limits the possibilities of client data being exposed, as opposed to a multi-tenanted design. Additionally, the single tenanted design also provides for logical and physical separation between non-production and production environments. Our iPaaS will be deployed based on client geographical requirements, with no component being deployed elsewhere. With the iPaaS integration or microservice payloads are dynamically managed so that best use is made from the available resource.
Rapid elasticity	The Capgemini Enterprise iPaaS can be scaled up or down in resource capacity, with a daily charging model applied.
Measured service	All resource usage is monitored and available via graphical monitoring operational tools.

Service Models

The following tables, organized by service offering, depict how Capgemini's cloud offerings adhere to the NIST service models.

BMC & ServiceNow (SaaS)

Service Models	BMC Remedy on Demand	ServiceNow
SaaS	BMC and ServiceNow operate their applications on a Software as a Service basis, providing the Purchasing Entities with the ability to use the provider's applications running on private cloud infrastructure that is hosted in the US. The applications are accessible from various client devices using any of the mainstream web browser interfaces. The Purchasing Entities are free from the need to manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage.	

Amazon, Microsoft, Virtustream (IaaS, PaaS, SaaS)

Service Models	Amazon AWS	Microsoft Azure	Virtustream
SaaS	AWS Software as a Service is in place to manage the services and underlying infrastructure for the cloud application that are developed to run on the platform. AWS provides the Business Applications Services and Solutions to provision business applications, including office software, messaging software, payroll processing software by partners like SAP or Oracle.	Azure SaaS provides a complete software solution that customers purchase on a pay-as-you-go basis Customers rent the use of an app for their organization, users connect to the application or service over the Internet, usually with a web browser. All of the underlying infrastructure, middleware, app software, and app data are located in the Azure data center facilities. Azure manages the hardware and	Capgemini is not offering Virtustream SaaS as part of this contract submission.

Service Models	Amazon AWS	Microsoft Azure	Virtustream
		software, will provide for availability and the security of the app and customer data as well. Azure SaaS allows customers to get quickly up and running with minimal upfront cost. Most common usage scenarios for SaaS are email, calendaring, and office tools such as Microsoft Office 365.	
IaaS	AWS provides Infrastructure as a Service, providing access to networking features, virtual computers, and data storage space. The AWS Infrastructure as a Service implementation provides customers with the highest level of flexibility and management control over their IT resources and is similar to existing IT resources that many IT departments and developers are familiar with today. AWS most common IaaS services are Amazon Elastic Compute Cloud (EC2), Amazon Simple Storage Service and Elastic Load Balancing.	Microsoft Azure Infrastructure as a service (IaaS) is an instant computing infrastructure to provision and managed resources. Allowing to quickly scale up and down on demand. Azure IaaS platform helps avoid the expense and complexity of buying and managing physical servers and other data center infrastructure. Each resource is offered as a separate service component, Azure manages the infrastructure, while the customer will purchase, install, configure, and manage their own software, operating systems, middleware, and applications. Azure IaaS provides a wide range of VM sizes, Containers, and images, enabling buyers to choose the best deployment options for their environment. Azure offers numerous networking services to build sophisticated network topologies and extend datacenters to the cloud. Azure delivers durable, highly available, and massively scalable storage options like Blob, Queue, File, and Disk—that keep pace with explosive data growth.	The Virtustream Service Description (SD) section contains information about Cloud Platform Services (CPS) that comprise the Infrastructure as a Service (IaaS) offering. The services that comprise the Virtustream IaaS offering allow customers to purchase one or more of the services for their cloud solution. IaaS Onboarding and Migration IaaS Compute IaaS Software IaaS Storage/ Backup IaaS Network IaaS Load Balancer IaaS Cloud and Storage Connect IaaS Ad Hoc Services
PaaS	AWS provides Platforms as a Service where it no longer is required to manage the underlying infrastructure (usually hardware and operating systems). This allows customers to focus on the deployment and management of cloud applications, eliminating the need to worry about application maintenance like resource	Azure Platform as a service (PaaS) is a complete development and deployment environment in the cloud, with resources that enable customers to deliver everything from simple cloud-based apps to sophisticated, cloud-enabled enterprise applications. Customers purchase resources from Azure on a pay-as-you-go basis and access them over a	Capgemini is not offering Virtustream PaaS as part of this contract submission.

Service Models	Amazon AWS	Microsoft Azure	Virtustream
	procurement, capacity planning, software maintenance, and patching. For example, AWS Elastic Beanstalk provides the PaaS services to automatically handle the deployment, from capacity provisioning, load balancing, auto-scaling to application health monitoring. At the same time, retaining full control over the AWS resources powering the web applications with access to the underlying resources at any time.	secure Internet connection. Azure PaaS allows customers to avoid the expense and complexity of buying and managing software licenses, the underlying application infrastructure, and middleware or the development tools and other resources. Customers manage the applications and services they develop, and Azure typically manages everything else. Most common PaaS services are App Services to develop and publish web applications and microservices, Azure Search and Azure CDN for content delivery to applications and end users.	

Capgemini (PaaS)

Capgemini Enterprise iPaaS (integration Platform as a Service) is our key offering to provide cloud agnostic API and hybrid integration platform services that support agile business processes and data and application integration.

Capgemini Enterprise iPaaS adheres to the NIST service models as identified in the following table.

Service Models	Capgemini Enterprise iPaaS
PaaS	Capgemini Enterprise iPaaS is provided as a PaaS Managed Service. 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.

Deployment Methods

The following tables depict how the Capgemini offerings adhere to their respective NIST deployment methods.

BMC and ServiceNow (SaaS)

BMC and ServiceNow operate their SaaS model from their own US-based data centers. Though hosted on shared infrastructure, the SaaS instances are discreet, separate, and solely occupied by each Purchasing Entity for their own exclusive use.

Deployment Method	BMC Remedy on Demand	ServiceNow
Private Cloud	BMC operates their SaaS model out of US-based, world-class, third-party data centers.	ServiceNow operates their SaaS model from their own US-based data centers.

Amazon, Azure, and Virtustream (IaaS, PaaS, SaaS)

Deployment Method	Amazon AWS	Microsoft Azure	Virtustream
Public Cloud	Amazon Public Cloud Implementation with Virtual Private Cloud (Amazon VPC) lets customers provision a logically isolated section of the AWS Cloud to launch AWS resources in a virtual network defined by customers. Customers have complete control over their virtual networking environment, including a selection of their own IP address range, the creation of subnets, and configuration of route tables and network gateways. Customers can use both IPv4 and IPv6 in their VPC for secure and easy access to resources and applications.	Microsoft Azure is Microsoft's public cloud computing platform. It provides a range of cloud services, including those for compute, analytics, storage, and networking. Users can pick and choose from these services to develop and scale new applications, or run existing applications, in the public cloud. Azure primarily uses a pay-as-you-go pricing model that charges based on usage.	Virtustream Public Cloud is excluded as part of this contract submission.
Community Cloud	The AWS Cloud provides secure, scalable, and cost-efficient solutions to support the unique requirements and missions of the US federal government. AWS cloud services can be employed to meet mandates, reduce costs, drive efficiencies, and increase innovation across civilian agencies, communities, and the Department of Defense. Cloud computing offers a pay-as-you-go model, delivering access to up-to-date technology resources that are managed by experts.	Microsoft Azure provides government and community cloud that offers hyper-scale compute, storage, networking, and identity management services, with world-class security. A physically and network-isolated instance of Microsoft Azure, operated by screened U.S. citizens, Azure Government provides standards-compliant IaaS and PaaS solutions.	Virtustream Community Cloud is excluded as part of this contract submission.
Private Cloud	Amazon Private Cloud implementation of Virtual Private Cloud (VPC) is a commercial cloud computing service that provides users a virtual private cloud, by provisioning a logically isolated section of Amazon Web Services Enterprise customers are able to access the Amazon Elastic Compute Cloud (EC2) over an IPsec based virtual private network. Unlike traditional EC2 instances which are allocated internal and external IP numbers by Amazon, the customer can assign IP numbers of their choosing from one or more subnets. By giving the user the option of selecting which AWS resources are public facing and which are not, VPC provides much more granular control over security.	Microsoft's private cloud offering focuses on the application lifecycle combined with automation. Microsoft's private cloud software is part of the System Center 2012 R2 offering. System Center incorporates several products under one umbrella including Virtual Machine Manager, Data Protection Manager, Endpoint protection and Operations Manager. This, coupled with a straightforward ability to create self-service portals based on mature IIS features, helps with the installation process. Leveraging the. NET framework does allow for	Virtustream's xStream Cloud Management Platform provides a unified, cloud-agnostic control plane for a private cloud that integrates infrastructure orchestration, enterprise application automation and a suite of business intelligence and service management tools. This powerful combination allows for the delivery of enterprise IT as a service, with true consumption economics,

Deployment Method	Amazon AWS	Microsoft Azure	Virtustream
		additional extensions and troubleshooting.	application service-level agreements, and best-in-class security and compliance features.
Hybrid Cloud	Hybrid cloud architecture is the integration of on-premises resources with cloud resources. By working closely with enterprises, AWS has developed the industry's broadest set of hybrid capabilities across storage, networking, security, application deployment, and management tools to make it easy for you to integrate the cloud as a seamless and secure extension of your existing investments. AWS has also created strategic partnerships with longtime leaders in on-premises platform providers such as VMware, Intel, Microsoft, SAP, and others to allow you to run your existing enterprise applications on AWS with full support and high performance.	Azure approach to the hybrid cloud allows computing environments to combine a public cloud and a private cloud by allowing data and applications to be shared between them. Enables on-premises and cloud environments work consistently across your entire organization. Increase developer productivity with a common approach to building applications and the flexibility to deploy those apps in the cloud or on-premises with Azure Stack. Increase end-user productivity using Azure Active Directory for single sign-on to both cloud and on-premises applications.	Virtustream Hybrid Cloud is excluded as part of this contract submission.

Capgemini (PaaS)

Deployment Methods	Capgemini Enterprise iPaaS
Public Cloud	Capgemini Enterprise iPaaS is typically deployed as a Private Cloud deployment (as defined by NIST), however, can also be extended to a Hybrid Cloud deployment.
Hybrid Cloud	

Cloud Services Descriptions

The following sections describe Capgemini's Cloud Service offerings initially grouped by Cloud Service Provider.

BMC Remedy on Demand and ServiceNow (SaaS)

The Service Integrator, whether that is the Purchasing Entity or a Third Party, is the primary operational interface between the Purchasing Entity and its IT Service Providers and is accountable for service performance.

The Service Integrator requires a service management toolset that is specifically configured to enable the integrated processes and tools providing the Service Integration foundation.

- Integrated toolset aligned with and configured for the comprehensive ITILv3 process stack
- Configured with Service Level Agreement (SLA) requirements to track the performance of all service providers against SLA targets

- Populated with foundation data specific to the Purchasing Entity's IT organization and operating environment
- Service Management Tool provides a single source of the truth

Capgemini has extensive experience delivering BMC Remedy on Demand and the ServiceNow toolsets configured for ITIL Service Management. Through this experience, we created a standardized set of configuration rules, content, and methods to configure either tool, supporting the Service Management Tool choice of most Purchasing Entities.

Capgemini is offering the resell of both service management configured tools, along with value-added services to conduct Service Management assessments, Service Management implementations, and post-deployment Service Management process coaching and training.

Amazon, Microsoft, and Virtustream (IaaS, PaaS, SaaS)

Capgemini partnered with Amazon, Microsoft, and Virtustream to bring customers solutions for Public, Community, Private and Hybrid Cloud computing, storage and advanced application environments. Simply stated, cloud computing enables organizations to consume compute resources, networking, storage, or cloud applications through a utility model rather than investing, building and maintaining computing infrastructures on premise. It offers rapid application transformation platforms versus the traditional data center environments that are typically designed, built, and decommissioned over a much longer timeline.

Amazon AWS Cloud Services (IaaS, PaaS, SaaS)

Capgemini will assist Purchasing Entities with the planning, set up, configuration, and the automated provisioning of the IaaS services, workloads, and applications. In this section, Capgemini addresses what AWS services we offer that meet the services, definitions, and deployment models identified in the Scope of Services, in Attachment D Capgemini is proud to be AWS Premier Consulting Partner. Amazon Web Services (AWS) provides a broad set of products and services you can use as building blocks to run sophisticated and scalable applications. Running your applications in the AWS Cloud can help you move faster, operate more securely, and save substantial costs; all while benefitting from the scale and performance of the cloud.

AWS provides Infrastructure as a Service, provides access to networking features, virtual computers, and data storage space. The AWS Infrastructure as a Service implementation provides customers with the highest level of flexibility and management control over their IT resources and is similar to existing IT resources that many IT departments and developers are familiar with today. AWS most common IaaS services are Amazon Elastic Compute Cloud (EC2), Amazon Simple Storage Service and Elastic Load Balancing.

Microsoft Azure Cloud Services (IaaS, PaaS, SaaS)

As a Managed Gold Certified Partner with Microsoft, Capgemini offers our clients a unique set of advantages with a presence in over 35 countries, and on 5 continents supported by 24,000 Microsoft skilled professionals with Accelerated Delivery Centers and Rightshore® Delivery options.

We are currently one of a handful of Microsoft partners to have achieved gold competencies in Microsoft Cloud Center of Excellence in Gold Cloud Platform, Gold Cloud Productivity, Gold Customer Relationship Management and Gold Datacenter. Achieving these competencies demonstrates Capgemini's enterprise-grade expertise, commitment, and the ability to deliver Microsoft Cloud solutions.

Capgemini leverages and provides a cloud transformation to our public and private clients through the design and build services described in our value-added services section on top of the resale of the Azure services. We can help Purchasing Entities choose the right Azure services to support their ecosystem given our experience in implementing multiple cloud solutions from different vendors. Microsoft Azure is an ever-expanding set of cloud services to help your organization meet your

business challenges. It's the freedom to build, manage, and deploy applications on a massive, global network using your favorite tools and frameworks.

Microsoft Azure Infrastructure as a service (IaaS) is an instant computing infrastructure to provision and managed resources. Allowing to quickly scale up and down on demand.

Azure IaaS platform helps avoid the expense and complexity of buying and managing physical servers and other data center infrastructure. Each resource is offered as a separate service component, Azure manages the infrastructure, while the customer will purchase, install, configure, and manage their own software, operating systems, middleware, and applications. Azure IaaS provides a wide range of VM sizes, Containers, and images, enabling buyers to choose the best deployment options for their environment. Azure offers numerous networking services to build sophisticated network topologies and extend datacenters to the cloud. Azure delivers durable, highly available, and massively scalable storage options like Blob, Queue, File, and Disk—that keep pace with explosive data growth.

Virtustream (IaaS)

Capgemini has partnered with Virtustream to fill a void in the cloud transformation journey for most State and Local government customers. As most SLED clients want to reap the benefits of elastic public cloud services, most have stranded datacenter facilities and assets that remain to protect legacy applications or sensitive high risk data that cannot easily be ported to a public or Hybrid environments. Capgemini offers Private Hosting utilizing Virtustream's IaaS platform that is purpose built

Virtustream's Enterprise-Class Cloud is based on being an "Enterprise Community Cloud" for multiple organizations, Federal Entities or State Agencies that require a high level of performance, security, compliance, and operational excellence. Virtustream's offer presents a private, dedicated resource pool within the larger Virtustream enterprise cloud, equipped to meet the performance, security, and SLA requirements for Purchasing Entities' applications. The Enterprise-Class Cloud provides Purchasing Entities with the benefits of a dedicated infrastructure model without the premium that is associated with it. Virtustream's product, xStream, has been developed over the course of the past 6 years with core Enterprise principles in mind: security, scalability, guaranteed performance, continuous operations, and true consumption.

Our offer supports the evolution to Cloud Computing services and EMC's wholly-owned subsidiary Virtustream, will be providing the IaaS-based security, application layer performance guarantee and multi-tenant cloud. Our offer meets the NIST definition of cloud capabilities described in Attachment D of the RFP.

VIRTUSTREAM µVM TECHNOLOGY

A key advantage of our IaaS offer is the micro-VM (µVM) technology we developed. With this unique intellectual property (IP) Virtustream is able to provide IaaS services for thousands of users around the world across six global nodes. The majority of our customers run mission-critical production applications in addition to test, development and quality assurance environments.

Virtustream' s µVM technology is proven to scale up and down thus providing Purchasing Entities the flexibility we expect they will seek when considering cloud services under the NASPO Cloud contract.

The µVM is the DNA at the heart of the IaaS offering as it is the fundamental building block of our IaaS solution. The µVM combines compute and memory and makes sure that storage and network needs are fully provisioned. In addition the µVMs can be combined to form the optimal virtual machine (VM) for each individual application – enabling application resources to dynamically increase or decrease thus enabling Participating Entities to be assured SLA's are met without interruption. Virtustream's µVM technology works with industry standards and most common virtualization software (like VMware, Red Hat, IBM PowerVM, CentOS, and XenServer) therefore Purchasing Entities will have the highest degree of choice and flexibility as they consider cloud services.

The µVM brings significant benefits: enabling application level performance SLAs – which the average VM cannot normally meet. Virtustream's µVMs eliminate wasted headroom in fixed size VMs,

generating significant efficiency improvements (up to 40% beyond traditional virtualization) and since our solution only charges by the µVM you only pay for the resources you actually consume, not what you might need, this typically results in lower costs to the end users. Using µVMs also enables applications to be used across multiple hypervisors, across multiple clouds

The Virtustream tab in the "SK18008 Capgemini Detailed Product Offering" contains brief information about the Virtustream IaaS offering. Each service offering is listed with its Service ID, name, and a short description. A list of terms specific to how these services are named and described are also defined in the same document for your convenience.

Capgemini

Capgemini Enterprise iPaaS (integration Platform as a Service) is a key offering to provide cloud agnostic API and hybrid integration platform service that supports agile business processes, data and application integration.

Different combinations of cloud-based and on-premises applications can be integrated as part of an evolving hybrid cloud environment. This allows the Purchasing Entity to provide new business services (composite applications) and APIs enabling the Purchasing Entity to unlock the data held within their business, foster innovation and accelerate speed to market.

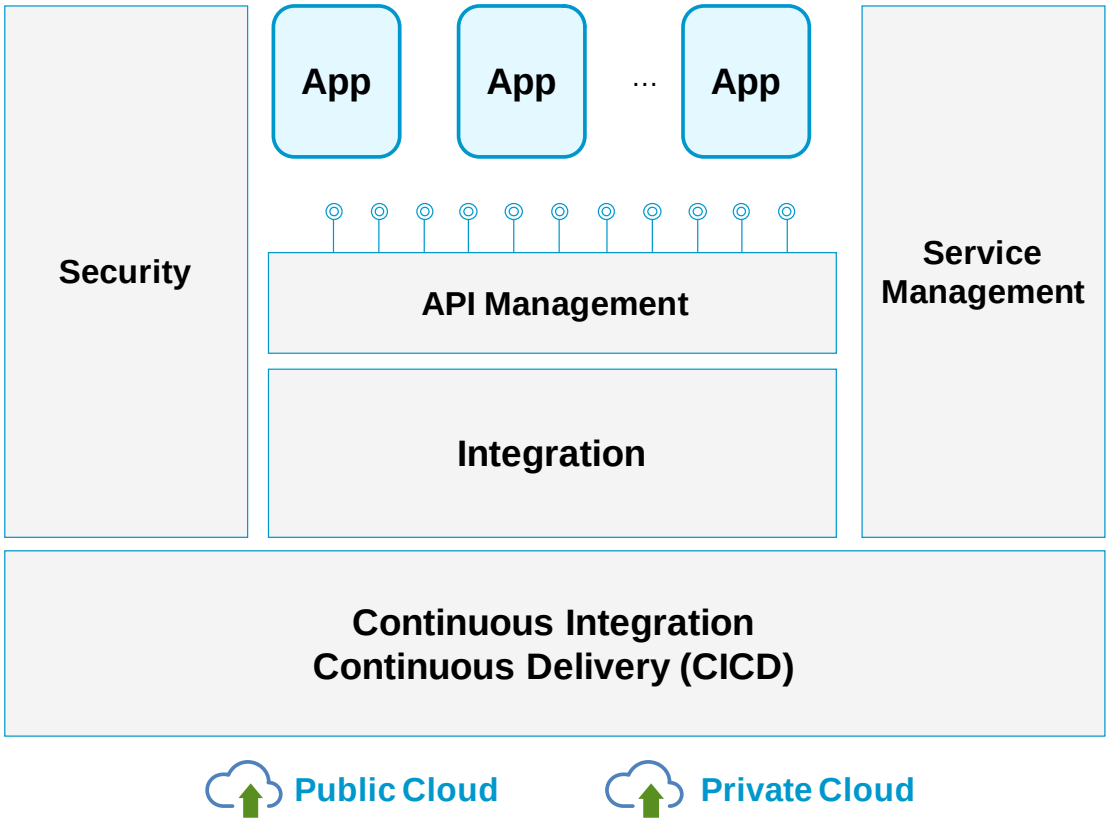

Figure 1: This diagram is for illustration only and does not represent any obligation or responsibility of Capgemini.

The **platform** is cloud-agnostic and can be ported across public and private cloud vendors. Capgemini Enterprise iPaaS gives the Purchasing Entity the flexibility to have their integration center of gravity on the Purchasing Entity's private cloud for reasons of security policy or regulation. Also, with a dedicated instance of the platform, the Purchasing Entity has the assurance of data isolation and provision of service levels. Built with open source products, the platform can be used to develop solutions ranging from pilots or prototypes to robust, secure, high-volume systems.

Capgemini Enterprise iPaaS is designed in a way that makes it simple and easy for the Purchasing Entity to buy, use and run this service.

- **Buy:** Capgemini Enterprise iPaaS can be bought on a pay-per-use model with a monthly subscription pricing and flexibility to avoid commercial lock-in. There are four predefined packages (large, medium, small, and micro) available. Once bought, the target for making the service available to the Purchasing Entity is 24 hours. Capgemini Enterprise iPaaS also allows the Purchasing Entity to scale between the package options on a daily pricing basis to efficiently manage any shifts in demand.
- **Use:** The Purchasing Entity can start using the Capgemini Enterprise iPaaS to securely publish APIs that unlock useful data hidden across the business, and to manage an API ecosystem. The Purchasing Entity can industrialize agile integrations with reusable integration patterns and pluggable components, deploy new integrations, and make use of user accelerators. The iPaaS can help continuous integration, to continuous delivery and continuous deployment.
- **Run:** The Purchasing Entity can provision additional platform environments on-demand, realize cost efficiencies from scaling environments daily based on demand, and use the Capgemini Enterprise iPaaS enterprise-grade service management to see how the Purchasing Entity's APIs and integrations are being used. The Supplier provides 24/7/365 support for production-ready environments and working hours for non-production environments.

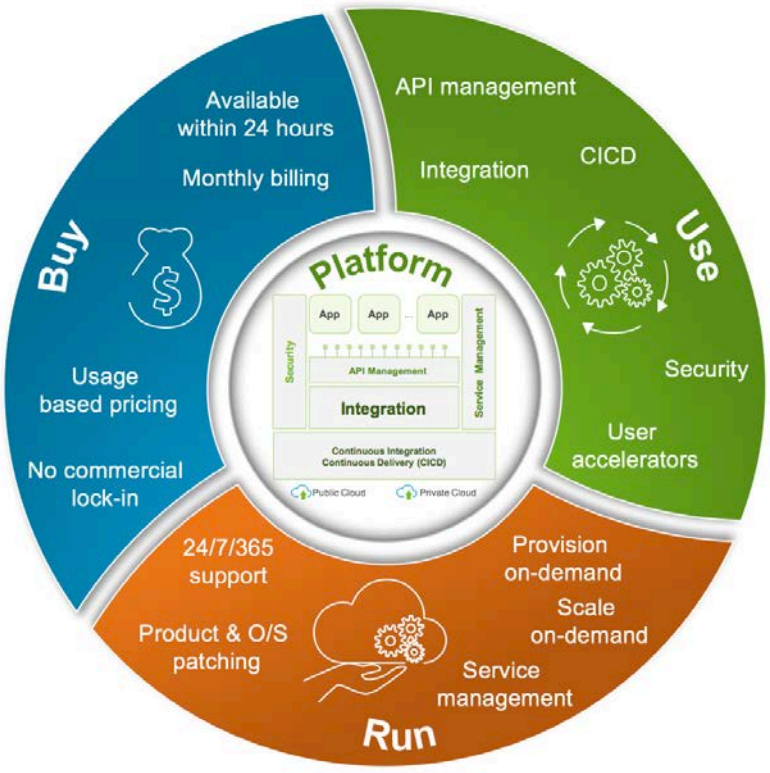

Figure 2: This diagram is for illustration only and does not represent any obligation or responsibility of Capgemini.

Business Need

Capgemini Enterprise iPaaS can help enterprise IT to deliver:

- Solutions to support digital transformations including cloud-first, mobile-first and Internet-of-Things strategies
- Faster integrations to connect an evolving hybrid IT landscape
- Reduced time to delivering composite applications and APIs
- Enhanced productivity for limited IT resources

- Improved ROI and optimized costs with a scalable consumption-based model
- Flexibility to avoid vendor lock-in, switch SaaS vendors and maintain reliable integration flows even when SaaS services are unavailable

These, in turn, can help the Purchasing Entity to meet the business demands for:

Increased speed to market

Continuous and rapid delivery of APIs and integration flows can reduce the time to deliver new services to a company's customers by swiftly reacting to the shift in customer demands and out-competing the competition.

Fostering innovation

Purchasing Entity can create strategic business value either by adopting new innovations through a SaaS product or by differentiating with innovative PaaS services connecting to the Purchasing Entity's core systems. Purchasing Entity can open up data hidden in business silos to develop innovative channels for growth and foster partner-led innovations.

Flexibility

Enable business to rapidly switch between cloud and SaaS vendors, scale to match market demands, and Purchasing Entity has the flexibility to move the Capgemini Enterprise iPaaS on the Purchasing Entity's private cloud if the on-premises infrastructure is the Purchasing Entity's center of gravity, and will remain so for reasons of security policy or regulation.

Cost efficiencies

Can reduce CapEx while optimizing OpEx through consumption-based pricing by adopting SaaS and PaaS models. Can realize cost savings from hosting development and test integrations in the public cloud, while moving production services into a private cloud.

Compliance with regulations

Can open up data from within the Purchasing Entity's business to comply with existing and emerging regulations. For example, in the banking industry, the use of APIs is fundamental to comply with the European Commission Revised Payment Services Directive (PSD2) and the UK Competition & Markets Authority Open Banking mandate.

Service Scope

The Capgemini Enterprise iPaaS is a black-box, fixed technology stack platform hosted in AWS US regions (East or West).

Use of the Capgemini Enterprise iPaaS is subject to the following terms:

- the open source software as described in Table C; and
- the Cloud Hosting Provider's terms available at AWS Reseller Customer License Terms via Amazons online portal.

Service Solution

The Enterprise iPaaS Platform is designed to support the development, execution, governance, and monitoring of application programming interfaces (APIs) and integration flows to connect both on-premises and cloud-based processes, services, applications, and data.

Enterprise iPaaS platform capabilities

The Enterprise iPaaS Platform can provide the capabilities set out in the below table.

Table A (iPaaS Capabilities)

Platform Component	Platform Capability	Definition
Shared Foundation Service	Central Function	Shared, centralized supporting services that enable the provisioning, management, and support of the Purchasing Entity Platform Instance by the iPaaS Support Team.
	Continuous Integration & Continuous Delivery	Enables the automated continuous integration and delivery of APIs and integrations.
API Management	API Design	APIs can be designed in the publishing user interface or by importing Swagger 2.0 definitions.
	API Publishing	APIs can be published to make them available for use. Access to specific APIs can be restricted to public and private groups of users.
	API Policies	Policies can be applied to throttle API requests through the API Gateway.
	Control Access and Enforce Security	APIs are secured using OAuth API access standard and support common OAuth grant profiles.
	API Discovery	Developers can browse, search and subscribe to published APIs that they are authorized to access via the API Store.
Integration	Integration Delivery	Integrations can be deployed and run.
Messaging	Message Queuing	Message queues to support integrations can be deployed and run.
Monitoring	Visual Analytics	Enables the display of data about use and performance of APIs and Integration deployed to the iPaaS.
User Accelerators	Worked Examples, Reusable Utilities	Developers will get access to worked examples of how to design, build and deploy APIs and Integrations to the iPaaS, along with reusable utilities that can be incorporated into the specific solutions being built by the Purchasing Entity.

Enterprise iPaaS platform sizes

Platform Instances can be provided in the sizes described in Table B (Platform Sizes) below.

The initial number(s) and size(s) of Platform Instances purchased by the Purchasing Entity as at the Effective Date shall be as described in Section Onboarding a New Purchasing Entity. Variations to the number(s) and size(s) of Platform Instances may be agreed from time to time in accordance with Section Platform Requests.

The “Estimated Capacity” figures are provided as an indication of the recommended platform size. No warranty is given by the Supplier that the size of platform ordered by the Purchasing Entity will meet the Purchasing Entity’s processing requirements and actual performance figures will depend on the nature of the Purchasing Entity specific APIs and Integrations.

Table B (iPaaS Platform Size)

	API Manager			Integrations				
Platform Size	API Gateway Cores	Estimated Capacity API Calls per second (https)	Included outbound data transfer allowance per month	Integration Cores	Estimated Capacity Integration Flows	Estimated Capacity Connections	Message Queuing Cores	Message Queueing Storage Allowance
Micro	2	25	0.25TB	2	5	15	2	30 GB
Small	4	100	1TB	6	20	60	6	90 GB
Medium	8	200	2TB	12	40	120	10	180 GB
Large	12	300	3TB	18	60	180	14	270 GB

Purchasing Entity personal data

Subject matter: Provision of the Enterprise iPaaS Platform.

Duration of the processing: The duration of the Subscription Term.

Purpose of the processing: The provision of the Enterprise iPaaS Platform and the performance by the Supplier of its obligations under the Agreement.

Nature of the processing: The Enterprise iPaaS Platform can process the following personal data:

- User information for subscribers to the API store;
- Login credentials for the iPaaS software listed in Table C;
- Any personal data within the APIs and integrations built by the Purchasing Entity. The Enterprise iPaaS is not the master source of the integration data.

Type of Purchasing Entity Personal Data: The personal data uploaded to the Enterprise iPaaS Platform or otherwise processed by the Supplier on the Purchasing Entity's behalf.

Categories of data subjects: The data subjects may include the Purchasing Entity's employees, customers, and end-users.

Open source software

Where indicated with a "Y" in Table C, the Purchasing Entity can be provided with user accounts for each end-user.

Table C (OSS Licenses)

No.	License	Software
1	Apache 2.0 (http://www.apache.org/licenses/LICENSE-2.0.html)	Grafana (Y)
		Kibana (Y)
		WSO2 API Manager (Y)
		WSO2 Carbon
		Subversion
		GoCD (Y)

No.	License	Software
		Rundeck (Y)
		Graphite / Carbon
		Whisper
		ElasticSearch
		Logstash
2	Mozilla 1.1 (https://www.rabbitmq.com/mpl.html)	RabbitMQ (Y)
3	Mozilla Public License, version 2.0 (https://www.mozilla.org/en-US/MPL/2.0/)	Nomad
		Consul Agent
		Consul Template
		Envconsul
		Consul Server
4	Oracle (http://www.oracle.com/technetwork/java/javase/terms/license/index.html)	Java SE
5	MIT (https://opensource.org/licenses/MIT)	Nomad UI (Y)
		Sensu Server
		Fabio
		Uchiwa
		Collected Plug-In
6	OpenVPN License (https://openvpn.net/index.php/license.html)	OpenVPN AS (Y)
7	Creative Commons (https://creativecommons.org/licenses/by-nc-nd/3.0/us/)	Nexus (Y)
8	Creative Commons Attribution 3.0 (https://www.centos.org/legal/trademarks/#license-and-attribution)	CentOS
9	BSD (https://redis.io/topics/license)	Redis
		Collected Plug-In
10	GNU GPL v2 (https://www.gnu.org/licenses/old-licenses/gpl-2.0.en.html)	Collected Plug-In
		HA Proxy
11	PostgreSQL License (https://www.postgresql.org/about/licence/)	Postgres SQL DB
12	CDDL-1.0	OpenDJ

No.	License	Software
	(https://opensource.org/licenses/CDDL-1.0)	
13	Software Specific License (https://github.com/tests-always-included/mo/blob/master/LICENSE.md)	Mustache
14	Software Specific License (http://nginx.org/LICENSE)	NginX

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.

Capgemini does not plan to use Subcontractors. Capgemini will provide iPaaS cloud service offering directly, and value-added professional services without the use of a subcontractor.

We intend to use our strategic and authorized partners Amazon, Microsoft, Virtustream, ServiceNow, and BMC to meet the NIST compliant cloud solutions detailed throughout this proposal.

Capgemini will resell our Partner's cloud services using our agreements with our Partners for the delivery of SaaS, PaaS and IaaS offerings and will pass through terms to Purchasing Entities. We break down the Attachment D criteria in section 8.1.3 in much greater detail.

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.

Capgemini does not plan to use Subcontractors. Capgemini will provide iPaaS cloud service offering directly, and value-added professional services without the use of a subcontractor.

We intend to use our strategic and authorized partners Amazon, Microsoft, Virtustream, ServiceNow, and BMC to meet the NIST compliant cloud solutions detailed throughout this proposal.

Capgemini will resell our Partner's cloud services using our agreements with our Partners for the delivery of SaaS, PaaS and IaaS offerings and will pass through terms to Purchasing Entities. We break down the Attachment D criteria in section 8.1.3 in much greater detail.

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 the 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.

Capgemini does not plan to use Subcontractors. Capgemini will provide iPaaS cloud service offering directly, and value-added professional services without the use of a subcontractor.

We intend to use our strategic and authorized partners Amazon, Microsoft, Virtustream, ServiceNow, and BMC to meet the NIST compliant cloud solutions detailed throughout this proposal.

Capgemini will resell our Partner's cloud services using our agreements with our Partners for the delivery of SaaS, PaaS and IaaS offerings and will pass through terms to Purchasing Entities. We breakdown the Attachment D criteria in section 8.1.3 in much greater detail.

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.

Capgemini and its Partners, collectively "Offeror" shall inform the Purchasing Entity of any security incident or data breach within the possession and control of Offeror and related to the service provided under the Master Agreement, Participating Addendum, or SLA. Such notice shall include, to the best of Offeror'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. Capgemini's cloud services Partners have required mechanisms in place to identify, respond and report data breaches within agreed timelines with Customers inline to regulatory compliance requirements applicable to Government organizations. The offeror shall support defined and agreed actions with Purchasing Entities during and after a data breach.

Offeror's Contract Manager will be involved with Purchasing Entities upon confirmed notification of data breach.

- **Incident Response**: Offeror will 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 will be handled on an urgent as-needed basis, as part of Offeror's communication and mitigation processes as mutually agreed, defined by law or contained in the Master Agreement, Participating Addendum, or SLA.
- **Security Incident Reporting**: Unless otherwise stipulated, the Offeror will immediately report a security incident related to its service under the Master Agreement, Participating Addendum, or SLA to the appropriate Purchasing Entity.
- **Breach Reporting**: After identification of confirmed data breach that affects the security of any Purchasing Entity data that is subject to applicable data breach notification law, the Offeror will (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.

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.

Capgemini and its Partners, collectively "Offeror" will not push adware, software, or marketing that is not explicitly authorized by the participating Purchasing Entity or the Master Agreement. The Information Technology (IT) assets to be used by Offeror's staff are security hardened, have stringent administrative policies enforced, and have adequate technical security controls (E.g. Anti-virus, anti-malware, web security, encryption, VPN) installed to detect and protect from adware, malware, and unauthorized malicious behavior.

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

Capgemini provides value-added consulting and professional services inclusive of advisory, implementation services and the tools to stand up the necessary testing, development, and staging environments in either AWS/Azure public or Virtustream private clouds. The development, testing, and staging environments are provisioned in accordance with the customer production environment.

For Capgemini iPaaS, physically and logically separated non-production environments are available based on client requirements and will be provisioned for user test/staging to assist in cross-system and application integration.

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

Capgemini's computer applications and websites are accessible to people with disabilities and comply with applicable policies and laws.

8.3.5 Offeror must describe whether its applications and content delivered through Web browsers are be accessible using currently released versions of multiple browser platforms (such as Internet Explorer, Firefox, Chrome, and Safari) at a minimum.

Capgemini certifies content delivered through Web browsers will be accessible using currently released versions of multiple browser platforms, like Internet Explorer, Firefox, Chrome, and Safari.

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.

Prior to the execution of Service Level Agreement, Capgemini's and its Partner's, collectively "Offeror's", Contract Manager will schedule a meeting with Purchasing Entity including Subject Matter Experts (SMEs) from Offeror and Purchasing Entities to review and 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 specific compliance obligations. The offeror will be responsible to provide appropriate security controls and assurance required to process and store sensitive and personal information.

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.

The Capgemini value-added professional services will leverage our long-established DELIVER™ methodology. This framework will make it possible to identify and manage the complexities of a successful implementation while transforming into a Cloud platform. This includes structured technical discovery, knowledge transfer, comprehensive migration planning and well documented using typical Microsoft Project Plans, deliverables charts, and a structured transition plan with a communication plan that drives all activities.

The standard phases are:

Advisory & Analysis

- Strategy, business case, roadmap & recommendations
- Baseline inventory & scope

Assessment & Design

- Assess landscape & services, sentencing, migration blueprints
- Scope definition, budgets & program plans

Planning & Management

- Planning of change batches and migration waves
- Stakeholder alignment & management of all migrations

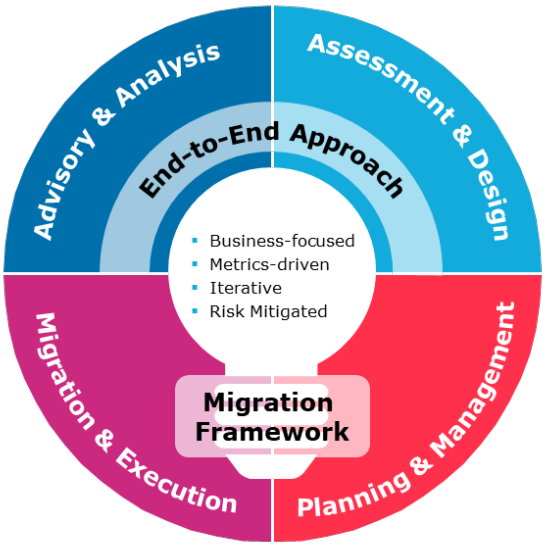

Migration & Execution Framework

- Fully tested implementation in Production for Cloud Environments
- Framework-based with a technical approach and delivery

Program Governance and Management

Capgemini is recognized in both the private and public sector for experience in the planning, structure, processes, and tools required to facilitate collaborative, flexible, and effective management for Purchasing Entities Cloud Transformation projects including:

- A project management approach founded on 40 years of experience that can reduce risk and promote support through leading methods
- A holistic change management approach that encompasses the phases of the project lifecycle to enable system acceptance and knowledge transfer
- A collaborative approach that fosters positive and productive working relationships and creates advocates for the project within the Purchasing Entities' staff
- Multi-year contracts with the State of Texas and State of Georgia delivering complex program governance and management from 2011 to present

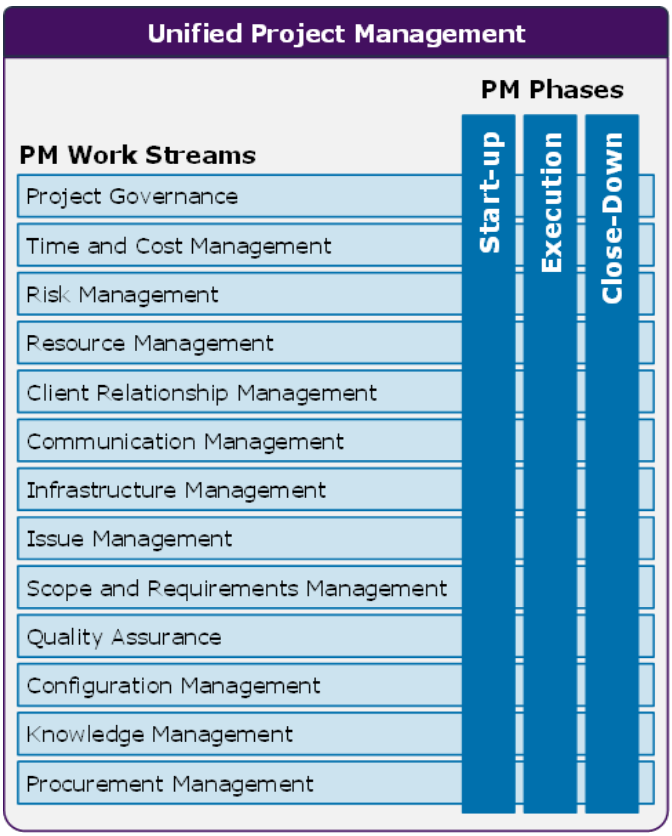

Capgemini views project management as a collaborative delivery function and will work closely with the Purchasing Entities teams throughout this the life of the engagements. Capgemini will use Unified Project Management (UPM) framework, illustrated on the right.

UPM is based on accepted project management standards (e.g., PMI PMBOK), IT management (e.g., ITIL®), and process improvement standards (e.g., CMMI and ISO). It has been used and enhanced over thousands of client engagements during our more than 40-year history. UPM helps drive quality and reduce program risk to promote the accomplishment of customers goals and objectives, delivering the oversight necessary to rapidly staff qualified resources and confirm that tasks are performed efficiently, accurately, on time, and in compliance with the requirements.

The UPM method is organized along streams or collections of common activities or themes across the life of the project.

Cloud Migration Framework

In the Cloud Advisory value-added professional services, Capgemini has developed a unique migration framework and tools to assist customers in migrating infrastructure and applications to the cloud. Capgemini implements a mature process reviewing current infrastructure and applications to define the decision criteria necessary to determine if a workload or application is compatible with public Cloud platforms and therefore viable to migrate. During this step, we determine if Infrastructure as a Service (IaaS), Platform as a service (PaaS) or Software as a Service (SaaS) patterns present suitable options.

The key steps include:

Run Capgemini analysis tools for identifying applications and dependencies; and rightsizing.

- Identify apps that can be retired and/or consolidated
- Determine appropriate migration pattern

The Cloud Migration timeline depicted in the diagram below is an example of a project timeline with actionable deliverables and outcomes.

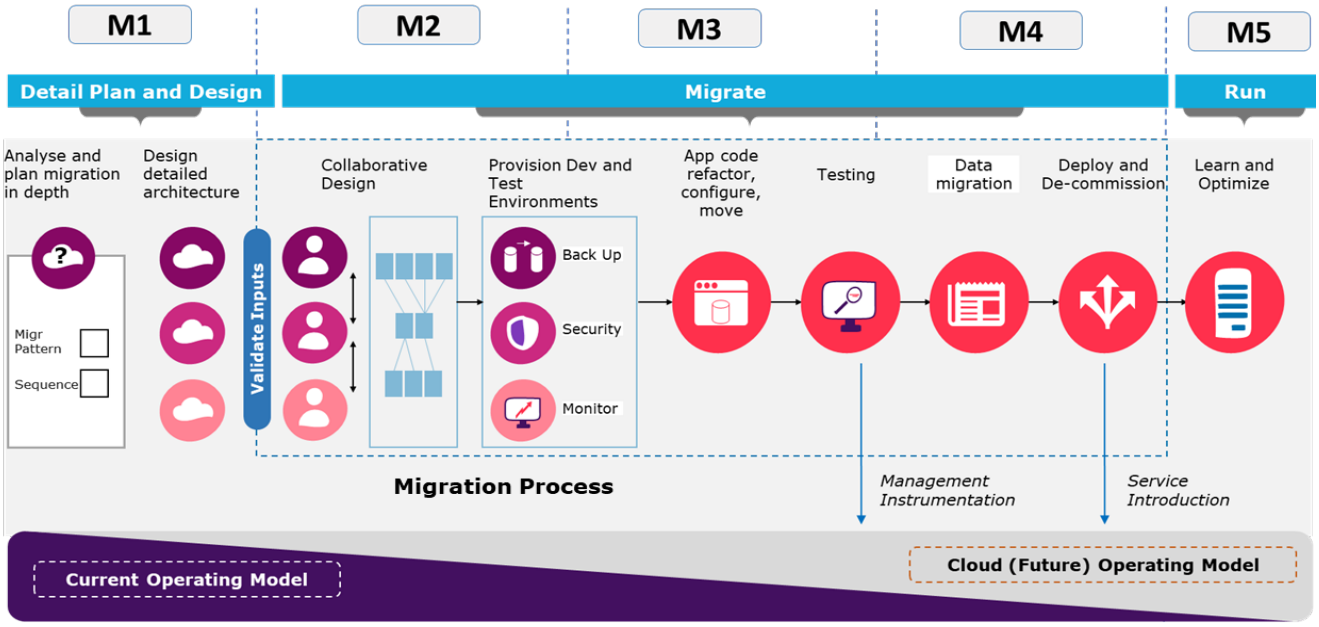

Figure 3: The Cloud Migration timeline

The timeline to implement a cloud solution will depend on the size, complexity of the customer environment, in addition factoring the time for development, testing, staging and migrating to production the cloud solutions to the new cloud environment.

8.3.8 The State of Utah expects Offeror to update the services periodically as technology changes. Offer must describe:

- How Offeror's services during Service Line Additions and Updates pursuant to section 2.12 will continue to meet the requirements outlined therein.
- How Offeror will maintain discounts at the levels set forth in the contract.
- How Offeror will report to the Purchasing Entities, as needed, regarding changes in technology and make recommendations for service updates.
- How Offeror will provide transition support to any Purchasing Entity whose operations may be negatively impacted by the service change.

AWS and Azure (IaaS, PaaS, SaaS)

Service Line Additions and Updates

Capgemini has a strategic relationship with AWS and Azure, based on that relationship, Capgemini obtains timely information on updates that will be implemented by the cloud providers in their environments. Capgemini has access to the upcoming roadmaps from the cloud providers, providing an early view of upcoming updates and enhancements to the cloud providers platforms. Capgemini evaluates the planned updates that can impact customers, establishing the necessary Change Management processes and proactive communication channels to inform customers of any updates.

Maintaining discounts

Established discounts for AWS and Azure are managed and audited by the Contract Manager, so that the discounts are maintained at the levels set forth in the contract.

Communicating needed changes in technology and making recommendations for service updates

Capgemini's Change Management processes focus on assisting the customer's IT organization in adopting a new way of working, along with gaining acceptance when changes occur across the organization. Change Management contributes to success by maximizing the traction and trajectory of the customer's cloud transformation projects and associated engagements.

Depending on the level of change, a fit for purpose Change Management plan will be designed and executed leveraging our leading practices and processes that take a customer-centric approach to reduce risk, provide for properly controlled releases to the cloud production environment.

Providing transition support when operations may be negatively impacted by service changes

Capgemini through its Change Management process can orchestrate failover plans in case a change negatively impacts the cloud production environment. In such event, proper communication channels are established to inform the customer of the potential downtime and restoration of services. A full root cause analysis can be performed to determine the cause of the error prior to making any future changes to the cloud production environment.

Virtustream (IaaS)

Service Line Additions and Updates

Capgemini has a strategic relationship with Virtustream and receives obtains timely information on service updates, enhancements, compliance certifications, and new capabilities implemented by the cloud provider to their environments. Capgemini's strong relationship with Virtustream provides insight into solution/service roadmaps, technology enhancements, and service level improvements allowing to better access potential impacts to our customer base. Capgemini evaluates the planned updates, establishing the necessary Change Management processes and proactive communication channels to inform customers of any modifications to their subscribed services.

Maintaining discounts

Established discounts for Virtustream are managed and audited by the Contract Manager so that the discounts are maintained at the levels set forth in the contract.

Communicating needed changes in technology and making recommendations for service updates

Capgemini's Change Management processes focus on assisting the customer's IT organization in adopting a new way of working, along with gaining acceptance when changes occur across the organization. Change Management contributes to success by maximizing the traction and trajectory of the customer's cloud transformation projects and associated engagements.

Depending on the level of change, a fit for purpose Change Management plan will be designed and executed leveraging our leading practices and processes that take a customer-centric approach to reduce risk, provide proper controlled releases to the cloud production environment.

Providing transition support when operations may be negatively impacted by service changes

Capgemini through its Change Management process can orchestrate failback plans in case a change negatively impacts the cloud production environment. In such event, proper communication channels are established to inform the customer of the potential downtime and restoration of services. A full root cause analysis can be performed to determine the cause of the error prior to making any future changes to the cloud production environment.

BMC Remedy on Demand (SaaS)

Service Line Additions and Updates

Included in the Remedy on Demand subscription services are periodic upgrades of the software version to the next major version. This usually takes place within a few months of the next newest release of the software. In between version upgrades, patches and hot packs are applied as required to maintain service levels and functionality.

Maintaining discounts

Discount levels are managed and audited by the Contract Manager so that the discounts are maintained at the levels set forth in the contract.

Communicating needed changes in technology and making recommendations for service updates

BMC assigns their own Account Executive and a Service Delivery Manager to each Purchasing Entity to communicate needed changes and for making recommendations.

Providing transition support when operations may be negatively impacted by service changes

BMC has developed a near zero downtime upgrade process for upgrading the underlying platform of Remedy ITSM so that upgrades to Remedy ITSM Version 9.1.4 and higher to minimize the impact of upgrades to operations. Even with a near-zero downtime upgrade event, there is a substantial amount of planning, preparing, functionality reviews, and testing, to conduct in addition to whatever development will go into the next version.

Further, BMC has developed a Zero Downtime Maintenance with Availability Zone architecture on its Linux-based Remedy on Demand infrastructure.

For these activities, Capgemini can easily assist in whatever capacity or level to assist with planning and executing an upgrade or transition into the Remedy on Demand.

ServiceNow (SaaS)

Service Line Additions and Updates

ServiceNow continuously updates their product functionality that is batched into releases that are roughly 7-10 months apart. ServiceNow will apply that functionality to their Purchasing Entity instances on a business as usual basis.

However, most Purchasing Entity will choose to upgrade their instances to the next release on their own schedule to provide an opportunity to deploy new functionality, and systematically test their customizations allowing an opportunity to make required updates.

Those Purchasing Entities requiring additional resources to design, develop, test, and deploy functionality or just to apply the release to the existing functionality can request assistance from Capgemini at published rate card rates.

Maintaining discounts

Discount levels are managed and audited by the Contract Manager so that the discounts are maintained at the levels set forth in the contract.

Communicating needed changes in technology and making recommendations for service updates

ServiceNow assigns an Account Exec and a Service Delivery Manager to each Purchasing Entity to communicate needed changes and for making recommendations.

Providing transition support when operations may be negatively impacted by service changes

The goal of ServiceNow's Nonstop Cloud is to be always operational for their "Clients" with no extended upgrade or maintenance windows, no single points of failure, and a focus on near-perfect availability with a Purchasing Entity accessible Real Availability Dashboard that shows availability of all instances running in the cloud. There is no downtime necessary for upgrades.

Capgemini Enterprise iPaaS (PaaS)

Maintaining discounts

The Capgemini Enterprise iPaaS is a Capgemini Service Offering and the discounts are managed and audited by the Contract Manager so that the discounts are maintained at the levels set forth in the contract.

Communicating needed changes in technology and making recommendations for service updates

The Capgemini Enterprise iPaaS includes a full managed service which incorporates informing our clients of any technical changes and agreeing on the application of these changes within the definition of the service.

Providing transition support when operations may be negatively impacted by service changes

All service changes are agreed with our clients and include a regression plan.

8.4 (E) CUSTOMER SERVICE

8.4.1 Offeror must describe how it will 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).

Quality Assurance

Capgemini's plan and approach to quality is based on a solid framework that promotes consistency and flexibility to provide alignment to customer's requirements. This ultimately provides the benefit of improved quality and predictability while providing a value. We have a robust quality and compliance group that covers functional areas such as Quality Management, Quality Management System (QMS) Administration, Training Administration, Internal Audit, Corrective Action and Preventive Action (CAPA), System Engineering Process Group, Process and Product Quality Assurance (PPQA), and SSAE18 Audit Reporting. Each of these areas contributes to providing product and process quality.

Capgemini has experience providing support for complex application and cloud transformations. These projects require in-depth quality assurance and may include the below model which would be agreed to in a SOW leveraging our value-added professional services:

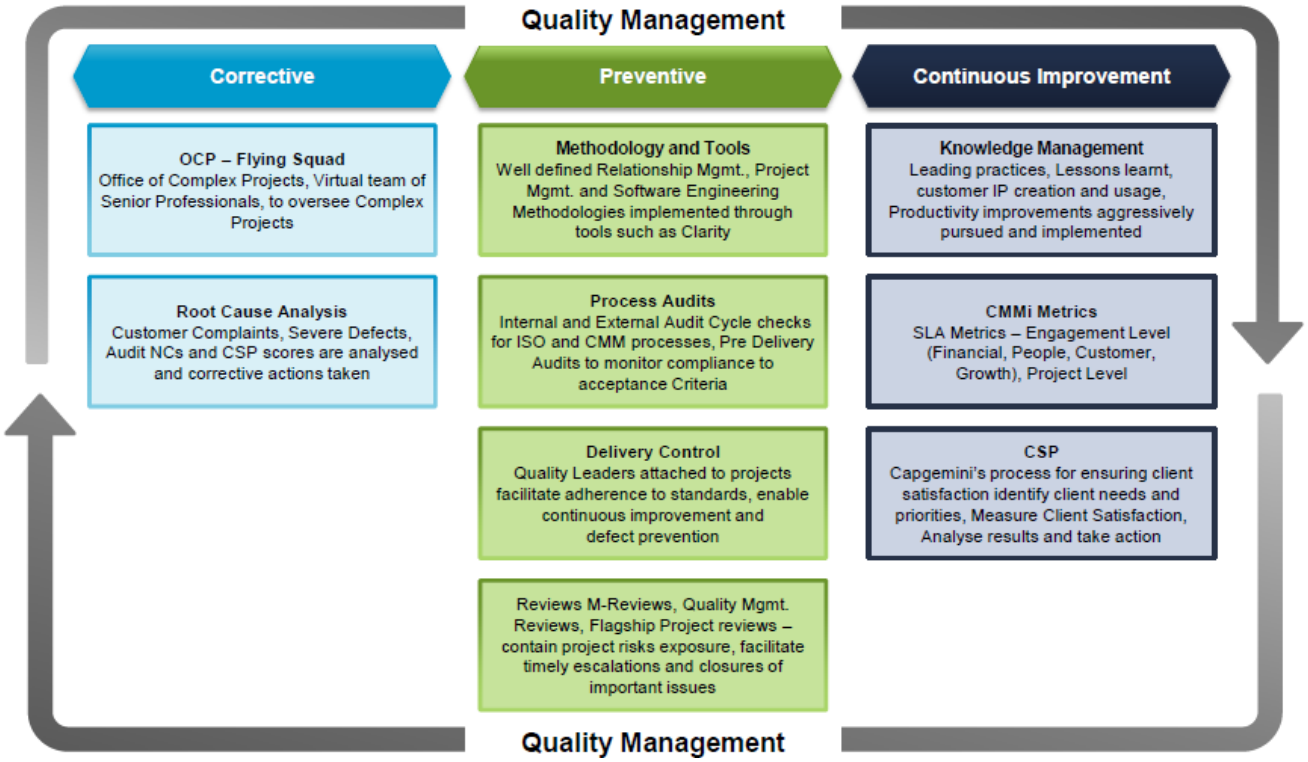

Figure 4: Quality Management

At a project level, Capgemini has a Quality Advisory/Risk and Delivery Management Process in place to assist in identifying, evaluating and monitoring risk/issues and quality of service so that we can provide our clients with quality deliverables. This helps the project in keeping a close eye on quality processes, the threshold on KPIs, compliances to Entry and Exit criteria, and most importantly defect monitoring, analysis, and corrective measures.

Capgemini is committed to delivering a high-quality design, implementations-transformation's to the Purchasing Entities. Our skilled Application, Cloud, Network and Security architects understand the complexity of moving three tiered and native cloud applications to Public, private and Hybrid cloud infrastructures. Our ultimate measure of quality is a happy customer which we measure with our Customer Satisfaction process, OTACE (On Time At Customer Expectations) and DELIVER™ methodology.

Capgemini's approach takes into consideration that there is a hierarchy of information needed to support the project: Enterprise, Operational and Individual.

- **Enterprise-level** themes are strategic or organizational in nature directed to most or all project stakeholders. These messages are managed through the project's Change Management team
- **Operation-level** messages are strategic and tactical in nature and are directed to employees by business function, process, or location. These messages are coordinated with project's Change Management team and executed through the Project Leadership and the Process Managers
- **Individual-level** information is tactical or transitional in nature and directed to individuals or small groups. These messages are coordinated with change network activities and are executed at the local level through change owners and leaders

Escalation Process

Below is an example of Capgemini's escalation process. We will work with each Purchasing Entity and tailor the escalation process to meet your business needs and requirements.

Escalation procedures have been developed so that the appropriate management attention is devoted to outstanding problems consistent with their impact. Each outstanding problem will be categorized with a priority level and assigned the appropriate level of resources consistent with its impact. This escalation procedure increases the level of resources as necessary to resolve problems effectively. Capgemini will clarify the problem situation and communicate action plans to the Purchasing Entity within a time frame appropriate to the severity of the pending problem.

Capgemini will take a pro-active approach to notify both Capgemini management and Purchasing Entity according to the below:

Escalation Criteria and Notification

Priority 1 Escalation Criteria	Person Notified
Priority 1 problem beyond defined problem resolution time.	Capgemini Contract Manager Purchasing Entity point of contact
Priority 2 Escalation Criteria	**Person Notified**
Priority 2 problem outstanding past defined problem resolution time.	Capgemini Contract Manager Purchasing Entity point of contact
Priority 3 Escalation Criteria	**Person Notified**
Priority 3 problem outstanding past defined problem resolution time.	Capgemini Contract Manager Purchasing Entity point of contact

The table below defines the priorities for Problems:

Problem Priority Definition

Priority	Definition
Priority 1	Priority 1 Critical Business Impact– System down or immediate work stoppage of a critical business service that threatens current and future productivity.
Priority 2	Priority 2 Significant Business Impact– Problem where the system or business service is proceeding but in a seriously impaired or in a restricted fashion and no acceptable workaround is possible.
Priority 3	Priority 3 Some Business Impact – Problem for which the impact is an inconvenience, which may require a workaround to restore functionality and productivity is not seriously impaired.
Priority 4	All other problems or requests.

Service Level Agreement (SLA)

Capgemini will align and include the performance metrics in the overall Service Level Management process. This will provide the insight into the Service Level Agreements that will be agreed upon between the Purchasing Entity and Capgemini. The Service Levels as well will be aligned with Capgemini's overall engagement metrics and reporting and be flowed down from the Cloud Service Providers.

These procedures include the approach that will be used to receive and document reporting requirements, assure the setup of necessary and appropriate supporting reporting environments, the ongoing monitoring and collection of report metrics across all levels of the organization, and processing and distribution of reports.

Service level management reporting allows the Purchasing Entity to be fully aware of Cloud Service Provider service commitments and performance at any given time. In turn, this allows the Purchasing Entity to meet its commitments to its internal business users and external citizens. Capgemini's Service Management provides a set of specialized organizational capabilities for providing value to the Purchasing Entity in the form of services. Capgemini's Service Management encapsulates the following:

- ITIL v3 processes
- ITIL v3 Incident, Problem, Change and Request Management
- A collaborative approach whereby Capgemini will work with the Purchasing Entity to agree what needs to be done and achieved and not just "Do It"
- Control the service based on a mutually agreed governance framework and a process quality driven approach.
- Event management, Service Level Management, Availability management, Knowledge management, Service Level Management

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

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

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

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

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

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

1. Capgemini will maintain a Lead Representative that is available to each Purchasing Entity that executes a Participating Addendum.
2. Upon execution of a Participating Addendum, Capgemini will have a Customer Service Representative, the Capgemini Contract Manager, available by phone or email 7 am to 6 pm Monday through Sunday.
3. The Customer Representative, Capgemini Contract Manager, will be able to respond to inquiries within 1 business day.
4. Capgemini will provide experienced design services for all the applicable categories listed in our product catalog.
5. Capgemini will provide Installation Services for the applicable categories listed in our product catalog.

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.

Capgemini follows “Security & Privacy by design” approach and works closely with our customers to identify and implement additional security controls required to store and process sensitive and personal information starting from data creation to its secure disposal inline to NIST guidelines by creating a SOW leveraging our value-added professional services.

Capgemini’s Partners included in our proposal, have implemented a broad range of data security controls and are certified, accredited to multiple industries and government standards and comply with security compliance regulations/requirements to hold, protect and disposal of data.

AWS and Azure (IaaS, PaaS, SaaS)

Data security measure provided by Amazon Web Services (AWS): AWS infrastructure has been designed to provide the highest availability while putting stringent safeguards in place regarding data security, privacy, and segregation. When deploying systems in the AWS cloud, AWS and its customers share the security responsibilities. Amazon follows “Security by design” approach to implement security and privacy controls in AWS. The centralized access, security, visibility and transparency of operating with the AWS cloud provides for increased capability for designing end-to-end security for all services, data, and applications in AWS. AWS customers retain controls and ownership of their data, and all data stored by AWS on behalf of customers have strong tenant isolation security and control capabilities.

Multiple options for data security, ranging from automated AWS encryption solutions to manual, client-side encryption options are available in AWS. Capgemini will work closely with Purchasing Entities to choose the right solutions based on the requirements and sensitivity of data to be protected, based on which AWS cloud services will be used to implement security controls to protect the data in AWS. For secure data storage, Encryption at rest is vital and is used for regulatory compliance so that sensitive data saved on disks is not readable by any unauthorized user or application. To comply with security regulations such as PCI DSS, HIPAA, data-at-rest options and key management solutions can be used leveraging AWS. Customers can also encrypt Amazon EBS volumes and configure Amazon S3 buckets for server-side encryption (SSE) using AES-256 encryption. Additionally, Amazon RDS supports Transparent Data Encryption (TDE).

AWS follows standard NIST guidelines for secure data disposal addressing the data remanence issues so that deleted data is not recoverable with reasonable efforts. More details about data security provisions are provided in Attachment A – Amazon Web Services.

Data security measures provided by Microsoft Azure: Microsoft Azure provide tools and security controls to safeguard data according to Customer's security and compliance needs. Encryption at rest is a common security requirement, in Azure, Customers have capabilities to achieve Encryption at rest without having the cost of implementation and management and the risk of a custom key management solution. Azure supports three scenarios for server-side encryption as below:

- Server-side encryption using Service Managed keys;
- Server-side encryption using customer-managed keys in the Azure key vault;
- Server-side encryption using computer-managed keys on computer-controlled hardware.

"Client-side" encryption can be performed by the service application in Azure, or by an application running in the customer data center. When leveraging client-side encryption model, the Azure Resource Provider receives an encrypted blob of data without the ability to decrypt the data in any way or have access to the encryption keys. In this model, the key management is done by the calling service/application and is opaque to the Azure service. Microsoft uses multiple encryption methods, protocols, and algorithms across its products and services to help provide a secure path for data to travel through the infrastructure, and to help protect the confidentiality of data that is stored within the infrastructure. Microsoft uses some of the strongest, most secure encryption protocols in the industry to provide a barrier against unauthorized access to your data. Proper key management is an essential element in encryption leading practices, and Microsoft helps provide that encryption keys are properly secured. Customers can leverage Microsoft Azure's data security capabilities such as:

- Transport Layer Security/Secure Sockets Layer (TLS/SSL), which uses symmetric cryptography based on a shared secret to encrypt communications as they travel over the network.
- Internet Protocol Security (IPsec), an industry-standard set of protocols used to provide authentication, integrity, and confidentiality of data at the IP packet level as it's transferred across the network.
- Advanced Encryption Standard (AES)-256, the National Institute of Standards and Technology (NIST) specification for a symmetric key data encryption that was adopted by the US government to replace Data Encryption Standard (DES) and RSA 2048 public key encryption technology.
- BitLocker encryption that uses AES to encrypt entire volumes on Windows server and client machines, which can be used to encrypt Hyper-V virtual machines when you add a virtual Trusted Platform Module (TPM). BitLocker also encrypts Shielded VMs in Windows Server 2016, so that fabric administrators can't access the information inside the virtual machine. The Shielded VMs solution includes the new Host Guardian Service feature, which is used for virtualization host attestation and encryption key release.
- Microsoft Azure Storage Service Encryption encrypts data at rest when it's stored in Azure Blob storage. Azure Disk Encryption encrypts your Windows and Linux infrastructure as a service (IaaS) virtual machine disks by using the BitLocker feature of Windows and the DM-Crypt feature of Linux to provide volume encryption for the operating system and the data disk.
- Transparent Data Encryption (TDE) encrypts data at rest when it's stored in an Azure SQL database.
- Azure Key Vault helps you easily and cost-effectively manage and maintain control of the encryption keys used by cloud apps and services via a cloud-based hardware security module (HSM).

More details about data security provisions available in Microsoft Azure are provided in Attachment B – Microsoft Azure.

Virtustream (IaaS)

Data security measures provided by Virtustream Private cloud: Virtustream's Private cloud is architected to the highest security and compliance standards and can significantly contribute to an enabler in achieving and maintaining compliance with FedRAMP, FISMA, SSAE16/SOC2/ISAE3402, PCI DSS 3.2, ISO 27001:2013, ISO 9001:2015, ISO 22301:2012, HIPAA/HITECH/HITRUST and other certification and compliance frameworks. In each of security and compliance the Virtustream cloud solutions, both Virtustream Private cloud, and xStream software follow the core tenants of cloud data security:

- Compliance
- Trust
  - Intel Trusted Execution Technology (TXT)
  - Two-factor authentication
  - Encryption
- Visibility
  - Auditing
  - Alerting
  - SIEM
- Control
  - Role-based access control
  - Virtual firewall

Data security controls for its entire lifecycle are implemented in a Virtustream Private Cloud platform. Data at rest and data in transit encryption is supported by Virtustream Cloud services. Various third-party products are used to secure data at rest and in motion as well as to authenticate the components of Virtustream cloud technologies stack. Utilizing FIPS-compliant cryptographic technology, Virtustream can support all major encryption requirements for file systems, database, and network transport protection. High-efficiency encryption technologies are deployed to protect the entire data lifecycle and are utilized throughout the Virtustream cloud environment to secure: data at rest including the entire virtual machine and its related data storage, transactional databases, data in the archive, data in motion, and the authentication of the various components of the xStream cloud stack. We can also maintain encryption and policies as data is moved and replicated. We support a wide variety of integrated key management options from sole responsibility to any flavor of shared responsibility. Virtustream's security and compliance certifications provide assurance of our security controls implemented in our data centers to protect our Customer's data starting from creation to secure disposal. Virtustream follows NIST guidelines for the secure disposal and media sanitization as applicable to the different category of data classification as "low risk", medium risk", and "high risk" data. More details about data security provisions provided by the Virtustream Private cloud are provided in Attachment C – Virtustream Private Cloud.

ServiceNow (SaaS)

Data security measures provided by ServiceNow: Data security is paramount to ServiceNow and ServiceNow have engineered their cloud services, the infrastructure that supports it, data encryption techniques, and security threat response processes so that customer's data is protected and secure at all times. Customer instances are hosted in a "private' environment that is dedicated to hosting ServiceNow's subscription-based cloud services. No other public cloud-hosting capabilities are used to deliver the service. ServiceNow's advanced multi-instance architecture provides that Customer's instance is logically separated from all other tenants in a cloud environment. This means that there is no co-mingling of customer data. ServiceNow adds new security properties to each release of the Now

Platform and provides specific recommendations for enabling these properties. Customers benefit from the services ServiceNow provides that support application security, malware protection, network security, system configuration, Identity and Access Management (IAM), security response, and data protection.

As a customer, browser-based sessions to your ServiceNow cloud instance(s) are encrypted over the internet via Transport Layer Security (TLS) using AES-128 or AES-256 block ciphers for data in transit security.

Everything inside the co-location spaces is owned, operated, and managed by ServiceNow. This includes the management of hard drives and server hardware. All hard drives are sanitized prior to leaving our private cages (per NIST 800-88 guidelines) so that your data is appropriately handled and protected. You can choose to further mitigate data exposure caused by the loss or theft of storage devices with AES-256 full-disk encryption of your data at rest.

All Customer data is hosted on solid-state or mechanical disks within ServiceNow's colocation spaces. No tapes or other forms of removable media are used to provide the service, including for backups. When functional storage devices reach their end-of-life or get reassigned to new customers, these are logically shredded using a process based on guidance from the U.S. National Institute of Standards and Technology (NIST). More details about data security provisions provided by ServiceNow are provided in Attachment D – ServiceNow.

BMC Remedy on Demand (SaaS)

Data security measures provided by BMC for Remedy on Demand: Safeguarding the privacy and security of personal information is a top priority for BMC Software in our data driven-economy. BMC has become the world's first enterprise IT management provider to secure EU accreditation for its Data Privacy Binding Corporate Rules (BCRs) as both a Controller and Processor of personal data. BCRs are considered to be the platinum standard for compliance in data privacy and personal data protection worldwide. BMC manages the security of customer's data by using the following guidelines:

- Strong physical security mechanisms are in place at all BMC OnDemand facilities based on SSAE16 (or equivalent) certified data centers;
- All external solution traffic is secured using encryption;
- You are allocated a dedicated environment that leverages virtualization for the mid-tier and application server components;
- Your database is dedicated to your data (data is not mixed among customers);
- The infrastructure and applications are configured to account for security standards using a hardening process to reduce security vulnerabilities;
- Monitoring is in place to alert you of any suspected or actual data breaches;
- Periodic penetration tests are performed to identify any potential or actual security issues;
- The operational and support organizations employ the separation of duties security principle so that only the resources required to support the solution have access to specific data;
- Periodic internal and external security audits are run on the systems to identify any vulnerabilities.

Remedy OnDemand customers retain ownership of their data at all times. Should a customer decide to leave the service, BMC will provide a file containing the Customer data in comma separated value (.csv) or database backup format upon customer request. Customer data is then destroyed via destruction of database encryption keys and data is overwritten with binary zeroes.

For all cloud-based services provided by Capgemini and its Partners, it is important for Purchasing Entities to understand some key points regarding the data ownership and management in the cloud, referred as "shared responsibility model":

- Purchasing Entities continue to be the owner of their data;
- Purchasing Entities decides the data security controls to be implemented as per data classification and criticality;
- Purchasing Entities choose the geographic location (s) to store their data;
- Purchasing Entities can download or delete their data whenever they intend to do so.

Capgemini Enterprise iPaaS (PaaS)

Because iPaaS does not store "Client" pass-through data, Capgemini's view is that security measures for Enterprise iPaaS are synonymous with managing strict control over changes to the environment. All components of the Capgemini Enterprise iPaaS are stored within a source control library. The Enterprise iPaaS Change Management process covers 3 release types: Critical Patches; Maintenance Releases; Major Upgrades. Each release is subject to a test assurance process. For each type of Release, the Purchasing Entity will be given access to release notes, and advance notice of the implementation containing risk assessments commensurate with the nature of the change. Service management tooling provides alerting and dashboarding of incidents, or Capgemini Enterprise iPaaS can be connected to the Purchasing Entity's existing tooling via the Additional Service Connect process. User incidents must be reported via the Capgemini Enterprise iPaaS Service Desk Portal. Users will be able to see the incidents the users have raised, service availability and performance stats.

Anti-virus and Intrusion detection is provided as standard on the Capgemini Enterprise iPaaS Platform. Alerts are reported via dashboards on the Enterprise iPaaS Platform.

If additional requirements are identified to comply with a Purchasing Entity's security policy or a request is made to provide a Security Management Plan and/or an Information Security Management System, these can be considered and priced in accordance with a SOW.

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

Capgemini and its Partners, collectively "Offeror", will comply with applicable laws, regulations and security standards related to data privacy and security. Identified security compliance requirements will be agreed upon and addressed in Master Service Agreement, Participating Addendum, or SLA with Purchasing Entities for the scope of services. The offeror will be responsible to provide required security controls to deliver services to Purchasing Entities in compliance with applicable laws and related data privacy and security requirements.

Capgemini is a security and privacy aware organization which is evident in our certifications, accreditations, and security program we manage internally. Our mandatory security guidelines are enforced with our Cloud services Partners as well as through the Master Service Agreements we have with our Partners to provide an appropriate level of security and privacy to our Customer's data we entrust to our Partners. Capgemini and our Partners comply with the applicable laws and related data privacy and security compliance requirements. More details about the Capgemini's and our Partner's security and compliance certifications & accreditations are provided in Capgemini's response to question in section 8.6.2.

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.

Capgemini will not access Purchasing Entity's user accounts or data unless it is needed to deliver the service (s) in scope and expressed by the terms of the Master Agreement, Participating Addendum, or SLA. Capgemini and our Partners will comply with Purchasing Entity's security policy, standards, and procedures and limit access to data relevant to job responsibilities and grant access required to deliver the responsibilities with adequate segregation of duties. Security controls are available to

enforce, monitor and report on access violations through Value add cybersecurity services. An additional layer of controls can be implemented to address the security requirements to process and access the high-risk data as described in Capgemini's response to question in section 8.6.3.

During a project or implementation, Capgemini will leverage (SAM), which is Citrix based Virtual Private Network (VPN). SAM is remote access model which enables remote support access to Capgemini's staff to the customer's network, systems, and applications via a secure, standard, two-tiered Citrix implementation with Transport Layer Security and multi-factor authentication. The Citrix based virtual desktop through which Capgemini's staff access Customer's network have stringent security controls implemented to prevent any transfer of data from Customer's network to Capgemini's network and vice-versa. Virtual desktops are secured and locked down by disabling clipboard functionality, local drive & printer access to block data transfer from virtual desktop to Capgemini network and vice-versa. Internet access from virtual desktops is also restricted to block any data upload to the internet and download any malicious data from internet to Customer's network. Any exceptions in the policy are reviewed and approved with the Customer.

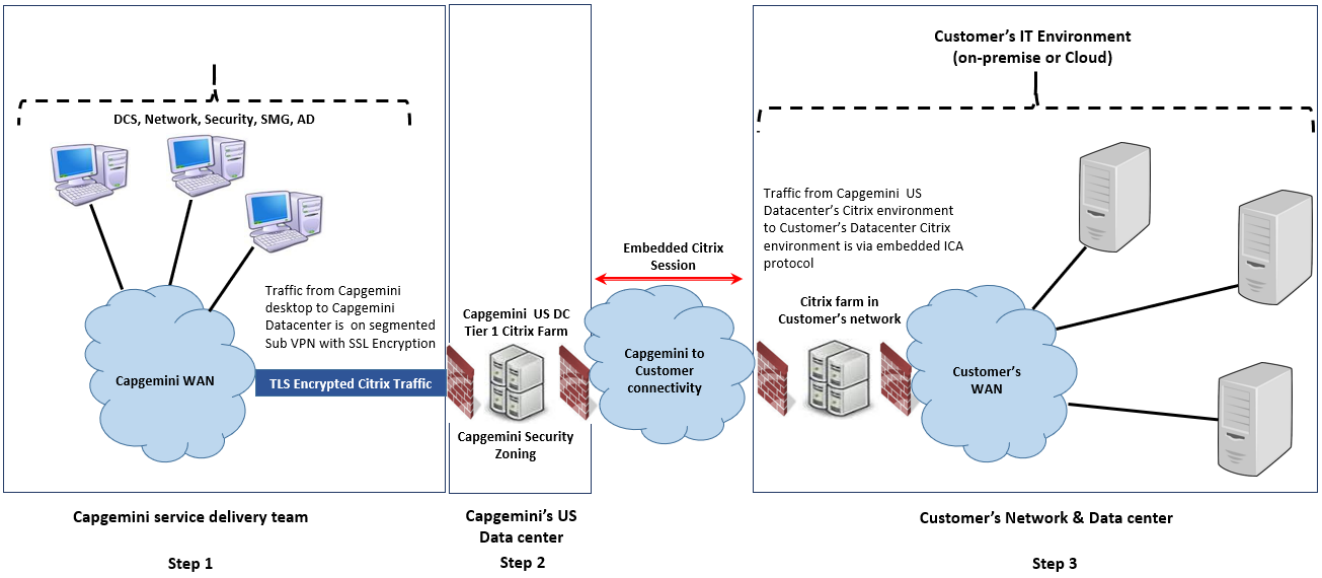

Figure 5: Customer Connectivity

Step 1: Login to Capgemini's network

- Level 1 authentication – corporate domain credentials (domain Login / password) + site access card;
- Traceability - Capgemini internal IT Logs and physical security logs;
- Access to Capgemini's office systems having hardened laptops, desktops and operating systems with corporate information security policies applied.

Step 2: Login to Capgemini's SAM Data center environment

- Level 2 authentication - CGADMIN authentication (User ID's password is different from corporate domain credentials) to connect to US Datacenter Citrix farm with multi-factor authentication;
- Traceability - Capgemini CGADMIN group membership, timestamp, and Citrix access logs;
- Access reporting and review;
- Use of software token for second-factor authentication.

Step 3: Access to Customer's Citrix farm

- Level 3 authentication - Customer specific credentials (which are different from Capgemini's corporate domain credentials and CGADMIN authentication) to access Citrix farm hosted in Customer's network (on-premise data center or Cloud) with multi-factor authentication;
- Traceability – Customer's Citrix farm access group membership, timestamp, and Citrix access logs;
- Access reporting and review;
- Use of software token for second-factor authentication.

Our Partners do not access customer's data, and customers are given choice as to how they store, manage and protect their data with our Partners while leveraging cloud services. Customers retain control and ownership of their data, and all the data stored by our Partners on behalf of our Customers have strong tenant isolation security and control capabilities. Customers must consider the sensitivity of their data if the how they will manage access permissions, how they will secure data while it is in transit and while it is at rest.

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.

Capgemini Cloud solutions and Capgemini's Cloud Service Partners are all committed to complying with and aligning to NIST definition of cloud computing as stated in NIST publication 800-145 for the scope of services described in Attachment D. Capgemini's response to the question in section 8.5.1, describes the security provisions Capgemini and its Cloud Service Partners have protected Purchasing Entity's data entrusted to Capgemini and its Cloud Service Partners. Capgemini and our Cloud Services Partners hold multiple security certifications and accreditations which meet the requirements and comply with the standards for Cloud computing and the scope of services described in Attached D. Capgemini's Cloud solutions leveraging our Partners for SaaS, PaaS and IaaS are capable of processing and storing State's different types of data from "Low-risk data" to "High-risk data" as categorized in FIPS publication 199.

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 the 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.

Capgemini Cloud solutions and Capgemini's Cloud Service Partners are all committed to complying with and aligning with standards organizations and security certifications.

> *Capgemini holds ISO/IEC 27001:2013 (Information Security Management) certification, Capgemini's service delivery is certified to ISO/IEC 20000-1 (Information Technology Service Management), ISO 9001 (Quality Management). Capgemini has also aligned our Cloud Infrastructure Services (CIS) organization with multiple other Cloud and Cybersecurity standards applicable to government and commercial Customers such as PCI-DSS, FISMA, GLBA, HIPAA, ISAE 3402/SSAE 16, NIST Cybersecurity framework, FedRAMP, GDPR, HITECH, CSA, FFTEC, COBIT, and ISO 22301. For value-added cyber security services, Capgemini will leverage cyber security professionals with the required certifications necessary to meet the demands of data privacy and data protection mechanisms.*

Below section provide details about the government, security, and cloud relevant certifications and accreditations our Cloud Services Partners hold and manage to secure Purchasing Entities data and applications.

AWS (IaaS, PaaS, SaaS)

Our Cloud services Partner Amazon have following certifications and accreditations applicable to all Cloud service models being offered by Amazon through Amazon Web Services.

Amazon Web Services' (AWS) Cloud and Security certifications and accreditations	
Certification and accreditations	DoD SRG, FedRAMP, FIPS, IRAP, ISO 9001, ISO 27001, ISO 27017, ISO 27018, MLPS Level 3, MTCS, PCI DSS Level 1, SEC Rule 17a-4(f), SOC 1, SOC 2, SOC 3, DNB [Netherlands], EU Model Clauses, FERPA, HIPAA, IRS-1075, ITAR, My Number Act [Japan], VPAT / Section 508, EU Data Protection Directive, CJIS, FedRAMP TIC, FISC, FISMA, GxP (FDA 21 CFR Part 11), IT-Grundschutz, MPAA, NERC, NIST, UK Cyber Essentials

Azure (IaaS, PaaS, SaaS)

Our Cloud Services Partner Microsoft has the following certifications and accreditations applicable to all Cloud service models being offered by Microsoft Azure. Azure meets a broad set of international and industry-specific compliance standards, such as General Data Protection Regulation (GDPR), ISO 27001, HIPAA, FedRAMP, SOC 1 and SOC 2, as well as country-specific standards, including Australia IRAP, UK G-Cloud, and Singapore MTCS. Rigorous third-party audits, such as those done by the British Standards Institute, verify Azure's adherence to the strict security controls these standards mandate.

Microsoft Azure' Cloud and Security certifications and accreditations	
Certifications and accreditations	CDSA, CJIS, CSA CCM, EU Model clauses, FDA 21 CFR 11, FedRAMP, FERPA, FIPS 140-2, HIPAA, IRAP, ISO/IEC 27018, ISO/IEC 27001:2013, MPLS, MTCS, PCI DSS, TCS CCCPPF, UK G-Cloud

Virtustream (IaaS)

Our Cloud services Partner Virtustream have following certifications and accreditations applicable to all IaaS Cloud Service model being offered by Virtustream Private Hosting.

Virtustream Private Hosting's Cloud and Security certifications and accreditations	
Certifications and accreditations	FedRAMP, FISMA, SSAE16/SOC2/ISAE3402, PCI DSS 3.2, ISO 27001:2013, ISO 9005:2015, ISO 22301:2012, HIAA/HITECH/HITRUST, CSA STAR, CJIS

ServiceNow (SaaS)

Our Cloud Services Partner ServiceNow has following certifications and accreditations for ServiceNow SaaS cloud service model being offered.

ServiceNow' Security certifications and accreditations	
Certifications and accreditations	ISO 27001:2013, ISO 27018, SSAE 16 SOC 1 Type 2, Section 508 Adherence, FedRAMP Moderate P-ATO Certification, DOD Impact Level 2, Pink Verify, FIPS 140-2

BMC Remedy on Demand (SaaS)

Our Cloud Services Partner BMC has following certifications and accreditations for Remedy on Demand SaaS cloud service model being offered.

BMC Remedy on Demand' s Security certifications and accreditations	
Certifications and accreditations	BMC's OnDemand offerings are designed based upon NIST (National Institute of Standards and Technology) NIST 800-53, Rev 4 controls and standards at a Moderate baseline in order to provide enterprise-grade security for our customers. In addition, BMC maintains a SOC 2 Type II for Security, Availability, and Confidentiality trust principles. BMC also has FedRAMP ATO.

Capgemini (PaaS)

The Capgemini's Enterprise iPaaS has following certifications and accreditations applicable to all PaaS cloud service model being offered.

Capgemini's Enterprise iPaaS certifications and accreditations	
Certifications and accreditations	▪ ISO/IEC 27001 ▪ HMG Cyber Essentials

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.

Capgemini's Cloud Services Partners are certified and accredited to multiple cloud and security compliance standards as described in Capgemini's response to question in section 8.6.2. Capgemini and Capgemini's Cloud Services Partners have implemented appropriate security controls and follow industry security practices to secure data and applications, including external and internal threats. Appropriate physical and technical security controls are implemented to provider required segregation of data and services for co-located customers for cloud-based services.

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).

As stated in Capgemini's response to question in Section 8.5.1, Capgemini and Capgemini's Cloud Service Partners have provisioned appropriate security controls to protect customer's data on the cloud to protect its confidentiality and access, modification from unauthorized personnel. Capgemini and Capgemini's Cloud Service Partners operate and manage the Information Security Management System (ISMS) and have certified the services to ISO 27001:2013 security standard. ISMS define the security policies and standards put together to protect the confidentiality, availability, and integrity of data by having applicable security controls implemented including the administrating controls such as segregation of duties, access management, and ongoing access reviews to make sure only authorized personnel to have access to customer data and applications.

Additional services are available from Capgemini's Cloud Service Partners to deliver 24x7x365 security monitoring services to monitor the Purchasing Entity's IT environment for malicious behavior, unauthorized access attempts and incident response.

- For the security of laptops, mobile, and other IT services, Capgemini maintains its security program by operating a Cybersecurity and Information Protection (CySIP) program which includes:
- Risk Management – assessing the risks to the business assets, treating the risks and identifying the mitigating controls to be applied if required;
- Controls Management – providing that the identified controls are correctly applied, and their function is effective in mitigating the risk; and
- Impact Management – covering Security Incident Management, confirming that all security breaches are detected, constrained and remediated;
- Overall there are 68 mandatory controls and specific strategic documents in place such as the Data Protection strategy and Data Privacy policy.

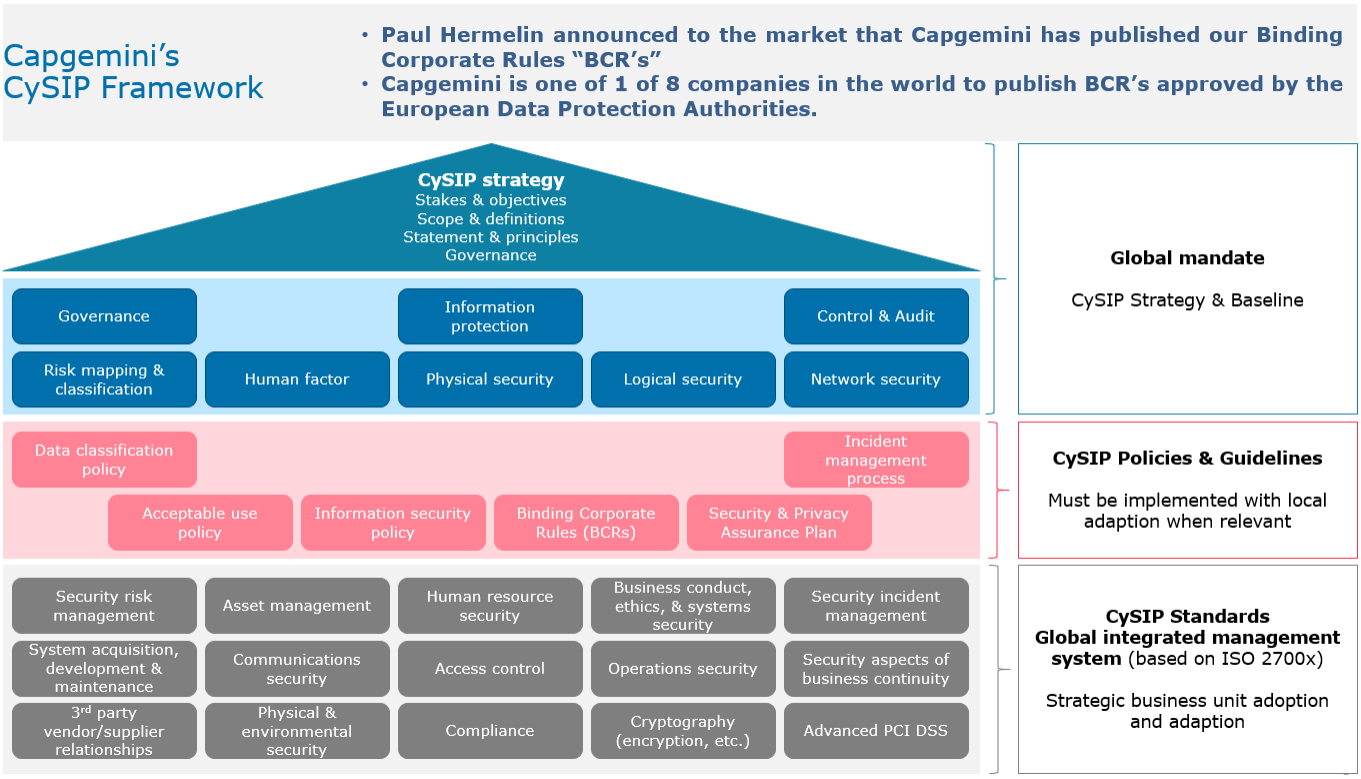

Figure 6: Capgemini CySIP Framework

The Information Security processes and security profiles are also validated by third-parties during a number of statutory external audits or via third-party security assessments. Additionally, our security controls are derived from our Customer's needs and their compliance requirements.

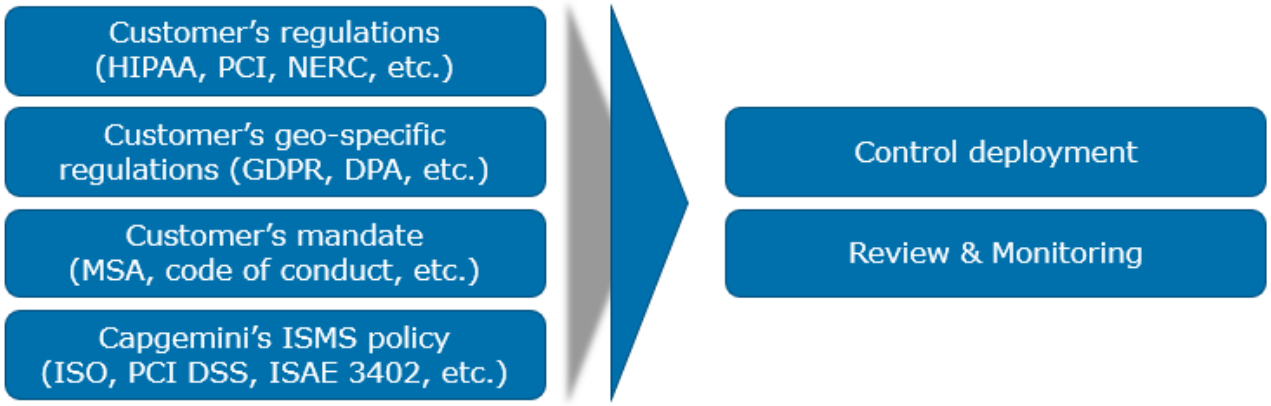

Figure 7: Security Control

Capgemini's security program and controls are derivatives of comprehensive risk assessment methodology, which are conducted at least on annual basis and whenever there is a significant change in business or technology environment. This provides a standardized information security framework to apply to all its business functions and our delivery centers. However, additional security layers are built based on client' specific requirements and the regulations which govern the client business such as the Data Protection Act, PCI-DSS, ISAE3402/SSAE 16 (SAS-70) etc.

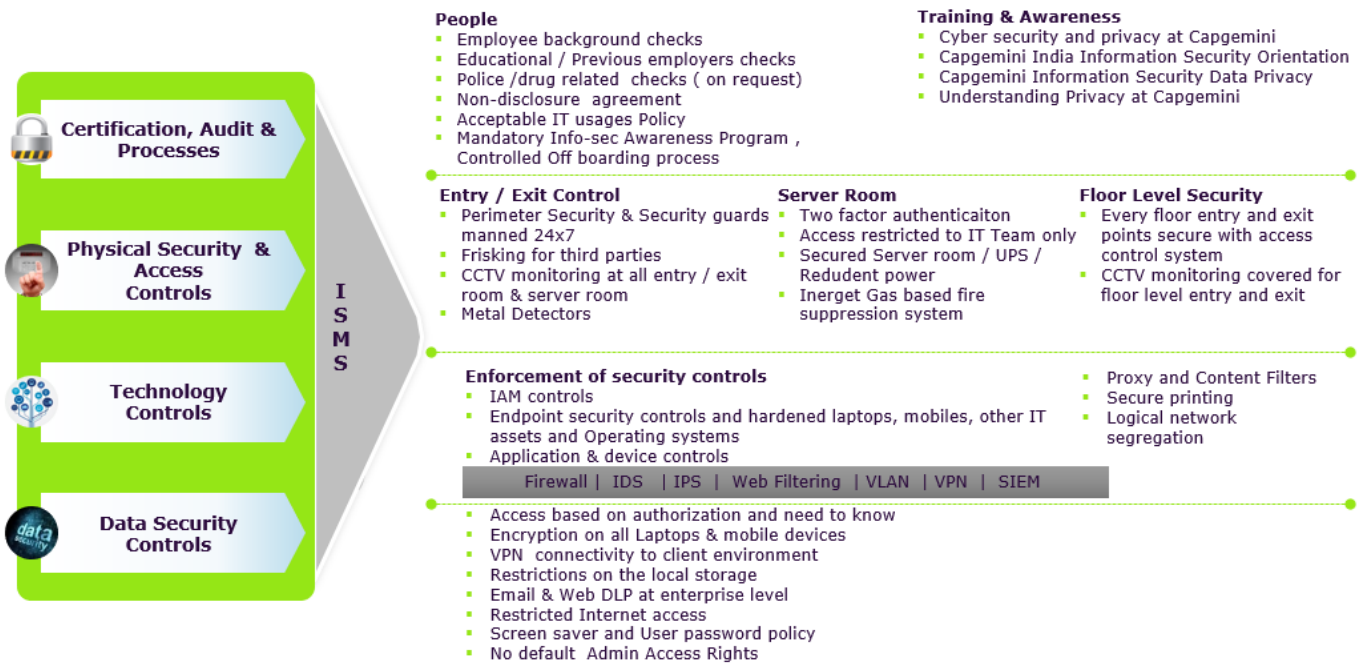

Figure 8: Capgemini Security Program

Capgemini and Capgemini's Cloud Service Partners will process client's personal data only as instructed by the client and will not use such personal data other than as necessary to perform the services. Capgemini and Capgemini's Cloud Service Provider's personnel are required to keep data confidential by agreeing to confidentiality obligations and signed Non-Disclosure Agreement (NDA) with personnel having access to sensitive and private information

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

Please refer to Capgemini's response to question in section 8.6.2 for the detailed list of the third-party attestations, reports, security credentials and other controls available for cloud services.

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.

Capgemini leverages the logging capabilities within Amazon, Azure, and Virtustream which provides a comprehensive log management and analysis strategy in Cloud which is mission critical, enabling Purchasing Entities to understand the relationship between operational, security, and change management events and maintain a comprehensive understanding of their Cloud infrastructure and applications.

Leveraging logging service capabilities from Capgemini's Cloud Service Providers, Purchasing Entities will have access to service-specific metrics and log files to insight into how each Cloud service is operating, and many services capture additional data, such as Application Programming Interface (API) calls, configuration changes, billing, and security events. Log files from web servers, applications, and operating systems will provide valuable data in different formats, and in a random and distributed fashion.

Cloud solutions offered by Cloud Service Partners, supporting a variety of devices and events types available for infrastructure monitoring, management, and security monitoring. Logging capabilities available with individual Cloud Service Partner vary and services are available as a base (legacy) and advanced (with additional service cost). Capgemini's Cloud Service Partners work with our customers to establish a standardized and consistent approach to security logging thus helps to mitigate risks

around the loss of integrity and control of applications and systems in the cloud, data loss and reputational impact to an acceptable level. Defined logging standard for our customers sets out the minimum requirements for the selection of logs, guidance to users regarding the implementation, deployment and operational practices for logging. Capgemini's Cloud Service Partners follow security log management guidelines recommended by NIST special publication 800-92 with adequate consideration to individual customer's requirements for log management planning and operational processes to comply with customer's information security policy. We work with our customers to diligently identify what logs need to be enabled, what log types need to be captured, processed, stored and analyzed to comply with certain security compliance requirements such as FISMA, HIPAA, SOX, GLBA, and PCI DSS. Each customer's business is different, requirements and different and hence no one size fits for all. We work with our customers to define logging policies that clearly define mandatory requirements and suggested recommendations for several aspects of log management, including the following:

- Log generation
  - Types of hosts, applications must or should perform logging;
  - Host components must or should perform logging (e.g., OS, service, application);
  - Types of events each component must or should log (e.g., security events, network connections, authentication attempts)
  - Data characteristics must or should be logged for each type of event (e.g., username and source IP address for authentication attempts)
  - How frequently each type of event must or should be logged (e.g., every occurrence, once for all instances in x minutes, once for every x instances, every instance after x instances)
- Log transmission
  - Types of hosts must or should transfer logs to a log management infrastructure
  - Types of entries and data characteristics must or should be transferred from individual hosts to a log management infrastructure
  - Log transfer mechanisms (e.g., which protocols are permissible), including out-of-band methods where appropriate (e.g., for standalone systems)
  - The frequency of logs transferred from individual hosts to a log management infrastructure (e.g., real-time, every 5 minutes, every hour)
  - Controls for confidentiality, integrity, and availability of each type of log data must or should be protected while in transit, including whether a separate logging network should be used
- Log storage and disposal
  - Log rotation frequency
  - Controls for confidentiality, integrity, and availability of each type of log data must or should be protected while in storage (at both the system level and the infrastructure level)
  - Log retention period preserved (at both the system level and the infrastructure level)
  - How unneeded log data must or should be disposed of (at both the system level and the infrastructure level)
  - Log storage management (at both the system level and the infrastructure level)
  - How to log preservation requests, such as a legal requirement to prevent the alteration and destruction of particular log records, must be handled (e.g., how the impacted logs must be marked, stored, and protected)
- Log analysis
  - Logs analysis mechanisms (at both the system level and the infrastructure level)

- Who must or should be able to access the log data (at both the system level and the infrastructure level), and how such accesses should be logged
- Log monitoring mechanism, process, and workflows
- How the confidentiality, integrity, and availability of the results of log analysis (e.g., alerts, reports) must or should be protected while in storage (at both the system level and the infrastructure level) and in transit
- How inadvertent disclosures of sensitive information recorded in logs, such as passwords or the contents of e-mails, should be handled.

Our Cloud Service Partners offers multiple capabilities for system, application, and security logging capabilities for the cloud services available for consumption through SaaS, PaaS, and IaaS cloud service models.

Capabilities available with our Cloud Service Partners have following features:

- Increased visibility leveraging cloud APIs for tight integration and richer context;
- Secure log storage;
- Easy administration;
- Notifications for log delivery with detailed information and time stamps;
- Flexibility to choose from multiple Vendor solutions what suits your requirements;
- Log file aggregation and correlation.

More details about the logging process and capabilities available with our Cloud Service Partners is provided in:

- Attachment A – Amazon Web Services
- Attachment B – Microsoft Azure
- Attachment C – Virtustream Private Cloud
- Attachment D – ServiceNow
- Attachment E – BMC Remedy on Demand

Capgemini's Cloud Service Partners have established Information Security Management System (ISMS) to manage and maintain the security controls implemented to comply with security certifications, compliance regulations/requirements, and standards. Capgemini's Cloud Service Partners are periodically audited and examined by the independent 3rd parties for compliance, applicability including surveillance audits to maintain the certifications and accreditations. Applicable certifications and accreditations are renewed after passing follow-up assessments and surveillance audits by recognized auditors/assessors and qualified 3rd parties.

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

Capgemini and Capgemini's Cloud Service Partners, collectively stated as "Offeror", can restrict visibility of cloud-hosted data and documents to specific users and groups. Data classification, system classification, and role-based identity & access controls are available and are enablers to implementing granular access controls enforced through segregation of duties, need to know, administrative and technical controls to protect unauthorized access and usage of data and documents. Capgemini's response to a question in section 8.5.2 provide details on controls as for how Capgemini's authorized users are controlled to access the Customer's platforms and applications with multiple layers of authentication enabled with multi-factor authentication controls.

Capgemini's Cloud Service Partners have stringent access controls available, additional cloud environments will be designed to provide isolation by implementing virtual networks and security

groups in such a way that will prevent data access from unauthorized resources. VPN connections will be established to protect access and network traffic to and from the Internet. Furthermore, security measures can be implemented to provide authentication and authorization to applications and associated data utilizing Active Directory applying Role Based Access Control. Identity and access management services available with our Cloud Service Partners can also be leveraged and implemented to have granular role-based access to cloud-based applications and systems for Purchasing Entity's users and customers. Integration options are also available for customers to integrate their organization's user identity databases such as Active Directory (AD) to cloud service providers IAM tools such AWS IAM, Microsoft AD to restrict visibility of cloud-hosted data, application, and systems.

8.6.8 Offeror must describe its notification process in the event of a security incident, including relating to timing, incident levels. The 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.

Capgemini and Capgemini's Cloud Service Partners, collectively "Offeror", follow ITIL aligned incident management process including security incident management. Timelines, communication mechanisms, and different response requirements will be agreed with Purchasing Entities addressing applicable laws and data being processed by Offeror in Master Service Agreement (MSA), additional addendums, or SLA for the scope of services.

On a high-level Offeror will follow below process for security incident management and it will be customized for individual Purchasing Entities to address specific compliance and legal obligations:

- The scope of the incident is taken into consideration by determining risk, severity, affected users, assets, and locations. A security incident will be assigned as "Priority Level 1" if the incident is characterized by the following:
  - The incident is one that has a high impact on the confidentiality, integrity, or availability of information and will increase in scope and impact if the incident is not mitigated.
  - The incident, because of the immediacy of its effect on critical business functions or information, requires a resolution on an immediate response basis.
- A security incident will be assigned as "Priority Level 2" if the incident is characterized by the following:
  - The incident can materially affect the Purchasing Entities, causing a substantial impact.
  - The effect of an incident is such that it does not require immediate resolution, but it does require a resolution (for example a change) is executed on a date and time in the near future specified by the Purchasing Entities.
- A security incident will be assigned as "Priority Level 3" if the incident is characterized by the following:
  - The incident does not materially affect the Purchasing Entities and does not cause a substantial impact on confidentiality, integrity, or availability. The incident does have the potential to do so if not resolved expeditiously.
  - The effect of the incident is such that it does not require an immediate resolution, but it does require that a resolution (for example a change) executed on a date and time mutually acceptable to both parties.
- A security incident will be assigned as "Priority Level 4" if the incident is characterized by the following:
  - The incident does not have an adverse impact on confidentiality, integrity, or availability because of either the nature of the fault or the small extent of the fault and the fact that the incident will not increase in impact over time.

- The effect of the incident is such that it does not require immediate resolution. A resolution (for example a change) may be required that can be planned for a date and time that is mutually acceptable to both parties.

For each security incident with different Priority levels, the Incident Response team will lead a defined action plan, follow the different response and resolution timelines and do all necessary escalations to defined stakeholders appointed by Purchasing Entity or using a pre-agreed escalation matrix between Offeror and Purchasing Entities.

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.

Capgemini's Cloud Service Partners have implemented appropriate security controls to provide necessary physical and virtual isolation for hosted infrastructure (servers) and data entrusted to our them and their support staff.

For the services in scope, Capgemini will design, during an implementation SOW, the cloud environment for Purchasing Entities with virtual networks, subnets and firewalls establishing rules to grant access to authorized users and resources, traffic that is not explicitly authorized will result in denying access to protected resources. Segregation of traffic will be established through network security groups, firewalls, establishing a DMZ for further security controls to grant explicit access to protected resources to and from the Internet. The least privileged security access settings will be established to prevent unauthorized access.

AWS (IaaS, PaaS, SaaS)

Our Cloud Service Partner Amazon's Virtual Private Cloud (Amazon VPC) lets customers provision a logically isolated section of the AWS Cloud where customers can launch AWS resources in a virtual network that customers define. Customers have complete control over the virtual networking environment; including a selection of customer owned IP address ranges, the creation of subnets, and configuration of route tables and network gateways. The customer can use both IPv4 and IPv6 in their VPC for secure and easy access to resources and applications.

The customer can easily customize the network configuration for the customer's Amazon VPC. For example, customers can create a public-facing subnet for web servers that have access to the Internet, and place backend systems such as databases or application servers in a private-facing subnet with no Internet access. The customer can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Additionally, customers can create a Hardware Virtual Private Network (VPN) connection between the customer corporate data center and their VPC and leverage the AWS Cloud as an extension of their corporate data center. Amazon AWS Virtual Private Cloud components are depicted in the diagram below.

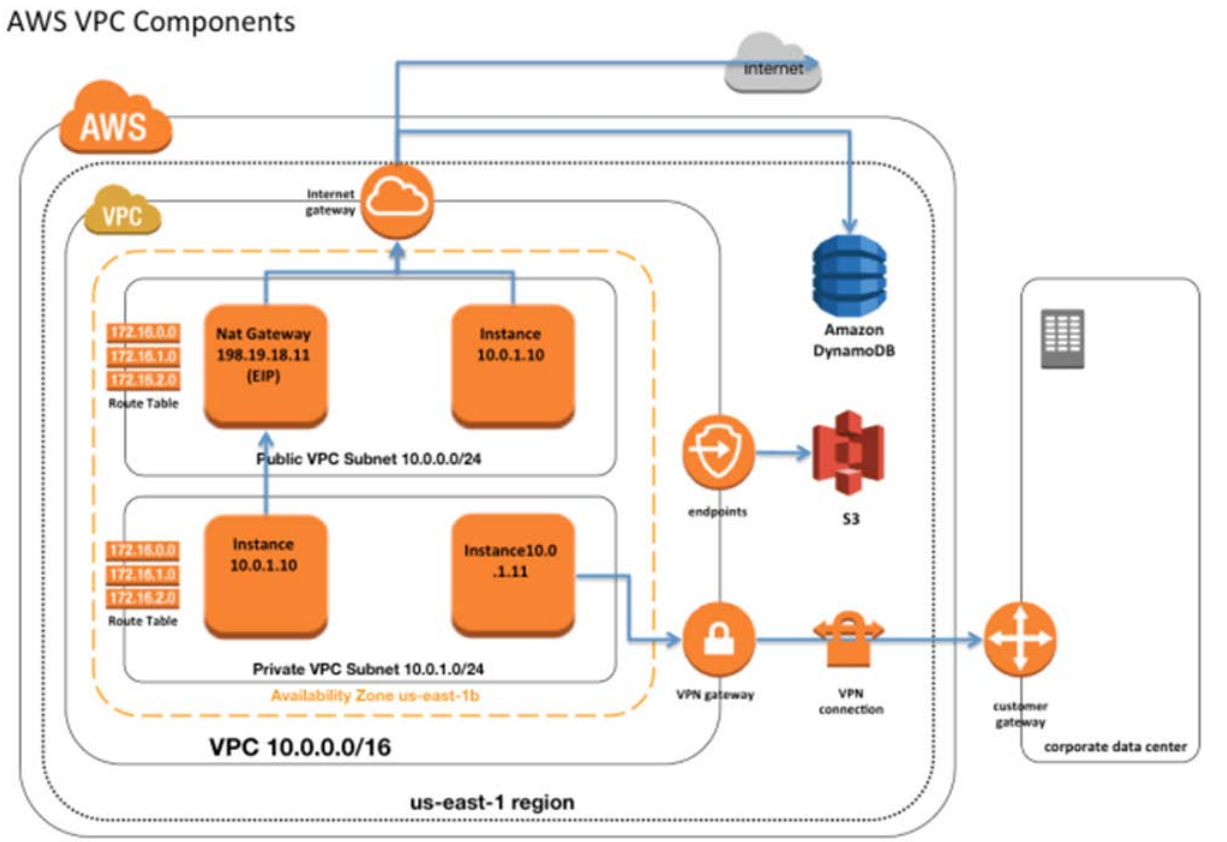

Figure 9: AWS VPC Components

Our Cloud Service Partner Azure's Virtual Network provides many types of Azure resources, such as Azure Virtual Machines (VM), to securely communicate with each other, the internet, and on-premises networks. Azure Virtual Network provides the following key capabilities:

Customers can implement multiple virtual networks within each Azure subscription and Azure region. Each virtual network is isolated from other virtual networks. For each virtual network customers can:

- Specify a custom private IP address space using public and private (RFC 1918) addresses. Azure assigns resources in a virtual network a private IP address from the address space that you assign;
- Segment the virtual network into one or more subnets and allocate a portion of the virtual network's address space to each subnet;
- Use Azure-provided name resolution, or specify your own DNS server, for use by resources in a virtual network.

All resources in a virtual network are controlled to communicate outbound to the internet with granular access controls. Customers can connect their on-premise It assets and networks to an Azure virtual network using any combination of the following options:

- Point-to-site virtual private network (VPN): Established between a virtual network and a single computer in the customer network. Each computer that wants to establish connectivity with a virtual network must configure its connection. The communication between the computer and a virtual network is sent through an encrypted tunnel over the internet.
- Site-to-site VPN: Established between the customers on-premises VPN device and an Azure VPN Gateway that is deployed in a virtual network. This connection type enables any on-premises resource that is authorized to access a virtual network. The communication between the on-premises VPN device and an Azure VPN gateway is sent through an encrypted tunnel over the internet.

- Azure ExpressRoute: Established between customers network and Azure, through an ExpressRoute partner. This connection is private. Traffic does not go over the internet.
- Customers can filter network traffic between subnets using either or both of the following options:
  - Network security groups: A network security group can contain multiple inbound and outbound security rules that enable customers to filter traffic to and from resources by source and destination IP address, port, and protocol.
  - Network virtual appliances: A network virtual appliance is a VM that performs a network function, such as a firewall, WAN optimization, or other network function.

Azure (IaaS, PaaS, SaaS)

Azure Virtual Network Cloud components are depicted in the diagram below, connecting a customer VNet and dedicated workload in a PaaS environment.

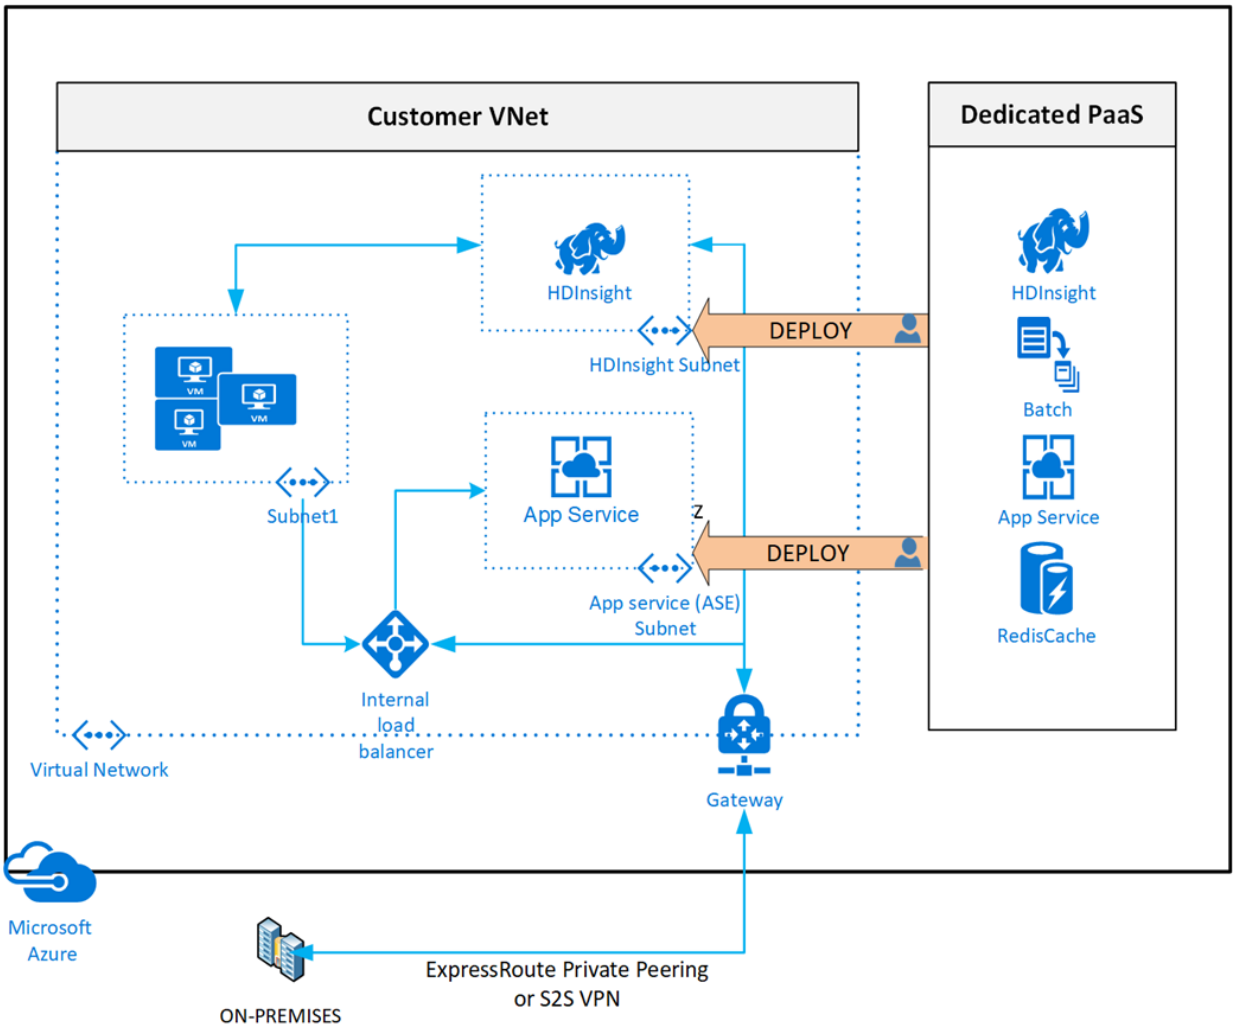

Figure 10: Azure Virtual Network Cloud components

Virtustream (IaaS)

Our Cloud Service Partner Virtustream's Private Cloud offerings provide the highest level of industry-leading security service available to state and government agencies with appropriate security controls implemented to provide physical and logical segregation for IT infrastructure and data hosted in the Virtustream Private cloud. Virtustream's environment has been designed and certified to meet or exceed the security requirements with the following core enterprise architecture, referenced standards, and frameworks:

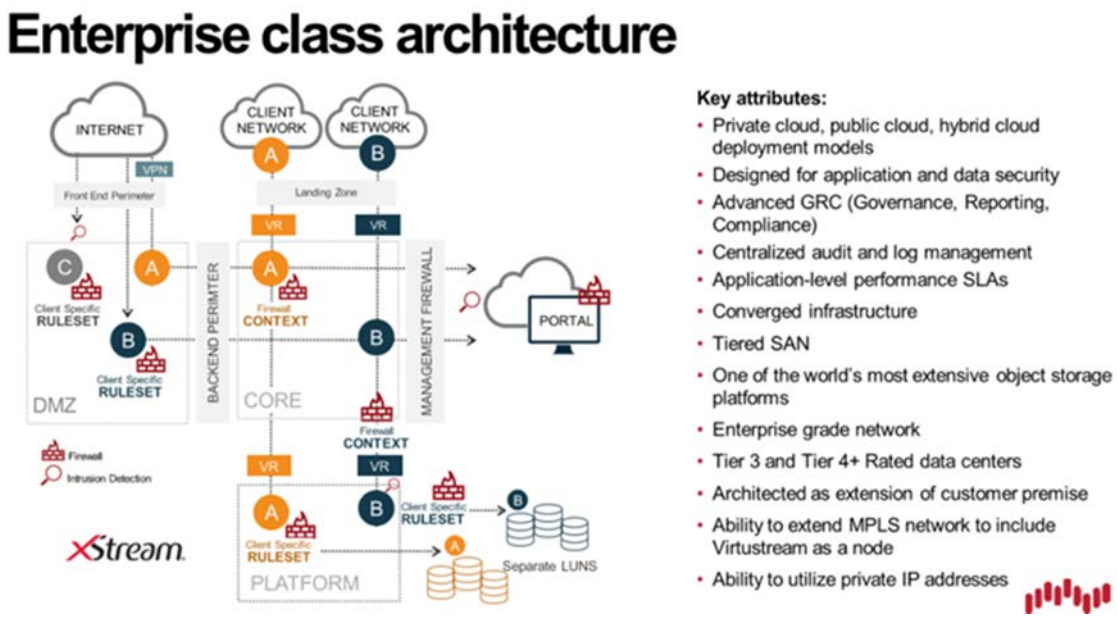

Figure 11: Enterprise Class Architecture

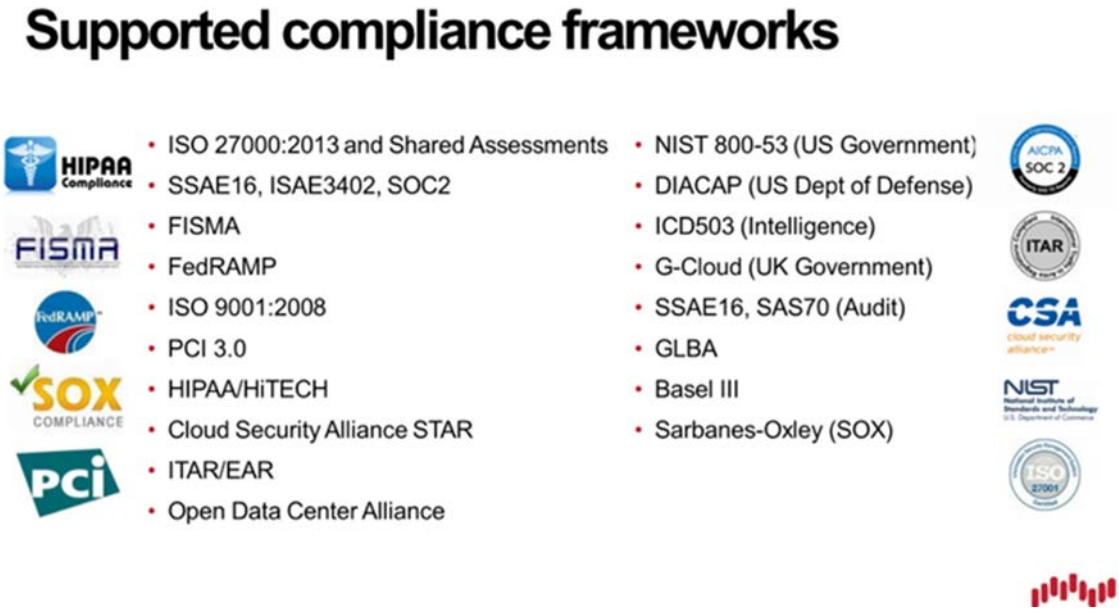

Figure 12: Supported Compliance Frameworks

ServiceNow (SaaS)

Capgemini's Cloud Service Partner ServiceNow have appropriate physical and logical security controls implemented to provide required segregation for hosted infrastructure and data. Please refer to Attachment D – ServiceNow for more details on ServiceNow capabilities.

BMC (SaaS)

Capgemini's Cloud Service Partner BMC have appropriate physical and logical security controls implemented to provide required segregation for hosted infrastructure and data. Please refer to Attachment E – BMC Remedy on Demand for more details on BMC capabilities for Remedy on Demand services.

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).

Capgemini will work closely with the Purchasing Entities to map out the technical reference architectures across IaaS, SaaS, and PaaS so that security is built into their cloud infrastructure – which includes selecting the right cloud deployment option, from the Cloud Service Provider who can offer the individual security measures embedded inside their architectures.

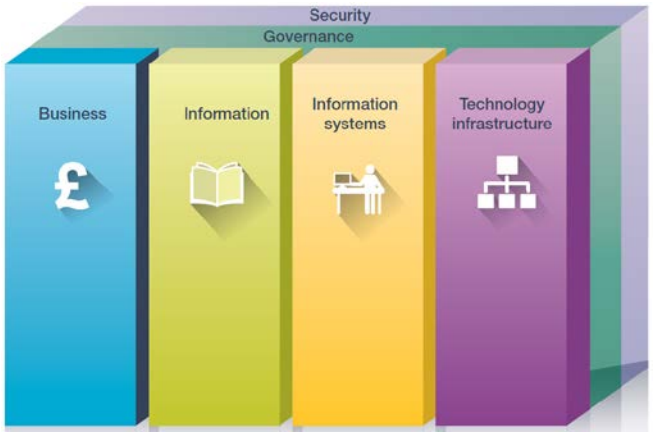

Capgemini offers Cloud strategy and consulting services to help the Purchasing Entities with their readiness to cloud adoption, plan the strategy for cloud transformation, execute transformation roadmap by implementing right cloud security architecture addressing business and security requirements. Purchasing Entities can purchase the cloud advisory and implementation service through Capgemini's **SK18008 – Capgemini - Detailed Product Offering Document**.

To best understand security considerations for cloud architecture for the Purchasing Entities, Capgemini uses four fundamental topics that are relevant for security measures within a cloud deployment model for SaaS, PaaS, and IaaS.

- **Data storage**
  - Assess data security and privacy requirements
  - Analyze data availability requirements
  - Evaluate connected applications and devices requirements
  - Options for virtualization
- **Location of data**
  - Define a location strategy
  - Data residency requirements choose the right location
  - Determine data security approach
  - Create governance framework with Cloud service provider
  - Encryption control requirements
- **Data security**
  - Data sensitivity
  - Encryption
  - Homomorphic encryption
  - Tokenization
  - Key management
  - Cloud Access Security Brokers (CASB)
- **Incident response planning**
  - Preparation
  - Detection and analysis
  - Containment, eradication, and recovery
  - Post-incident activities

Putting in place the right enterprise architecture framework makes it easier to do this: the framework contains processes, products, tools, and techniques needed to create a complete IT systems

architecture for all infrastructure – not only cloud. Capgemini's integrated architecture framework (IAF) is an example of such framework other examples include SABSA and TOGAF.

The framework helps to provide that security and governance are an integral part of the overall cloud architecture, positioning the organization to create the right controls (including monitoring). It also helps to limit cost: it is almost always less expensive to build security in from the start than to retrofit it. Control of non-financial costs such as those associated with damage to legislative compliance, reputation and customer confidence is another driver for adopting a framework. Following Capgemini's architecture framework has another important advantage: it makes it possible to trace characteristics of the technical implication back to either requirements or risks – something that is often vital for business governance and legal compliance on the cloud. Capgemini and its Architects have been active contributors in the development of industry standard architecture frameworks such as The Open Group Architecture Framework (TOGAF) and Cloud Security Alliance (CSA) reference architecture. In line with industry standard Cloud security reference architectures and guiding principles, Capgemini has created the below Cloud Security Reference Model (CSRM) to help our Customers to have end-to-end visibility and risk management of services for secure cloud adoption for SaaS, PaaS and IaaS service models applicable to our Cloud Service Partners.

Reference Security architecture design for the cloud: The figure below details Capgemini's approach for security architecture design for cloud-based services.

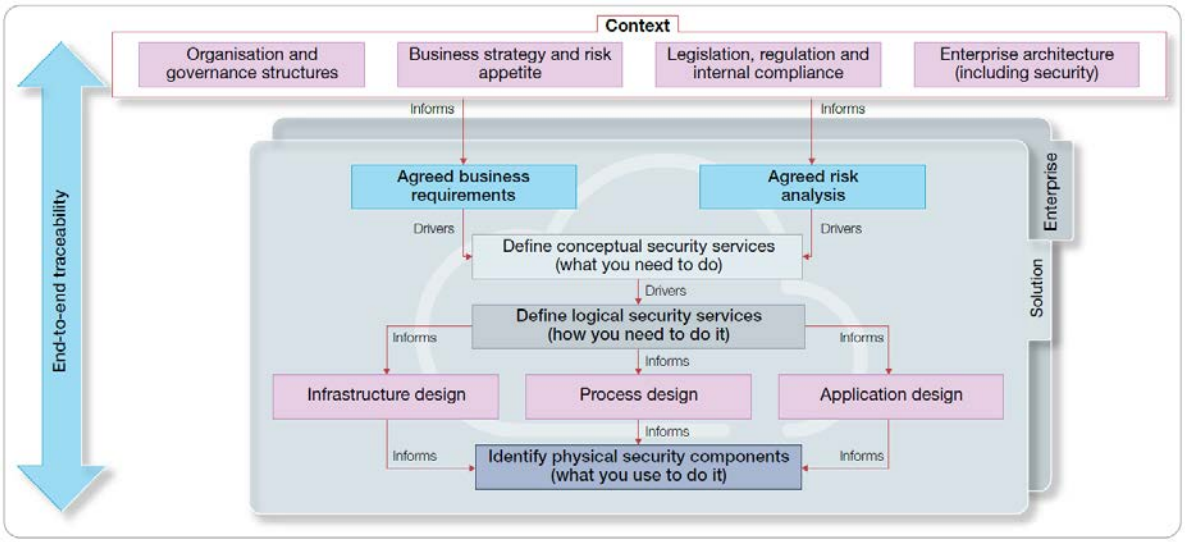

Figure 13: Approach to security architecture design

The design entails four views: the business context, the conceptual security architecture, the logical security architecture and the physical security components.

Business Context - The context consists of:

- High-level business requirements;
- Business needs;
- Regulations and legislation;
- High-level risk assessments;
- The division of responsibility between the consumer, CSP and third parties.

Conceptual security architecture. Based on the context, the architect defines the business context, security domains and the interconnections between the domains, which together constitute the conceptual security architecture. This architecture must be agreed by all the relevant stakeholders before proceeding with the logical controls selection.

Logical security architecture. The logical security architecture for the cloud is then derived from the conceptual security architecture. The architect selects controls from applicable standards, compliance and legal regulations such as:

- Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM) - the Security, Trust, and Assurance Registry (STAR) program provides varying assurance requirements and maturity levels of CSPs and consumers, and is used globally by customers, providers, industries, and governments.
- ISO/IEC 27001 – the requirements for information security management systems (ISMS).
- ISO/IEC 27018:2014 – an auditable standard for CSPs who process personal data. Includes around 70 controls from different international data protection laws.
- Cyber Essentials – a UK government-backed, industry supported scheme to help organizations protect themselves against common cyber-attacks. It provides a clear statement of the basic controls all organizations should implement to mitigate the risk from common internet-based threats, within the context of the UK government's 10 Steps to Cyber Security10.
- PCI-DSS – standard for protecting credit card data.

Physical security components. The next step is the selection of the physical components that will be used to implement the logical security controls just identified. In-depth knowledge of existing services and offers and their underlying security controls is important here, especially during the bid or proposal stage.

Security reference model – Capgemini leverages security reference model (SRM) during security architecture design for cloud services for our customers which illustrates a series of conceptual security

services grouped into high-level areas such as hosting, security governance, and access management.

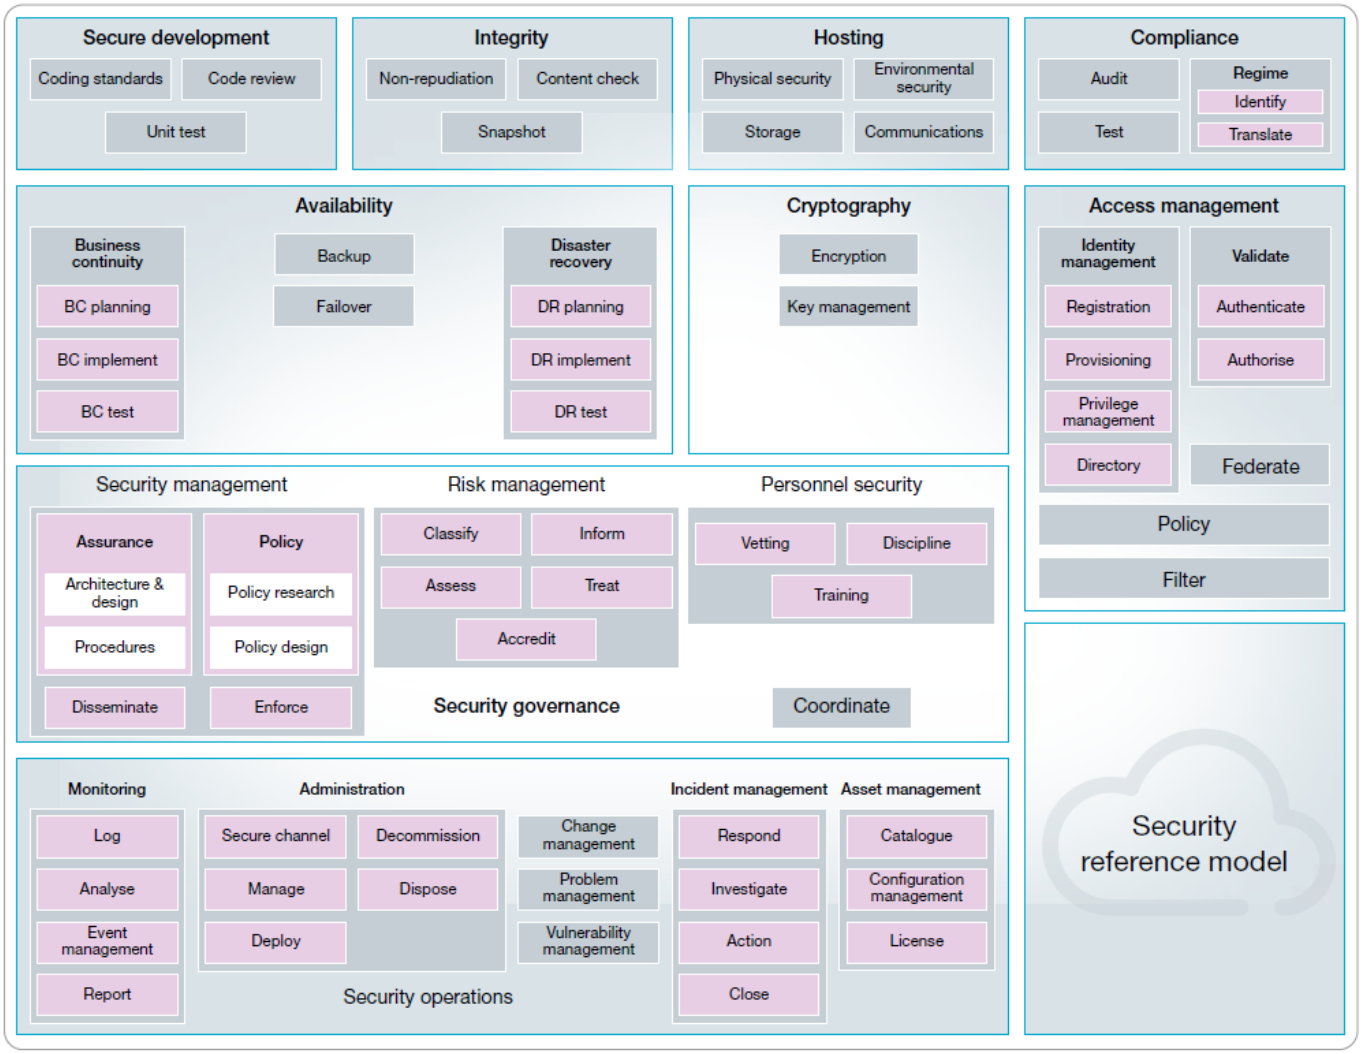

Figure 14: Generic security reference model (SRM) for cloud services

Capgemini leverages security reference model to identify services that are in scope for various security areas for our customers. The model can be used during procurement process to conduct due diligence of the security controls claimed by Cloud Service Providers (CSPs), the model can also show the division of responsibilities for the different cloud service models SaaS, PaaS and IaaS. The conceptual security architecture is used as the basis for a logical security architecture. Conceptual and logical security architecture (together with the wider business and non-functional requirements for Purchasing Entities) to identify appropriate physical components and then implement appropriately secure solutions for Purchasing Entities. Following are some important aspects Capgemini considers while designing cloud services architectures for our customers.

- **Functional and non-functional requirements**. Gather the requirements and group similar ones together.
- **Risk assessments**. Identify assets, threats and the business impacts of compromise, and again the need for certain conceptual security services may become obvious. Risk assessment methodologies and tools leverage and customers may adopt include:
  - ISO/IEC 27005:2011 – an international standard providing guidelines for information risk management.
  - Information Security Forum (ISF) – The Standard of Good Practice for Information Security is comprehensive and compatible with other well-known standards. It is particularly geared

towards ISF's own Information Risk Analysis Methodology (IRAM) and automated tool, Risk Analyst Workbench (RAW).

- CESG Information Standard 1/2 – with its supporting documents, this provides a suite of information risk management guidance for use, predominantly, by central government departments, the wider public sector, and their suppliers. However, it can be used by any organization to assess and manage technical risks.
- US National Institute of Standards and Technology (NIST) Special Publication 800-30 – this is the US government's preferred risk assessment methodology and is mandated for US government agencies. The methodology is usable by organizations of all sizes in both the private and public sectors. It is designed to be consistent with the ISO standards, and flexible enough to be used with other risk management frameworks.

- **Architecture frameworks**. Leverage architecture frameworks like TOGAF 9.1, IAF and SABSA.
- **Interactions between on-premise and on-cloud**. Identify whether services are joint or provider responsibilities. This leads to a better understanding of the interface (management, contractual and security) between the consumer and cloud service provider (CSP).
- **Physical realization**. After identifying the logical requirements of the service, we start to identify the technologies or processes that best fit our customer's needs – including the best deployment models.
- **Traceability and defensibility**. Identify and address any compliance and governance aspects. Traceability from either requirements or risks through to technical implementation is fundamental here – and convenient when the auditors visit to ask their usual questions.

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

Capgemini's hiring practices and procedures mandated by Capgemini's Group Human Resource (HR) organization and implemented by Capgemini's regional HR organization in compliance to regional/local labor laws and compliance/regulatory/Federal requirements. Capgemini performs pre-employment background checks as defined by Capgemini's HR policy and additional background screening & checks are performed for the employees hired for Regulatory/Government Customers accessing sensitive/classified data. Capgemini has "Security Cleared" resources who have clearance to deliver services to our Government Customers.

Some of the pre-employment screening checks are sourcing screening, recruiter screenings, technical interview, functional interview, in-person interview, educational & background verification. Capgemini also has provisions for "Post-employment screening" where additional resource screening is requested by our Customers before resources are deployed at Regulated/Government Customer's premises. Some of the post-employment checks are residence checks, additional employment checks, criminal record checks, drug checks, identity checks, global database checks, additional educational checks.

The Capgemini's employees delivering services to our Customers from Capgemini's service delivery locations, employee activities are logged, monitored, reviewed as per procedures defined inline to the applicable laws and regulations without compromising the privacy of our employees. Capgemini's service delivery centers are ISO 27001:2013 certified and have administrative, technical and physical security controls implemented. Please refer to section 8.6.4 for more details on the high-level details of controls implemented.

All Capgemini employees' remote access to Capgemini network and our Customer's network is controlled by separate access credentials (as described in Capgemini's response to the question in section 8.5.3) and additional security access controls are implemented where needed through closed/secured workspaces as described in Capgemini's response to question in section 8.6.3.

AWS (IaaS, PaaS, SaaS)

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

Microsoft (IaaS, PaaS, SaaS)

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

Virtustream (IaaS)

Security procedures followed by Capgemini's Cloud Services Partner Virtustream: All Virtustream employees must pass a rigorous background investigation. At a high-level screening process include to review criminal records, validate past employment, perform drug screening, review Office of Foreign Asset Control, review sexual offender status. Additional security screening is available for the employees being onboard for government customers.

ServiceNow (SaaS)

Security procedures followed by Capgemini's Cloud Services Partner ServiceNow: ServiceNow has implemented capabilities to protect against insider threats and data exfiltration. They have a program called Controlled Access to provide access to customer instances and data is logged and monitored, and that sufficient preventative controls are in place to protect customer data.

BMC (SaaS)

Security procedures followed by Capgemini's Cloud Services Partner BMC: All BMC personnel complete pre-employment vetting procedures prior to joining BMC. This may include background checks, employment history, drug testing, OFAC List, criminal history, and more, in accordance with local laws and regulations. BMC OnDemand staff undergo annual corporate and role-based training, including compliance and ethics. Security rights for system access are controlled by BMC's SaaS security team, with authorization based on role-based permissions. BMC enforces a least privilege policy whereby only those personnel with the need to access data are permitted to access it. For example, some data is restricted to DBAs only, some to administrators only, some to BMC OnDemand Support personnel only.

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

Capgemini and Capgemini's Cloud Services Partners have appropriate security controls implemented to secure our Customer's data at rest and in transit to protect its confidentiality, integrity, and availability. Capgemini and Capgemini's Cloud Services Partners are ISO 27001:2013 certified and have security controls implemented at multiple layers. Capgemini is leveraging our Cloud Service Partners to deliver Cloud services to Purchasing Entities which have following security controls available to protect our Customer's data at rest in databases, servers, and data in transit using encryption and digital certificates. Standard and strict access controls are available to enable "need to know" and "segregation of duties" to enforce security controls on personnel to prevent intentional and

unintentional malicious/unauthorized activities. Capgemini and Capgemini's Cloud Services Partners leverage and comply with Cybersecurity control guidelines provided by NIST cybersecurity framework for our government Customers along with ISO 27001:2013 security controls and other applicable compliance requirements.

AWS (IaaS, PaaS, SaaS)

Security measures and controls available in Amazon Web Services for data security at rest and in transit:

AWS delivers a secure, scalable cloud computing platform with high availability, offering the flexibility to build a wide range of applications. AWS offers several options for encrypting data at rest, for an additional layer of security, ranging from completely automated AWS encryption solution to manual client-side options. Encryption requires 3 things - Data to encrypt, Encryption keys, Cryptographic algorithm method to encrypt the data.

AWS provides different models for Securing data at rest on the following parameters

- Encryption method
  - Encryption algorithm selection involves evaluating security, performance, and compliance requirements specific to your application
- Key Management Infrastructure (KMI)
  - KMI enables managing & protecting the encryption keys from unauthorized access
  - KMI provides
  - Storage layer that protects plain text keys
  - Management layer that authorizes key usage
- Hardware Security Module (HSM)
  - A common way to protect keys in a KMI is using HSM
  - An HSM is a dedicated storage and data processing device that performs cryptographic operations using keys on the device.
  - An HSM typically provides tamper evidence, or resistance, to protect keys from unauthorized use.
  - A software-based authorization layer controls who can administer the HSM and which users or applications can use which keys in the HSM

HTTPS with digital certificates can be used to encrypt data in transit, remote desktop, file transfer and remote shell protocols. Please refer to Attachment A – Amazon Web Services for more details and capabilities available for data security in Amazon Web Services for SaaS, PaaS and IaaS cloud service models.

Azure (IaaS, PaaS, SaaS)

Security measures and controls available in Microsoft Azure for data security at rest and in transit:

Azure Storage Service Encryption for Data at Rest helps customers protect their data to meet organizational security and compliance commitments. With this feature, Azure Storage automatically encrypts data before persisting it to Azure Storage and decrypts the data before retrieval. The handling of encryption, encryption at rest, decryption, and key management in Storage Service Encryption is transparent to users. All data written to Azure Storage is encrypted through 256-bit AES encryption, one of the strongest block ciphers available.

Storage Service Encryption is enabled for all new and existing storage accounts and cannot be disabled. Because customer data is secured by default, there is no need to modify code or applications to take advantage of Storage Service Encryption.

The feature automatically encrypts data in:

- Both performance tiers (Standard and Premium).
- Both deployment models (Azure Resource Manager and classic).
- All of the Azure Storage services (Blob storage, Queue storage, Table storage, and Azure Files).

Storage Service Encryption does not affect Azure Storage performance. Customers can use Microsoft-managed encryption keys with Storage Service Encryption or use their own encryption keys. HTTPS with SSL/TLS certificates can be used to encrypt data in transit, remote desktop, file transfer and remote shell protocols. These security services are offered by Microsoft Azure either with base services or with additional service cost. Please refer to Attachment B – Microsoft Azure for more details and capabilities for data security controls available in Microsoft Azure for SaaS, PaaS, and IaaS cloud service models.

Virtustream (IaaS)

Security measures and controls available in the Virtustream Private cloud for data security at rest and in transit:

Virtustream cloud is architected to the highest security and compliance standards and can significantly contribute as an enabler in achieving and maintaining compliance with FedRAMP, FISMA, SSAE16/SOC2/ISAE3402, PCI DSS 3.2, ISO 27001:2013, ISO 9001:2015, ISO 22301:2012, HIPAA/HITECH/HITRUST and other certification and compliance frameworks. Architected with security and compliance standards in mind, Virtustream storage cloud is certified against applicable HIPAA/HITECH, PCI, ISO, SOC2, Cloud Security Alliance (CSA)-STAR, and CJIS criteria. In each of security and compliance the Virtustream cloud solutions, both cloud IaaS and xStream software, follow the core tenants of cloud security:

- Compliance
- Trust
  - Intel Trusted Execution Technology (TXT)
  - Two-factor authentication
  - Encryption
- Visibility
  - Auditing
  - Alerting
  - SIEM
- Control
  - Role-based access control
  - Virtual firewall

Encryption a foundation of confidentiality – is supported by Virtustream Cloud services. Various third-party products are used to secure data at rest and in motion as well as to authenticate the components of Virtustream cloud technologies stack. Utilizing FIPS-compliant cryptographic technology, Virtustream can support all major encryption requirements for file systems, database, and network transport protection. Please refer to Attachment C – Virtustream Private Cloud for more details and capabilities available for Virtustream's Private Cloud services for IaaS cloud service model.

ServiceNow (SaaS)

Security measures and controls available in ServiceNow for data security at rest and in transit:

ServiceNow has implemented stringent security controls facing ability to prevent and mitigate security threats, protect data and help customers to comply with growing number of global security compliance mandates. To this end, ServiceNow has made significant investments in technology, processes, and expertise so that cloud services meet the most stringent of standards for security, availability, scalability, privacy, and compliance.

ServiceNow has engineered their cloud services, infrastructure that supports it, data encryption techniques, and security threat response processes so that customer's data is protected and secure at all times. As a customer, browser-based sessions to your ServiceNow cloud instance(s) are encrypted over the internet via Transport Layer Security (TLS) using AES-128 or AES-256 block ciphers. These ciphers are subject to the browser versions in use and may be influenced by customer's Internet proxy infrastructure. Customers can also force specific cipher suites via their own browser or proxy if desired. All end-user access to a ServiceNow instance is always automatically redirected to HTTPS if attempted over HTTP.

ServiceNow can apply encryption to integrations, such as LDAP and Web Services, as well as commonly used file transfer methods. In the case of an LDAPS over SSL connection, Customers can conveniently store certificates for specific LDAP servers within a ServiceNow instance for use in signing instance-bound Web Service requests. ServiceNow also supports certificate-based mutual Web Services security authentication with external endpoints for all ServiceNow instances. Data can be securely transferred to Customer's ServiceNow instances using pre-defined file transfer integration methods. Customers can also use clear text protocols such as FTP or HTTP to transfer data or support specific tasks, such as an approval or status request. Please refer to Attachment D – ServiceNow for more details and capabilities available for ServiceNow's services for SaaS cloud service model.

BMC (SaaS)

Security measures and controls available in BMC for data security at rest and in transit:

Remedy on Demand complies with FIPS 140-2 and is activated with the application of one of two levels of encryption:

- BMC Remedy Encryption Performance Security — When this option is activated, AR System encrypts network traffic by using AES CBC with a 128-bit key for data encryption and a 1024-bit modulus for the RSA key exchange, and SHA-1 for message authentication.
- BMC Remedy Encryption Premium Security — When this option is activated, AR System encrypts network traffic by using AES CBC with a 256-bit key for data encryption and a 2048-bit modulus for the RSA key exchange, and SHA-1 for message authentication.

Further, BMC designed Remedy on Demand so that all Purchasing Entity instances that are separate and apart from each other. Because the Purchasing Entity data is separated by the database, there is no possible way for a Purchasing Entity to gain access to another's data. This database separation of Purchasing Entities also facilitates the ability of BMC Support Personnel to maintain data separation through all of the activities required to maintain an optimally performing system. Please refer to Attachment E – BMC Remedy on Demand for more details and capabilities available for Remedy on-demand services for SaaS cloud service model.

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.

Capgemini will follow the process described in response to question in section 8.3.1. Below sections provide details about the high-level process and procedures followed by our Cloud Services Partner. Final process and the process will be agreed with Purchasing Entities in Master Services Agreement, Addendums, or SLA for individual contracts with Purchasing Entities.

AWS (IaaS, PaaS, SaaS)

Policies and procedures regarding notification available in Amazon Web Services are as below:

AWS Customers retain the responsibility to monitor their cloud 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. The staff operates 24x7x365 coverage to detect incidents and manage the impact to resolution. AWS utilizes a three-phased approach to manage incidents:

1. Activation and Notification Phase: Incidents for AWS begin with the detection of an event. This can come from several sources including:
    i. Metrics and alarms - AWS maintains an exceptional situational awareness capability, most issues are rapidly detected by 24x7x365 monitoring and alarming of real-time metrics and service dashboards. The majority of incidents are detected in this manner. AWS utilizes early indicator alarms to proactively identify issues that may ultimately impact Customers.
        a. Trouble ticket entered by an AWS employee
        b. Calls to the 24X7X365 technical support hotline. If the event meets incident criteria, then the relevant on-call support engineer will start an engagement utilizing AWS Event Management Tool system to start the engagement and page relevant program resolvers (e.g. Security team). The resolvers will perform an analysis of the incident to determine if additional resolvers should be engaged and to determine the approximate root cause.
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. Postmortem 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 SOC, PCI DSS, ISO 27001 and FedRAMP compliance.

Azure (IaaS, PaaS, SaaS)

Policies and procedures regarding notification available in Microsoft Azure are as below:

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

The Azure Incident Response process follows five main phases:

4. Identification – System and security alerts may be harvested, correlated, and analyzed. Events are investigated by Microsoft operational and security organizations. If an event indicates a security issue, the incident is assigned a severity classification and appropriately escalated within Microsoft. This escalation will include product, security, and engineering specialists.
5. Containment – The escalation team evaluates the scope and impact of an incident. The immediate priority of the escalation team is to make sure that the incident is contained and data is safe. The escalation team forms the response, performs appropriate testing, and implements changes. In the case where the in-depth investigation is required, content is collected from the subject systems using best-of-breed forensic software and industry leading practices.
6. Eradication – After the situation is contained, the escalation team moves toward eradicating any damage caused by the security breach and identifies the root cause of why the security issue occurred. If a vulnerability is determined, the escalation team reports the issue to product engineering.
7. Recovery – During recovery, software or configuration updates are applied to the system and services are returned to a full working capacity.
8. Lessons Learned – Each security incident is analyzed so that the appropriate mitigations applied to protect against future reoccurrence.

Virtustream (SaaS)

Policies and procedures regarding notification available in Virtustream Private Cloud services are as below:

Virtustream follows the Incident Response Plan (IRP) if there is an issue detected which caused a breach. Notification would be made to affected customers within the timeframes contractually agreed upon. Customers should extend their incident response within their tenant space to align with Virtustream IRP if an incident is detected. Virtustream will not take an action on an incident within the customer tenant space without input from the customer.

ServiceNow (SaaS)

Policies and procedures regarding notification available in ServiceNow are as below:

Unless notification is delayed by the actions or demands of a law enforcement agency, ServiceNow would report to the State or Cardholders the unauthorized acquisition, access, use, disclosure or destruction of Customer Data (a “Breach”) promptly following determination by ServiceNow that a Breach occurred. The initial report would be made to the security contact(s) designated by the Customer in ServiceNow’s customer support portal.

ServiceNow would take reasonable measures to promptly mitigate the cause of the Breach and shall take reasonable corrective measures to prevent future Breaches. As information is collected or otherwise becomes available to ServiceNow and unless prohibited by law, ServiceNow would provide

information regarding the nature and consequences of the Breach that are reasonably requested to allow States or Cardholders to notify affected individuals, other government agencies and/or credit bureaus. States and Cardholders are solely responsible for determining whether to notify impacted Data Subjects and for providing such notice, and for determining if regulatory bodies or enforcement commissions need to be notified of a Breach.

BMC (SaaS)

Policies and procedures regarding notification available in BMC Remedy on demand are as below:

Security event preparedness and response events are as described in the BMC Incident Response Plan. A security incident, or offense, is defined as an event generated by the Security Information and Event Management (SIEM) system based on correlating logs and events from incoming log sources. Customer notification for a security incident is as soon as possible (following assessment and validation of the event) but within 24 hours. Customers are notified via phone call or email, depending on the criticality of the event. BMC will disclose all information related to the security event that is related to individual Purchasing Entity's environment.

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 de-provisioning it before the Offeror is no longer contractually obligated to maintain the service, include planned and unplanned activities. An Offeror's response should include detail on how an Offeror maintains the 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.

The Purchasing Entity is primarily responsible for creation, usage, maintenance and destruction of its owned data when purchasing cloud storage capabilities from Capgemini Cloud Service providers.

We will work with the Purchasing Entities IT resources, Application owners, and stakeholders to create and manage detailed plans for service shutdown and data center exit procedures as applicable.

Final snapshot reports will be created and provided to appropriate Purchasing Entity stakeholders. This report will validate the following:

- Service Shutdown report
- Application data verified for correctness and decommissioning
- Metadata (App, version, decommission time/date, etc.) is documented
- A copy of backup data to be provided
- Application data categorization by data retention policies
- Application data stored is encrypted for security
- Data can be retrieved and restored in a timely manner

Capgemini will provide resources necessary for planning, labor, and technical expertise necessary for the virtual data center decommissioning or service decommissioning. Capgemini resources will manage the entire process and dispose of data information resources.

Through the value-added professional services catalog, Capgemini can provide a team of project managers and technical resources, through an agreed upon statement of work, that have abundant experience with highly complex migrations to private and public cloud architectures, including the necessary service shut down processes and decommissioning of an exit from any cloud-based virtual data centers and in addition legacy data centers that may have been used in a Hybrid architecture.

Program management support will be provided in helping drive the final consolidation/re-deployment plan. We will provide project manager's and technical resources experienced in data information management/decommissioning activities to support service shutdown in the cloud virtual data center, hybrid data centers, PaaS platforms, and SaaS applications with private information or configuration items which will coordinate the work with our Cloud Service providers.

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.

Capgemini will not have possession of data, the Purchasing Entity is responsible for destruction of its own data.

Through the value-added professional services catalog, Capgemini will have delegated administrative access to the Purchasing Entity data that resides in the cloud. Capgemini will define the correct processes to securely copy and transfer the requested server images and data from the cloud service provider over to media or transfer data via VPN connection to an onsite repository dictated by the Purchasing Entity.

8.8 (E) SERVICE OR DATA RECOVERY

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

a. Extended downtime.

b. Suffers an unrecoverable loss of data.

c. Offeror experiences a system failure.

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

e. Describe your Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

Capgemini, through an agreed upon statement of work, will work with each Purchasing Entity to create an architecture and solution which will address each specific BC/DR requirement, inclusive of RTO and RPO specifications.

Extended Downtime

Capgemini will configure the runbook to failover the production environment over to the Disaster Recovery (DR) site or region in case of a disaster or an unplanned outage is experienced. The key stakeholders will be informed of the failover and will enforce all change management rules and protocols in order to failover to the DR region and restore operations in accordance to the RTO and RPO's defined in the business continuity plan.

Further investigations will take place to determine the cause of failure, efforts will be made to resolve any issues with the production environment, after resolving the issues that caused the outage the production environment is reverted back to its original state.

Unrecoverable Losses of Data

Capgemini will implement daily backups to retrieve data in the event of any data loss is incurred. Data replication to other regions can be implemented in order to provide a shorter recovery time window allowing for data restoration.

System Failures

Cloud Service Providers in this offering, AWS, Microsoft Azure, ServiceNow, and BMC provide SLA's in an event of a system failure. AWS and Azure have multiple data centers in the US that provide data replication and redundancy in the event of a failure.

If a failure is incurred by the Cloud Service Provider, then a protocol is followed to restore service within their published SLA's. Capgemini will work directly with the Cloud Service Provider to understand the nature of the failure, communicate with the customers of any remediation actions to restore service.

Ability to Recover and Restore Data Within 4 Business Hours

Capgemini will plan the DR recovery plan and runbook to meet the requirement to recover and restore data within four business hours. The runbook will be configured and implemented in the Cloud Service Provider environment. The DR strategy can be defined by configuring data replication recovery time objectives (RTO) and recovery point objectives (RPO) within organizational limits.

For example, Azure Site Recovery helps provide business continuity by keeping business apps and workloads running during outages. Site Recovery replicates workloads running on physical and virtual machines (VMs) from a primary site to a secondary location. When an outage occurs at your primary site, you fail over to a secondary location, and access apps from there. After the primary location is running again, you can fail back to it. Site Recovery provides continuous replication for Azure VMs and VMware VMs, with replication frequency as low as 30 seconds. Similar DR strategy and replication can be configured in AWS environment.

Recovery Point Objectives and Recovery Time Objectives

Capgemini can configure the proper DR strategy and runbooks with the Cloud Service Providers in order to comply with the business continuity metrics, applying the desired RTOs and RPOs to provide the acceptable recovery times.

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.

Method of data backups

Capgemini will configure the Cloud Service Providers backup services to backup Virtual Machines, applications, and data on a daily weekly, monthly and quarterly basis. Specific methods, snapshots, the archive will be dependent on each solution.

Method of server image backups

A backup agent is installed on all Virtual Machines, then a snapshot of the OS is taken and backed up to a secondary storage location in Cloud Service Providers that can be used to restore the Virtual Machines image when required.

Digital location of backup storage (secondary storage, tape, etc.)

Cloud Service Providers provide capabilities to store backups in US data centers, up to six copies of backed up data can be stored in a geographical restore storage configuration.

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.

Cloud offerings included in this proposal from Capgemini's Cloud Service Partners provide standard encryption technologies to protect sensitive data for SaaS, PaaS and IaaS cloud service models.

For non-browser-based access, Secure Shell, Secure FTP, and IPSEC VPN services are available from Capgemini's Cloud Service Partners. Additionally, various security APIs from different cloud service Vendors and our Cloud Service Partners are available for Purchasing Entities to transact with cloud environments securely.

Data security provisions and capabilities available with Capgemini's Cloud Service Partners to protect customer's data are provided a response to a question in section 8.5.1. Capgemini follows "defense in depth" security architecture framework to protect our Customer's data and applications on Cloud. Leveraging our Cloud Service Partner's security and encryption capabilities, Capgemini can design and implement a layered security architecture with "encryption anywhere" approach to enable data security for Purchasing Entity's data wherever it goes. Capgemini applies a security first approach for data at rest and data in transit for data that will reside in the Cloud Provider's Public and Private Cloud environments, proposed cloud services and solutions account and enforce security standards and configurations leveraging Cloud Service Provider's capabilities.

AWS (IaaS, PaaS, SaaS)

Encryption technologies/controls available for the security of data in transit or data at rest in Amazon Web Services for SaaS, PaaS, and IaaS cloud service models.

AWS offers multiple security capabilities for data security to protect data at rest and data in transit. The table below summarizes the available options for encrypting data at rest across AWS. Capgemini recommends that Purchasing Entities determine which encryption and key management model is most appropriate for their data classifications in the context of the AWS service Purchasing Entities are using.

	Encryption Method and KMI			
	Model A		Model B	Model C
AWS Service	Client-Side Solutions Using Customer-Managed Keys	Client-Side Partner Solutions with KMI for Customer-Managed Keys	Client-Side Solutions for Customer-Managed Keys in AWS CloudHSM	Server-Side Encryption Using AWS-Managed Keys
Amazon S3	Bouncy Castle, OpenSSL, Amazon S3 encryption client in the AWS SDK for Java	SafeNet ProtectApp for Java	Custom Amazon VPC-EC2 application integrated with AWS CloudHSM client	Amazon S3 server side encryption, server-side encryption with customer provided keys, or server-side encryption with AWS Key

	Encryption Method and KMI			
	Model A		Model B	Model C
				Management Service
Amazon Glacier	N/A	N/A	Custom Amazon VPC-EC2 application integrated with AWS CloudHSM client	All data is automatically encrypted using server-side encryption
AWS Storage Gateway	Linux Block Level: ▪ Loop-AES, dm-crypt (with or without LUKS), and TrueCrypt Linux File System: ▪ eCryptfs and EncFs Windows Block Level: ▪ TrueCrypt Windows File System: ▪ BitLocker	Trend Micro SecureCloud, SafeNet StorageSecure	N/A	Amazon S3 server side encryption
Amazon EBS	Linux Block Level: ▪ Loop-AES, dmcrypt+ LUKS and TrueCrypt Linux File System: ▪ eCryptfs and EncFs Windows Block Level: ▪ TrueCrypt Windows File System: ▪ BitLocker, EFS	Trend Micro SecureCloud, SafeNet ProtectV	Custom Amazon VPC-EC2 application integrated with AWS CloudHSM client	Amazon EBS Encryption with AWS Key Management Service
Oracle on Amazon RDS	Bouncy Castle, OpenSSL	CipherCloud Database Gateway and Voltage SecureData	Custom Amazon VPC-EC2 application integrated with AWS CloudHSM client	Transparent Data Encryption (TDE) and Native Network Encryption (NNE) with optional Oracle Advanced Security

	Encryption Method and KMI			
	Model A		Model B	Model C
				license TDE for Microsoft SQL Server
Microsoft SQL Server on Amazon RDS	Bouncy Castle, OpenSSL	CipherCloud Database Gateway and Voltage SecureData	Custom Amazon VPC-EC2 application integrated with AWS CloudHSM client	N/A
Amazon Redshift	N/A	N/A	Encrypted Amazon Redshift clusters with your master key managed in AWS CloudHSM or on premises Safenet Luna HSM	Encrypted Amazon Redshift clusters with AWS-managed master key
Amazon EMR	eCryptfs		Custom Amazon VPC-EC2 application integrated with AWS CloudHSM client	S3DistCp using Amazon S3 server side encryption to protect persistently stored data

AWS offers data encryption services to protect data in transit using SSL or by using client-side encryption. Purchasing Entities can also leverage multiple other Vendors providing encryption services for data in transit available in AWS. The table below provides an overview of capabilities supported by AWS.

Data in transit	Secure protocol use	AWS
Web Access	HTTPS	Supported
File transfer	FTPS, SFTP, SCP, WebDAV over HTTPS	Supported
Remote Shell	SSH2 terminal	Supported
Remote desktop	radmin, RDP	Supported

Azure (IaaS, PaaS, SaaS)

Encryption technologies/controls available for the security of data in transit or data at rest in Microsoft Azure for SaaS, PaaS, and IaaS cloud service models.

Leveraging Capgemini's Cloud service Partner's services, Capgemini can configure the Azure data at rest security through the Azure cloud services that provide an Asymmetric AES256 data encryption algorithm. Data in transit can be secured by applying security rules and standards for all

communications over the Internet for data residing in the Azure Public Cloud. HTTPS with SSL certificates can be used to encrypt data in transit, remote desktop, file transfer and remote shell protocols. Furthermore, HTTPS with SSL certificates can be used to encrypt data in transit that will communicate with the Cloud Providers Public Cloud environments, remote desktop, file transfer and remote shell protocols to access data hosted by the Cloud Providers can be enabled. The protocols that can be leveraged are defined in the table below.

Data in transit	Secure protocol use	Azure
Web Access	HTTPS	Supported
File transfer	FTPS, SFTP, SCP, WebDAV over HTTPS	Supported
Remote Shell	SSH2 terminal	Supported
Remote desktop	radmin, RDP	Supported

Virtustream (IaaS)

Encryption technologies/controls available for the security of data in transit or data at rest in Virtustream Private Cloud for IaaS cloud service model.

Virtustream's cloud solutions are designed to meet the exacting requirements of the complex enterprise, government customers. In each area of security and compliance with the Virtustream cloud solutions, follow the core tenets of cloud security: Trust, Visibility, Control, Compliance, Defense in Depth. High efficiency encryption technologies are available to protect the entire data lifecycle and are utilized throughout the Virtustream cloud environment to secure: data at rest including the entire virtual machine and its related data storage, transactional databases, data in the archive, data in transit, and the authentication of the various components of the cloud stack. Virtustream can also maintain encryption and policies as data is moved and replicated and support a wide variety of integrated key management options from sole responsibility to any flavor of shared responsibility.

In addition to the standard data in transit security measures described, Virtustream Storage Cloud is independently audited as part of an ongoing third-party certification and self-assessment program. The combination of Dell EMC Data Domain and/or Data Protection Suite software and Virtustream Storage Cloud enables secure, efficient, and cost-effective long-term retention (LTR) of backups with a subscription-based object storage service.

Today, Virtustream Storage Cloud is certified against applicable HIPAA, PCI, ISO, SOC 2, and Cloud Security Alliance (CSA)-STAR, and CJIS criteria. All security reports, whether independent or self-asserted, are available to customers under NDA who request it through their business development contact.

ServiceNow (SaaS)

Encryption technologies/controls available for the security of data in transit or data at rest in ServiceNow for SaaS cloud service model.

Browser-based sessions to Purchasing Entity's ServiceNow cloud instance(s) are encrypted over the internet via Transport Layer Security (TLS) using AES-128 or AES-256 block ciphers. These ciphers are subject to the browser versions in use and may be influenced by your Internet proxy infrastructure.

ServiceNow can apply encryption to integrations, such as LDAP and Web Services, as well as commonly used file transfer methods. In the case of an LDAPS over SSL connection, you can conveniently store certificates for specific LDAP servers within a ServiceNow instance for use in signing instance-bound Web Service requests. We also support certificate-based mutual Web Services security authentication

with external endpoints for all ServiceNow instances. Data can be securely transferred to your ServiceNow instances using pre-defined file transfer integration methods. You can also use clear text protocols such as FTP or HTTP to transfer data or support specific tasks, such as an approval or status request.

ServiceNow provides the option to perform column-level encryption on fields and attachments. This feature is available to both delivered services as well as custom built applications developed on the Now Platform. ServiceNow supports AES-128, AES-256, and 3DES encryption algorithms, and apply your choice to encrypt data. To mitigate the possible compromise of encrypted customer data, they re-encrypt (wrap) your keys with a secondary key. In some cases, data stored in fields and attachments that is encrypted cannot be searched for or reported on.

Everything inside the co-location spaces is owned, operated, and managed by ServiceNow. This includes the management of hard drives and server hardware. All hard drives are sanitized prior to leaving our private cages (per NIST 800-88 guidelines) so that Purchasing Entity data is appropriately handled and protected. You can choose to further mitigate data exposure caused by the loss or theft of storage devices with AES-256 full-disk encryption of your data at rest. Full-disk encryption is available at additional cost.

BMC (SaaS)

Encryption technologies/controls available for the security of data in transit or data at rest in BMC Remedy on Demand for SaaS cloud service model.

Remedy on Demand complies with FIPS 140-2 and is activated with the application of one of two levels of encryption:

- BMC Remedy Encryption Performance Security — When this option is activated, AR System encrypts network traffic by using AES CBC with a 128-bit key for data encryption and a 1024-bit modulus for the RSA key exchange, and SHA-1 for message authentication.
- BMC Remedy Encryption Premium Security — When this option is activated, Action Request (AR) System encrypts network traffic by using AES CBC with a 256-bit key for data encryption and a 2048-bit modulus for the RSA key exchange, and SHA-1 for message authentication.

BMC offers two options for encryption of data at rest:

9. The entire database can be encrypted at rest upon request. With the exception of customers in the FedRAMP data center, encryption is not performed by default, so Purchasing Entites must notify BMC SaaS Operations of this requirement in advance of system provisioning. BMC utilizes Microsoft's Transparent Data Encryption (TDE) which performs real time I/O encryption and decryption of the data and log files utilizing a symmetric database encryption key (DEK).
10. Purchasing Entities may encrypt only certain character fields. This option utilizes AES 128-bit encryption.

Purchasing Entities need to consider that encrypted fields are not searchable, so option 2 has to be used intelligently. For option 1, data in use is not data at rest, and therefore a field tagged in a global search index would be active and searchable (assuming the field-level encryption flag is not also active). A customer's specific use case(s) would determine whether they need enterprise encryption (requires a BMC-managed key for the database) or field-level encryption (requires a customer-managed key for the application).

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.

Capgemini will sign, after an appropriate review, a mutually agreeable business associate agreement (BAA) or any other agreement that may be necessary to protect a Purchasing Entity's data.

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. The offeror shall not use the government data or government-related data for any other purpose including but not limited to data mining. The offeror or its subcontractors shall not resell or otherwise redistribute information gained from its access to the data received as a result of this RFP.

Capgemini and Capgemini's Cloud Service Partners will only use data for the purposes defined in the Master Services Agreement, a participating addendum, or a related SLA. Please see our response to question 8.5.1 wherein Capgemini described the measures available to protect Purchasing Entity's data.

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.

Capgemini will work each Purchasing Entity to negotiate appropriate SLAs for Capgemini Cloud offerings provided directly or through Capgemini's Cloud Service providers.

> *The Cloud Service Providers typically offer SLA's that are nonnegotiable for public and community-based services and Hybrid and Private services can typically be negotiated depending on the measurement and severity.*

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.

The Cloud Service Provider SLA's and Capgemini's iPaaS offering will be included as the following Attachments:

- Amazon SLAs
- Azure SLAs
- BMC SLAs
- ServiceNow SLAs
- Virtustream SLAs
- Capgemini iPaaS

AWS (IaaS, PaaS, SaaS)

Capgemini resells the cloud services provided by AWS and Azure. Each Cloud Service Provider publishes their SLA's, the cloud SLA's are policies governing service level objectives by the cloud service providers. Capgemini can assist customers through an agreed upon statement of work to configure the IaaS, PAAS, SaaS cloud services in accordance with the cloud service providers SLA's.

Different types of cloud SLA's are IaaS SLAs (with a distinction between compute and storage services), PaaS SLAs, and SaaS SLAs. In general, service level objectives vary across service models.

The following table provides a list of the common AWS cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for AWS Cloud Services will be provided as an attachment.

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
EC2 Instances	< 99.95%	10%	100% – EC2 or	"Region Unavailable" and "Region

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
& EBS Volumes	< 99%	30%	Amazon EBS was in a state of "Region Unavailable"	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.
S3 (storage services)	< 99.9%	10%	100% – Error Rates	"Error Rate" means: Total No. of internal server errors / Total number of requests for application request during that five minute period
	< 99%	30%		
S3 (storage services) Standard – Infrequent Access (Standard-IA)	< 99%	10%	100% – Error Rates	
	< 98%	30%		
CloudFront (CDN service)	< 99.9%	10%	100% – Error Rates	
	< 99%	25%		
	< 99%	25%		
Route 53 (DNS service)	5 – 30 minutes	1 day	Duration was not 100% available	
	31 mins – 4 hours	7 days		
	More than 4 hours	30 days		

Azure (IaaS, PaaS, SaaS)

The following table provides a list of the common Azure cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for Azure Cloud Services will be provided as an attachment.

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Virtual Machines	< 99.99%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.
	< 99%	25%		
	< 99%	25%		
Storage Service LRS, ZRS, GRS, and RA-GRS	< 99.9%	10%	100% – Average Error Rate	"Average Error Rate" for a billing month is the sum of Error Rates for each hour in the billing month
	< 99%	25%		

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
(write requests) Accounts				divided by the total number of hours in the billing month.
Storage Service RA-GRS (read requests) Accounts	< 99.99%	10%	100% – Average Error Rate	GRS – Geographically Redundant Storage LRS – Locally Redundant Storage RA-GRS – Read Access Geographically Redundant
	< 99%	25%		
Storage Service – LRS, GRS and RA-GRS (write requests) Blob Storage Accounts (Cool Access Tier)	< 99%	10%	100% – Average Error Rate	"Cool Access Tier" is an attribute of a Blob Storage Account indicating that the data in the account is infrequently accessed and has a lower availability service level than data in other access tiers.
	< 98%	25%		
Storage Service – RA-GRS (read requests) Blob Storage Accounts (Cool Access Tier)	< 99.9%	10%	100% – Average Error Rate	
	< 98%	25%		
	< 99%	25%		
	< 99%	25%		
Traffic Manager (DNS service)	< 99.95%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated Deployment Minutes, across all Profiles deployed by you in a given Microsoft Azure subscription, during which the Profile is unavailable.
	< 99%	25%		
CDN Service	< 99.9%	10%	Downtime: To assess Downtime, Microsoft will review data from any commercially reasonable independent measurement system used by the customer.	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).
	< 99.5%	25%		
	Encrypted: >6 hours Monthly RTO	100%		
	< 99%	25%		

Virtustream (IaaS)

Virtustream offers 99.999% availability SLA on any cloud workloads designed and deployed into a high availability (H/A) solution, with H/A available at all layers of the cloud architecture.

Performance assurance

Resource modeling and predictive analytics enable proactive infrastructure optimization and industry-leading SLAs

- Application performance SLAs guarantee IOps vs. simple availability
- 99.999% availability SLA for H/A workloads
- Storage performance SLAs
- Application performance SLA for SAP
- Managed services response time SLA
- Disaster Recovery SLA

Figure 15: Performance Assurance

Availability & resiliency

High availability at every level of the architecture

- Up to 99.999% availability for cloud deployments
- Added storage replication and DR protection with dual cloud/ hybrid cloud deployments
- Geographically dispersed data centers
- Storage-based replication
  - 15 Min RPO – Non HANA
  - 30 Min RPO – HANA
  - 1 Hour RTO – Non HANA
  - 2 Hour RTO - HANA

Feature/ Offering	Enterprise Basic µVM	Enterprise Core µVM	Studio Basic µVM	Studio Core µVM
Storage replication	NO	YES	NO	YES
DR Protection	NO	YES	NO	YES
Site Resources	Single Site	Dual Sites	Single Site	Dual Sites
Availability	99.5%	99.999%	99.4%	99.9%

90% of Virtustream Customers subscribe to DR Services for Production workloads

Figure 16: Availability & resiliency

All Incidents that are reported to Virtustream or that Virtustream otherwise becomes aware of will initially be assigned a priority by Virtustream as set forth below. Internal escalation for Incidents resources shall be determined by Virtustream based on the priority level assigned to the Incident by Virtustream. The priority/severity level may be adjusted as agreed to by the parties.

Priority / Severity	Definition	Time to Respond [Note 1]	Customer Communication Interval	Level of Effort
1	**A major part of the Services is unavailable / not operating correctly, affecting multiple users. No workarounds are in place, and business operations are not possible.** **OR** Incident has a critical impact on the business (e.g., loss of the Exchange production server impacting all users).	30 minutes	Every 30 minutes	Immediate and continuous effort until the issue is resolved or a workaround is developed
2	**Part of the Services is unavailable / not operating correctly, affecting users in a single function. No workarounds are in place, and business operations in this function are not possible / severely impacted.** **OR** Incident has a serious impact on part of the business (e.g., a configuration change is impacting a small subset of users).	60 minutes	Every 60 minutes	Continuous effort until the issue is resolved or a workaround is developed
3	**Part of the Services is unavailable / not operating correctly, affecting users in a single function. Workarounds are in place, but business operations are impacted, although not severely.** **OR** Incident has a temporary impact on users and is non-critical or is a development issue (e.g., email is slow to deliver).	4 hours	Updates provided as available	Work until the issue is resolved or a workaround is developed during business hours
4	**The incident that is causing inconvenience to the business, but not impacting operations.** **OR** Incident has a minor impact on users or business, or issue is a request for further information.	1 US business day	Not applicable	Will be addressed during the next general update to the services

Virtustream shall provide Customer with the Service Level Credit if the Services fail to satisfy any of the Service Levels set out herein. The Service Level shall commence thirty (30) days following commencement of the Steady State phase of the SOW. Each of the Service Level Credits shall be based on the fees paid for the applicable service, as set out in this SOW. The aggregate Service Level Credits for all Service Levels in any month shall not exceed 15% of the total monthly recurring charges

("MRCs"). The Service Level Credit shall be Customer's sole and exclusive remedy and Virtustream's sole and exclusive liability for Unscheduled Downtime.

µVM Availability on Enterprise Core µVM	µVM Availability on Enterprise Basic µVM	µVM Availability on DMZ Core VM	µVM Availability on DMZ Basic VM	Service Level Credit [Note 2]
99.95% – 99.999%	99% - 99.5%	99.5% - <99.9%	99% - <99.4%	1%
99.5% - 99.94%	98% - 98.99%	98% - <99.5%	98% - <99%	3%
95% - 99.4%	95% - 97.99%	95% - <98%	95% - <98%	5%
90% - 94.99%	90% - 94.99%	90% - <95%	90% - <95%	10%
Below 90%	Below 90%	Below 90%	Below 90%	15%

ServiceNow (SaaS)

As a cloud-based solution, ServiceNow customers can connect to the application via any supported web browser. ServiceNow is fully responsible for the availability of the application. ServiceNow Operations organization provide 24 hour x 7 days per week operation and management of the ServiceNow NOW Platform of applications.

ServiceNow provides 99.8% Service Availability commitment for Production Environments. Following is a sample of the Technical Support Response targets maintained by ServiceNow.

CUSTOMER SUPPORT RESPONSE TARGETS		
Severity	**Time Schedule**	**Initial Response Targets**
P1	Continuously, 24 hours per day, 7 days per week	30 Minutes
P2	Continuously, but not necessarily 24 hours per day, 7 days per week	2 Hours
P3	As appropriate during normal business hours	1 Business Day
P4	Varies	N/A

Service Credits can be issued

BMC (SaaS)

As a cloud-based solution, Remedy on Demand customers is able to connect to the application via any supported web browser. BMC is fully responsible for the availability of the application. The BMC Operations organization provides 24 hour x 7 days per week operation and management of the BMC Remedy On Demand service.

Remedy OnDemand provides 99.9% Service Availability commitment for Production Environments. Following is a sample of the Technical Support Response targets maintained by BMC Operations.

TECHNICAL SUPPORT RESPONSE TARGETS		
Severity	**Time Schedule**	**Initial Response Targets**
S1	24 x 7 (Includes published holidays)	15 Minutes
S2	Local Business Hours: 7am-7pm, M-F (Excludes published holidays)	30 Minutes
S3		4 Business Hours
S4		16 Business Hours

Service Credits are not provided for missed Targets.

Capgemini iPaaS Service Offering (PaaS)

- Prioritise all Support Requests based on its reasonable assessment of the severity level of the Incident reported
- Respond to all Support Requests in accordance with the Response Times specified in the table set out below:

Severity level of Incident	Definition	Service Level Response
1 **Critical Extensive**	Business Critical: An outage of a Small, Medium or Large Platform Instance leading to total loss of service of all APIs and/or integrations deployed on the Enterprise iPaaS Platform.	2 hours
2 **High** **Significant**	Business Impact: Impact on key functionality and/or performance degradation of the Enterprise iPaaS Platform but Purchasing Entity critical business functions are still operational. No acceptable workaround is available.	6 hours
3 **Medium** **Moderate**	Operational Impact: Moderate impact on usage of Enterprise iPaaS Platform but remains operational. A workaround is available to improve the situation until the issue is fully resolved.	12 hours
4 **Low** **Minor**	Minor Impact: Minor impact on usage of Enterprise iPaaS Platform. Includes minor, cosmetic, or documentation related issues. No impact on Enterprise iPaaS Platform features.	24 hours

On receipt of a Support Request by the Supplier iPaaS Support team, the validity and severity of the Incident are assessed, with a response provided within the above timescales to either accept or return the Incident and amend the Incident severity in line with the definitions if required.

The maximum number of Support Requests which the Purchasing Entity can raise in each month is as follows:

Number of Platform Instances	Purchasing Entity Incident Allowance Per Month
1-4	60
5-10	120
11-15	180

Actual resolution time will depend on the nature of the fault underlying the Incident.

8.11 (E) DATA DISPOSAL

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

The Purchasing Entity is primarily responsible for creation, usage, maintenance and destruction of its owned data when purchasing cloud storage capabilities from Capgemini Cloud Service providers.

Capgemini understands that disposal of data assets must be authorized by the Purchasing Entities, Capgemini can engage through an agreed statement of work to properly manage the disposal process by working with the Cloud Service Providers. Capgemini can provide asset reporting, including Confirmation of Disposal and Data Security, and reports to the Purchasing Entities for the authorized disposal of data.

AWS (IaaS, PaaS, SaaS)

AWS customers manage access to their data and content, customers control user access to AWS services and resources residing in the Public or Community Cloud. AWS provides an advanced set of access, encryption, and logging features to help customers do this effectively. AWS does not access or use custom content for any purpose without customer consent. Customers can place their data in any AWS region. Customers at any time can download and delete their data.

Azure (IaaS, PaaS, SaaS)

Microsoft is governed by strict standards and follows specific processes for removing cloud customer data from systems under their control, overwriting storage resources before reuse, and purging or destroying decommissioned hardware. When customer data is hosted in the Azure multitenant cloud services, Microsoft takes careful measures to logically separate customer data. This helps prevent one customer's data from leaking into that of another customer, which also helps to block any customer from accessing another customer's deleted data. Customers manage their data and can place their data in any Azure region. Customers can also at any time download and delete their data. If a subscription has expired or is terminated, all associated customer data (including storage accounts) is held for a 90 day period before it is deleted in order to recover from accidental subscription cancellation.

Virtustream (IaaS)

Virtustream customers manage access to their data and content, customers control user access to Virtustream services and resources residing in the Private Cloud. Virtustream provides access, encryption, and logging features to help customers do this effectively. Virtustream does not access or use custom content for any purpose without customer consent. Customers can place their data in any Virtustream region. Customers at any time can download and delete their data.

Capgemini through the agreed statement of work can assist Purchasing Entities to erase all data from the Cloud Service Provider's storage repository for any IaaS, PAAS or SaaS data residing in the Cloud Providers environments. Capgemini will inform the Purchasing Entity that data has been erased as instructed. Capgemini does not have access to the physical media or storage once the subscription is terminated with the Cloud Service Provider.

ServiceNow (SaaS)

All Customer data is hosted on solid-state or mechanical disks within ServiceNow's colocation spaces. No tapes or other forms of removable media are used to provide the service, including for backups. When functional storage devices reach their end-of-life or get reassigned to new customers, these are logically shredded using a process based on guidance from the U.S. National Institute of Standards and Technology (NIST) 800-88. Records of the destruction are maintained.

BMC (SaaS)

Remedy OnDemand customers retain ownership of their data at all times. Should a customer decide to leave the service, BMC initiate's its Customer Departure process that includes the return of the Customer's data and destruction of the data from storage and backup media.

BMC provides a file containing the Customer data in comma separated value (.csv) or database backup format upon customer request. Customer data is then destroyed via destruction of database

encryption keys and data is overwritten with binary zeroes. Confirmation of data destruction is performed and logged in the Customer service records.

8.12 (E) PERFORMANCE MEASURES AND REPORTING

8.12.1 Describe your ability to guarantee reliability and uptime greater than 99.5%. Additional points will be awarded for 99.9% or greater availability.

AWS (IaaS, PaaS, SaaS)

Capgemini resells the cloud services provided by AWS and Azure. Each Cloud Service Provider publishes their SLA's, the cloud SLA's are policies governing service level objectives by the cloud service providers. Capgemini can assist customers through an agreed upon statement of work to configure the IaaS, PAAS, SaaS cloud services in accordance with the cloud service providers SLA's.

Different types of cloud SLA's are IaaS SLAs (with a distinction between compute and storage services), PaaS SLAs, and SaaS SLAs. In general, service level objectives vary across service models.

The following table provides a list of the common AWS cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for AWS Cloud Services will be provided as an attachment.

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
EC2 Instances & EBS Volumes	< 99.95%	10%	100% – EC2 or Amazon EBS was in a state of "Region Unavailable"	"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.
	< 99%	30%		
S3 (storage services)	< 99.9%	10%	100% – Error Rates	"Error Rate" means: Total No. of internal server errors / Total number of requests for application request during that five minute period
	< 99%	30%		
S3 (storage services) Standard – Infrequent Access (Standard-IA)	< 99%	10%	100% – Error Rates	
	< 98%	30%		
CloudFront (CDN service)	< 99.9%	10%	100% – Error Rates	
	< 99%	25%		
	< 99%	25%		
Route 53 (DNS	5 – 30 minutes	1 day	Duration was not	

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
service)	31 mins – 4 hours	7 days	100% available	
	More than 4 hours	30 days		

Azure (IaaS, PaaS, SaaS)

The following table provides a list of the common Azure cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for Azure Cloud Services will be provided as an attachment.

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Virtual Machines	< 99.99%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.
	< 99%	25%		
	< 99%	25%		
Storage Service LRS, ZRS, GRS, and RA-GRS (write requests) Accounts	< 99.9%	10%	100% – Average Error Rate	"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. GRS – Geographically Redundant Storage LRS – Locally Redundant Storage RA-GRS – Read Access Geographically Redundant "Cool Access Tier" is an attribute of a Blob Storage Account indicating that the data in the account is infrequently accessed and has a lower availability service level than data in other access tiers.
	< 99%	25%		
Storage Service RA-GRS (read requests) Accounts	< 99.99%	10%	100% – Average Error Rate	
	< 99%	25%		
Storage Service – LRS, GRS and RA-GRS (write requests) Blob Storage Accounts (Cool Access Tier)	< 99%	10%	100% – Average Error Rate	
	< 98%	25%		
Storage Service – RA-GRS (read requests) Blob Storage Accounts (Cool Access Tier)	< 99.9%	10%	100% – Average Error Rate	
	< 98%	25%		
	< 99%	25%		
	< 99%	25%		
Traffic	< 99.95%	10%	((Maximum Available	Downtime: The total accumulated

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Manager (DNS service)	< 99%	25%	Minutes – Downtime) / Maximum Available Minutes) x 100	Deployment Minutes, across all Profiles deployed by you in a given Microsoft Azure subscription, during which the Profile is unavailable.
CDN Service	< 99.9%	10%	Downtime: To assess Downtime, Microsoft will review data from any commercially reasonable independent measurement system used by the customer.	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).
	< 99.5%	25%		
	Encrypted: >6 hours Monthly RTO	100%		
	< 99%	25%		

Virtustream (IaaS)

Virtustream offers 99.999% availability SLA on any cloud workloads designed and deployed into a high availability (H/A) solution, with H/A available at all layers of the cloud architecture.

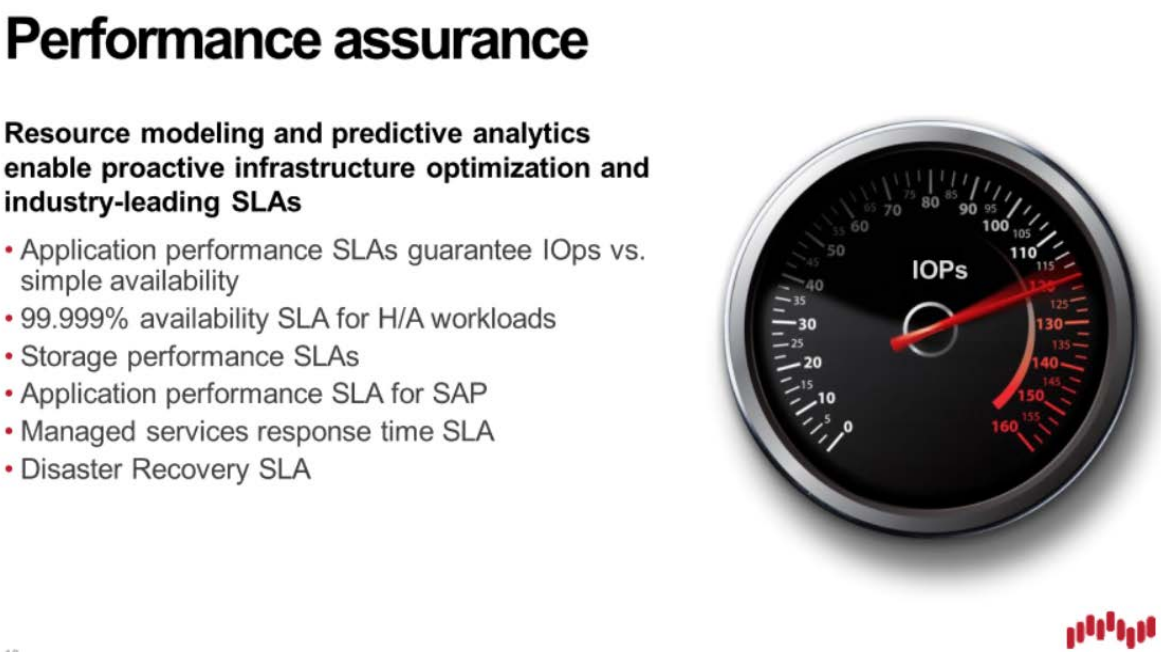

Figure 17: Performance Assurance

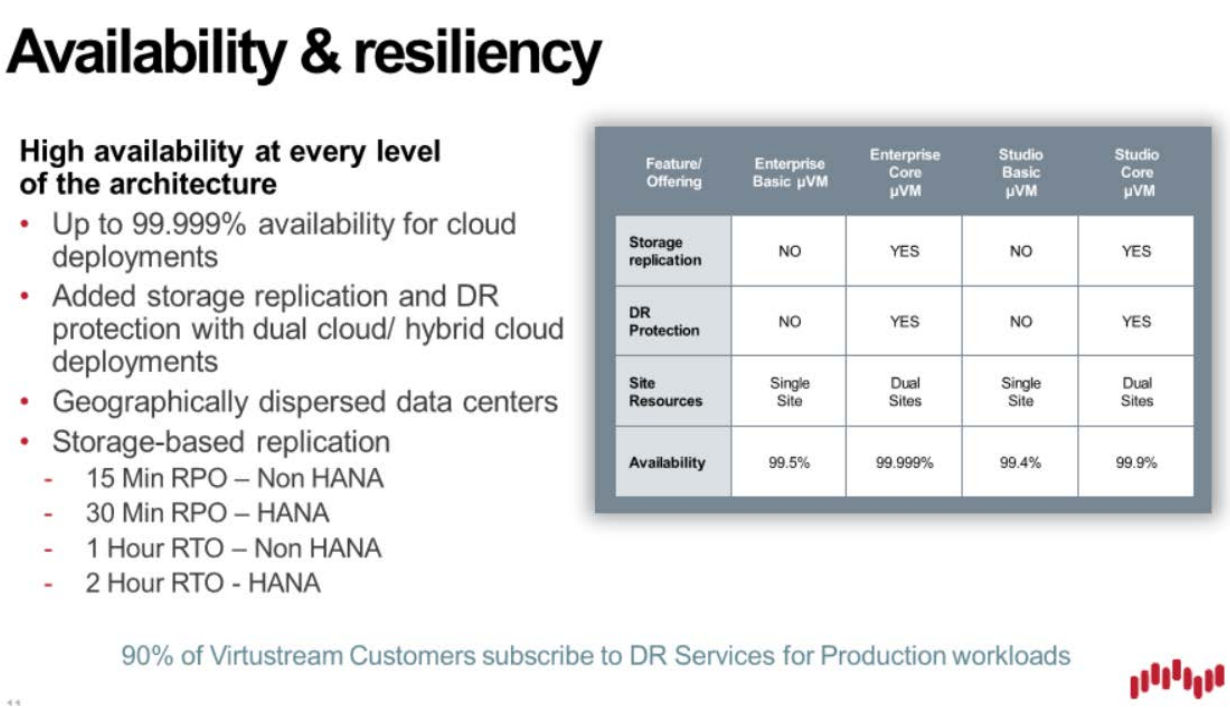

Figure 18: Availability and Resiliency

BMC (SaaS)

BMC OnDemand services carry a 99.9% service availability commitment for production environments across all “Remedy on Demand” subscription services.

ServiceNow (SaaS)

ServiceNow operates under a 99.8% service availability commitment.

BMC and ServiceNow can offer these uptime commitments because of the highly available, redundant private cloud architecture on which they have built their services.

Capgemini Enterprise iPaaS (PaaS)

Subject to the exclusions set out in Section Subscription and Support Services Exclusions, the Supplier shall use reasonable commercial endeavors (subject to planned outages detailed below) so that the Enterprise iPaaS Platform is Available 99.95% of the time during the Support Hours.

The actual Availability shall be stated as a percentage and calculated as follows:

$$\left(\frac{(x-y)}{x}\right) \times 100$$

where:

x = the total number of minutes during the Measurement Period less the Scheduled Downtime.

y = the total number of minutes during the Measurement Period where the Enterprise iPaaS Platform is Unavailable (other than unavailability due to Scheduled Downtime or an Exclusion).

Any planned outages as described in section Patching, Maintenance and Upgrades. will be excluded from the above instances and Capgemini will be relieved of its service level obligations during any instance of Force Majeure.

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

Each Cloud Service Provider publishes their SLA's. We have included the SLAs as separate Attachments to this response titled as follows:

- Amazon SLAs
- Azure SLAs
- BMC SLAs
- ServiceNow SLAs
- Virtustream SLAs
- Capgemini iPaaS SLAs

AWS (IaaS, PaaS, SaaS)

The services offered by AWS and Azure are organized in part by SLA, which means to attain a specific SLA outcome the Purchasing Entity will want to procure services that match the specific SLA requirements. Capgemini can, through an agreed upon statement of work, assist the Purchasing Entities with procuring and then with configuring those cloud services that meet SLA requirements.

Specific SLA metrics differ based on the types of cloud services; IaaS (with a distinction between compute and storage services), PaaS, and SaaS.

The following table provides a list of the common AWS cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for AWS Cloud Services will be provided as an attachment.

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
EC2 Instances & EBS Volumes	< 99.95%	10%	100% – EC2 or Amazon EBS was in a state of "Region Unavailable"	"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.
	< 99%	30%		
S3 (storage services)	< 99.9%	10%	100% – Error Rates	"Error Rate" means: Total No. of internal server errors / Total number of requests for application request during that five minute period
	< 99%	30%		
S3 (storage services) Standard – Infrequent Access (Standard-IA)	< 99%	10%	100% – Error Rates	
	< 98%	30%		
CloudFront (CDN service)	< 99.9%	10%	100% – Error Rates	
	< 99%	25%		
	< 99%	25%		

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Route 53 (DNS service)	5 – 30 minutes	1 day	Duration was not 100% available	
	31 mins – 4 hours	7 days		
	More than 4 hours	30 days		

Azure (IaaS, PaaS, SaaS)

The following table provides a list of the common Azure cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for Azure Cloud Services will be provided as an attachment.

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Virtual Machines	< 99.99%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.
	< 99%	25%		
	< 99%	25%		
Storage Service LRS, ZRS, GRS, and RA-GRS (write requests) Accounts	< 99.9%	10%	100% – Average Error Rate	"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. GRS – Geographically Redundant Storage LRS – Locally Redundant Storage RA-GRS – Read Access Geographically Redundant "Cool Access Tier" is an attribute of a Blob Storage Account indicating that the data in the account is infrequently accessed and has a lower availability service level than data in other access tiers.
	< 99%	25%		
Storage Service RA-GRS (read requests) Accounts	< 99.99%	10%	100% – Average Error Rate	
	< 99%	25%		
Storage Service – LRS, GRS and RA-GRS (write requests) Blob Storage Accounts (Cool Access Tier)	< 99%	10%	100% – Average Error Rate	
	< 98%	25%		
Storage Service – RA-GRS (read requests) Blob Storage Accounts (Cool Access Tier)	< 99.9%	10%	100% – Average Error Rate	
	< 98%	25%		
	< 99%	25%		
	< 99%	25%		

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Traffic Manager (DNS service)	< 99.95%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated Deployment Minutes, across all Profiles deployed by you in a given Microsoft Azure subscription, during which the Profile is unavailable.
	< 99%	25%		
CDN Service	< 99.9%	10%	Downtime: To assess Downtime, Microsoft will review data from any commercially reasonable independent measurement system used by the customer.	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).
	< 99.5%	25%		
	Encrypted: >6 hours Monthly RTO	100%		
	< 99%	25%		

Virtustream (IaaS)

All Incidents that are reported to Virtustream or that Virtustream otherwise becomes aware of will initially be assigned a priority by Virtustream as set forth below. Internal escalation for Incidents resources shall be determined by Virtustream based on the priority level assigned to the Incident by Virtustream. The priority/severity level may be adjusted as agreed to by the parties.

Priority / Severity	Definition	Time to Respond [Note 1]	Customer Communication Interval	Level of Effort
1	**A major part of the Services is unavailable / not operating correctly, affecting multiple users. No workarounds are in place, and business operations are not possible.** **OR** Incident has a critical impact on the business (e.g., loss of the Exchange production server impacting all users).	30 minutes	Every 30 minutes	Immediate and continuous effort until the issue is resolved or a workaround is developed
2	**Part of the Services is unavailable / not operating correctly, affecting users in a single function. No workarounds are in place, and business operations in this function are not possible / severely impacted.** **OR** Incident has a serious impact on part of the business (e.g., a configuration change is impacting a small subset of users).	60 minutes	Every 60 minutes	Continuous effort until the issue is resolved or a workaround is developed
3	**Part of the Services is unavailable / not operating**	4 hours	Updates provided	Work until the issue is resolved or a

Priority / Severity	Definition	Time to Respond [Note 1]	Customer Communication Interval	Level of Effort
	correctly, affecting users in a single function. Workarounds are in place, but business operations are impacted, although not severely. **OR** Incident has a temporary impact on users and is non-critical or is a development issue (e.g., email is slow to deliver).		as available	workaround is developed during business hours
4	**The incident that is causing inconvenience to the business, but not impacting operations.** **OR** Incident has a minor impact on users or business, or issue is a request for further information.	1 US business day	Not applicable	Will be addressed during the next general update to the services

BMC (SaaS)

BMC OnDemand services carry a 99.9% service availability commitment for production environments across all “Remedy on Demand” subscription services.

The following table defines BMC’s severity definitions for classifying reported incidents.

Severity	Severity Criteria
S1	Critical service impact—The issue critically affects a primary business service, major application, or mission-critical system. Customer resources should be available and willing to work with BMC 24 hours a day, 7 days a week to resolve the issue. The following characteristics describe a Severity 1 issue: ▪ Business service is not operational ▪ Production system crashes ▪ Data integrity is at risk ▪ Production backup and recovery operations fail
S2	Significant service or implementation impact—A business service, major application, or system is seriously affected, or implementation is stopped. No acceptable workaround is available.
S3	Moderate service impact—The business service, major application, or system is moderately impacted, no data has been lost, and the business service, application, or system is still functioning. The issue may be temporarily circumvented by using an available workaround.
S4	No Service Impact—Noncritical issues, general questions, enhancement requests, or documentation issues

ServiceNow (SaaS)

ServiceNow holds out a 99.8% service availability. If Customer’s production instances of the Subscription Service fall below the Availability SLA of ninety-nine and eight-tenths percent (99.8%) during a calendar month, the Participating Entity’s options for failure of the Subscription Service to meet the Availability SLAs is either (1) to request that the affected Subscription Term be extended for the number of minutes the Subscription Service was not Available in the month, or (2) to request that

ServiceNow issue a service credit to Customer for the dollar value of the number of minutes the Subscription Service was not Available in the month as calculated on a per minute rate for ServiceNow charges to Customer for the month. The credit would be applied to the next invoice for subscription fees.

ServiceNow's definition of Severity

Priority	Definition
P1	Any defect that causes an instance to be unavailable.
P2	Any defect that causes a mission-critical function to fail.
P3	Any request or defect that is significantly impeding work or progress.
P4	Any request or defect that is important but not significantly impeding work or progress.

Capgemini iPaaS (PaaS)

Capgemini shall use reasonable commercial endeavours (subject to planned outages detailed below) to ensure that the Enterprise iPaaS Platform is Available 99.95% of the time during the Support Hours.

The actual Availability shall be stated as a percentage and calculated as follows:

$$\left(\frac{(x-y)}{x}\right) \times 100$$

where:

x = the total number of minutes during the Measurement Period less the Scheduled Downtime.

y = the total number of minutes during the Measurement Period where the Enterprise iPaaS Platform is Unavailable (other than unavailability due to Scheduled Downtime or an Exclusion).

Any planned outages will be excluded from the above instances. Capgemini will be relieved of its service level obligations during any instance of Force Majeure.

Capgemini shall prioritize all Support Requests based on its reasonable assessment of the severity level of the Incident reported; and respond to all Support Requests in accordance with the Response Times specified in the table set out below.

Severity level of Incident	Definition	Service Level Response
1 Critical Extensive	Business Critical: Outage of a Small, Medium or Large Platform Instance leading to total loss of service of all APIs and/or integrations deployed on the Enterprise iPaaS Platform.	2 hours
2 High Significant	Business Impact: Impact on key functionality and/or performance degradation of the Enterprise iPaaS Platform but Purchasing Entity critical business functions are still operational. No acceptable workaround is available.	6 hours
3 Medium Moderate	Operational Impact: Moderate impact on usage of Enterprise iPaaS Platform but remains operational. A workaround is available to improve the situation until the issue is fully resolved.	12 hours

Severity level of Incident	Definition	Service Level Response
4 Low Minor	Minor Impact: Minor impact on usage of Enterprise iPaaS Platform. Includes minor, cosmetic, or documentation related issues. No impact on Enterprise iPaaS Platform features.	24 hours

On receipt of a Support Request by the Offeror iPaaS Support team, the validity and severity of the Incident is assessed, with a response provided within the above timescales to either accept or return the Incident and amend the Incident severity in line with the definitions if required.

The maximum number of Support Requests which the Purchasing Entity can raise in each month is as follows:

Number of Platform Instances	Purchasing Entity Incident Allowance Per Month
1-4	60
5-10	120
11-15	180

Actual resolution time will depend on the nature of the fault underlying the Incident.

The following matters are excluded from, and do not form part of, the Subscription and Support Services:

- enhancements to the Enterprise iPaaS Platform that are not part of a planned Release
- the provision of Enterprise iPaaS Software other than as specified in the Service Solution
- the resolution of Incidents due to improper Use or a failure by the Purchasing Entity to comply with a Purchasing Entity Responsibility
- support of any version of the Enterprise iPaaS Platform which has been discontinued by Capgemini
- support of any version of the Enterprise iPaaS Platform if the Purchasing Entity refuses to accept an Upgrade Release in respect of a Platform Instance within the forty (40) Business Days' window;
- any Incident caused by the Cloud Hosting Provider services on which the relevant Platform Instance is hosted (and for these purposes, the Purchasing Entity acknowledges that where any Incident relates to any services supplied by the Cloud Hosting Provider, Capgemini's ability to respond and resolve any Incidents will be limited by the service response and support offered by the Cloud Hosting Provider and the service levels as set forth in the Cloud Hosting Provider's terms of service will apply)
- each, an "**Exclusion**"

Where the root cause of an Incident, in the reasonable opinion of the Offeror, is attributable to any one or more of the exclusions set out in Section Service Solution (above), the Service Levels shall not apply to the Incident.

Capgemini may reasonably determine that any services are Out-of-scope Services. If Capgemini makes any such determination, it shall promptly notify the Purchasing Entity of that determination. For the purposes of this Document, "**Out-of scope Services**" shall mean any services provided by Capgemini in connection with any apparent problem regarding the Enterprise iPaaS Platform reasonably determined by Capgemini not to have been caused by a fault, but rather by a Purchasing Entity cause (including, without limitation, any improper use, misuse or unauthorised alteration of the

Enterprise iPaaS Platform by the Purchasing Entity) or a cause outside Capgemini's control (including any investigational work resulting in such a determination).

The Purchasing Entity acknowledges that Capgemini is not obliged to provide Out-of-scope Services.

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.

Upon execution of a Participating Addendum, Capgemini will be able to have available a Customer Service Representative (CSR) available by phone or email 7 am to 6 pm Monday through Sunday. The CSR can assist the Purchasing Entity with securing from the Cloud Provider the specific support actions required to address the need.

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

AWS and Azure (IaaS, PaaS, SaaS)

AWS and Azure provide service credits to customers if they fail to meet their published SLAs for their cloud services.

The following table provides a list of the common AWS cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for AWS Cloud Services will be provided as an attachment.

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
EC2 Instances & EBS Volumes	< 99.95%	10%	100% – EC2 or Amazon EBS was in a state of "Region Unavailable"	"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.
	< 99%	30%		
S3 (storage services)	< 99.9%	10%	100% – Error Rates	"Error Rate" means: Total No. of internal server errors / Total number of requests for application request during that five minute period
	< 99%	30%		
S3 (storage services) Standard – Infrequent Access (Standard-IA)	< 99%	10%	100% – Error Rates	
	< 98%	30%		
CloudFront (CDN service)	< 99.9%	10%	100% – Error Rates	
	< 99%	25%		
	< 99%	25%		

AWS Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Route 53 (DNS service)	5 – 30 minutes	1 day	Duration was not 100% available	
	31 mins – 4 hours	7 days		
	More than 4 hours	30 days		

The following table provides a list of the common Azure cloud services SLA's with monthly uptime, service credits and monthly uptime calculations. The complete list of SLA's for Azure Cloud Services will be provided as an attachment.

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Virtual Machines	< 99.99%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no External Connectivity.
	< 99%	25%		
	< 99%	25%		
Storage Service LRS, ZRS, GRS, and RA-GRS (write requests) Accounts	< 99.9%	10%	100% – Average Error Rate	"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. GRS – Geographically Redundant Storage LRS – Locally Redundant Storage RA-GRS – Read Access Geographically Redundant "Cool Access Tier" is an attribute of a Blob Storage Account indicating that the data in the account is infrequently accessed and has a lower availability service level than data in other access tiers.
	< 99%	25%		
Storage Service RA-GRS (read requests) Accounts	< 99.99%	10%	100% – Average Error Rate	
	< 99%	25%		
Storage Service – LRS, GRS and RA-GRS (write requests) Blob Storage Accounts (Cool Access Tier)	< 99%	10%	100% – Average Error Rate	
	< 98%	25%		
Storage Service – RA-GRS (read requests) Blob Storage Accounts (Cool Access Tier)	< 99.9%	10%	100% – Average Error Rate	
	< 98%	25%		
	< 99%	25%		
	< 99%	25%		

Azure Services	Monthly Uptime percentage	Service Credit	Monthly Uptime calculation	Comments
Traffic Manager (DNS service)	< 99.95%	10%	((Maximum Available Minutes – Downtime) / Maximum Available Minutes) x 100	Downtime: The total accumulated Deployment Minutes, across all Profiles deployed by you in a given Microsoft Azure subscription, during which the Profile is unavailable.
	< 99%	25%		
CDN Service	< 99.9%	10%	Downtime: To assess Downtime, Microsoft will review data from any commercially reasonable independent measurement system used by the customer.	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).
	< 99.5%	25%		
	Encrypted: >6 hours Monthly RTO	100%		
	< 99%	25%		

Virtustream (IaaS)

Virtustream shall provide the Purchasing Entity with the Service Level Credit if the Services fail to satisfy any of the Service Levels set out herein. The Service Level shall commence thirty (30) days following commencement of the Steady State phase of the SOW. Each of the Service Level Credits shall be based on the fees paid for the applicable service, as set out in this SOW. The aggregate Service Level Credits for all Service Levels in any month shall not exceed 15% of the total monthly recurring charges ("MRCs"). The Service Level Credit shall be Customer's sole and exclusive remedy and Virtustream's sole and exclusive liability for Unscheduled Downtime.

µVM Availability on Enterprise Core µVM	µVM Availability on Enterprise Basic µVM	µVM Availability on DMZ Core VM	µVM Availability on DMZ Basic VM	Service Level Credit [Note 2]
99.95% – 99.999%	99% - 99.5%	99.5% - <99.9%	99% - <99.4%	1%
99.5% - 99.94%	98% - 98.99%	98% - <99.5%	98% - <99%	3%
95% - 99.4%	95% - 97.99%	95% - <98%	95% - <98%	5%
90% - 94.99%	90% - 94.99%	90% - <95%	90% - <95%	10%
Below 90%	Below 90%	Below 90%	Below 90%	15%

BMC and ServiceNow (SaaS)

BMC and ServiceNow make every reasonable effort to meet their target response times and to resolve issues as soon as possible. However, neither provider offers service credits should those reasonable efforts fail to meet published targets.

Capgemini iPaaS (PaaS)

The Purchasing Entity must inform the Offeror of Unavailability of the Enterprise iPaaS Platform within 10 days of the end of the month in which the Purchasing Entity determines the Enterprise iPaaS Platform was Unavailable.

Subject to the Service Credit Cap (in excess of which, no Service Credits shall be payable) and the notice period as defined in the above paragraph, if during any Measurement Period the actual Availability is lower than 99.95% then the Offeror shall:

- credit the Purchasing Entity's account; or
- if the Subscription Charges have been paid by the Purchasing Entity in advance, reimburse the Purchasing Entity,

- in each case, by an amount calculated as follows:

$$\frac{(99.95 - a)}{100} \times b$$

where:

a = the actual Availability achieved (expressed as a percentage);

b = means: (i) the total Subscription Charges payable during the relevant Measurement Period; or (ii) where the Subscription Charges have been paid by the Purchasing Entity in advance of the services, the proportion of those Subscription Charges paid in advance which would have been payable (on a pro-rata basis) had the Subscription Charges been payable monthly and arrears.

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

Cloud Service Providers have scheduled weekend maintenance windows and strive for zero downtime maintenance. They have deployed functionality that allows restarts of services without impacting the active instances. If planned downtime is required they provide notifications in advance. Notification of planned downtime will be communicated as far in advance as possible, with a week in advance of such changes being the targeted minimum. Notifications contain the actions that will be taken, any impact on services, and the scheduled date and times.

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

Cloud Service Providers will plan with the customer-specific remediation actions during the transition and stabilization periods required for cloud migration projects. The DR strategy to be implemented will define the failover and failback plans to restore service in an event of an unplanned outage.

If a disaster occurred resulting in transfer of operations to an alternate data center and either or both the Recovery Time Objective and/or the Recovery Point Objectives were not met; the Vendor will issue negotiated service credits based on several factors including the amount of extra work required while system was down, missed delivery of benefits to constituents, and damage to State and Cardholder reputation.

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.

Cloud Service Providers can provide performance reports, these originate from their respective web portals, see samples below. The reports can be generated online and can be configured to report performance metrics in real time. Performance data can be downloaded for further analysis by customers.

The Cloud Service Providers make available performance reports of their cloud services, which allows the clients to monitor key performance analytics through dashboard provided by the cloud service providers, for example, IaaS performance metrics can be readily available like CPU, memory, storage, network, etc. The dashboards provided by the cloud service providers can be adapted for either monitoring or problem resolution purposes. The section below provides performance sample reports for AWS, Azure, and Virtustream.

AWS Sample Performance Report

The figure below provides an AWS sample report.

Figure 19: AWS sample report

Azure Sample Performance Report

The figure below provides an Azure sample report.

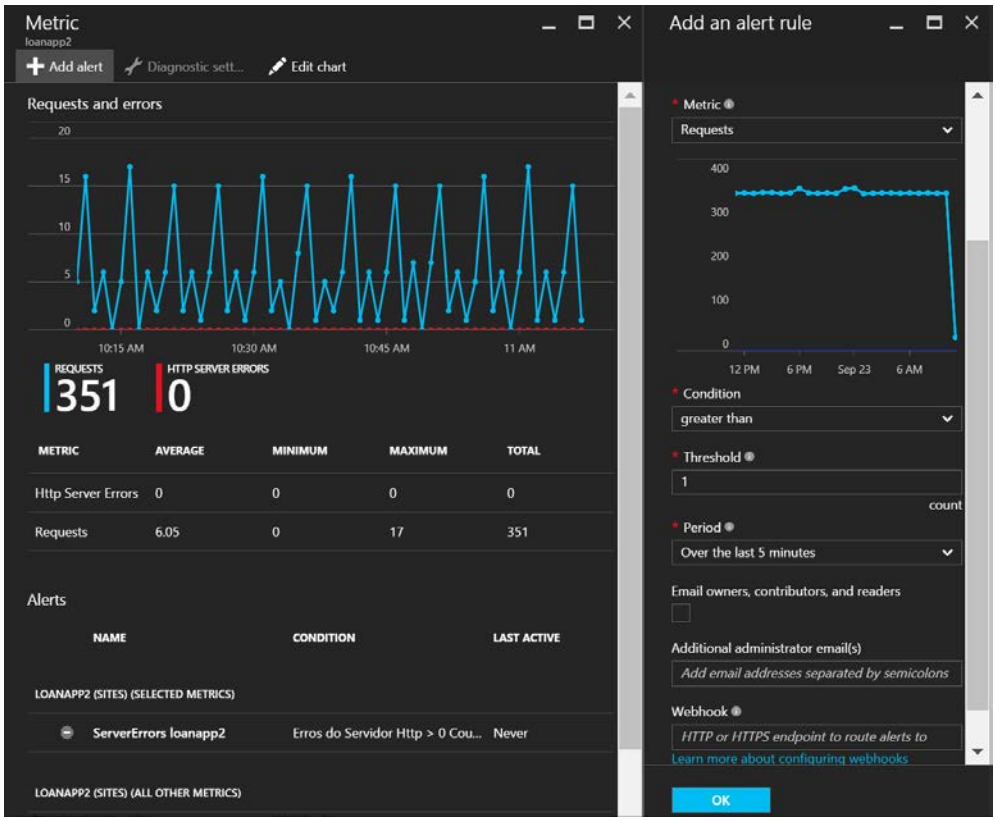

Figure 20: Azure sample report

Virtustream Sample Performance Report

The figure below provides a Virtustream sample report.

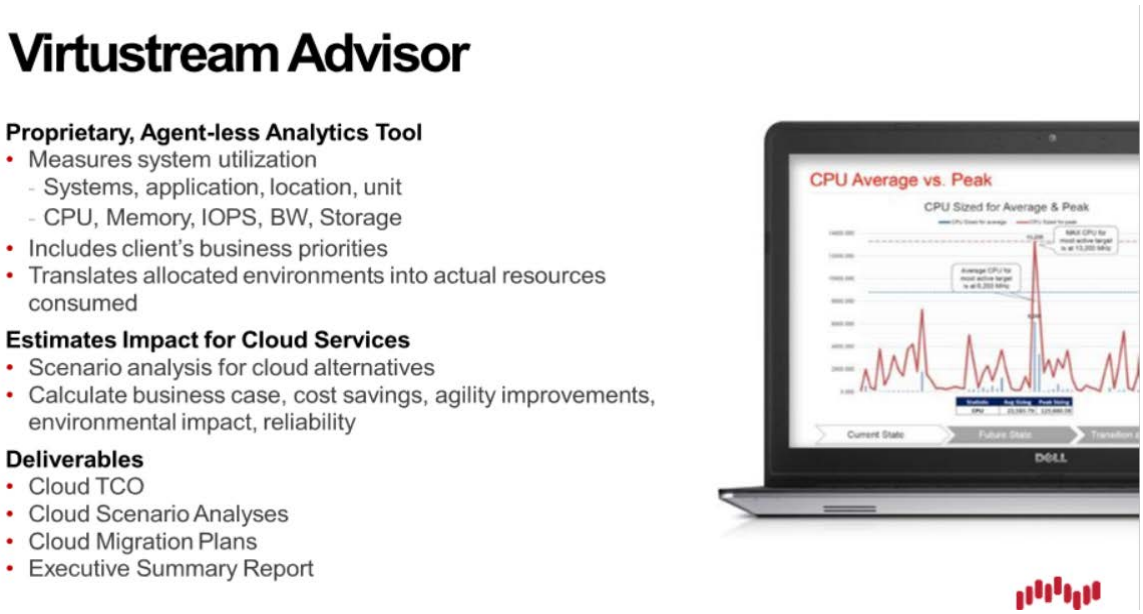

Figure 21: Virtustream sample report

BMC Sample Performance Report

Remedy OnDemand publishes a service status dashboard for its customers through the i.onBMC.com support portal. This dashboard provides a real-time view of the following information and key performance indicators (KPIs):

- Production platform version
- Number of open support requests and a link to the support request report
- "Active" metrics including:
  - Number of Fixed, Floating and Read licenses in use
  - Number of emails in the queue
- "Last hour" metrics including:
  - Number of users logged in
  - Last login response time
  - Number of incidents, service requests, change requests and work orders submitted
- "Last 30-day" metrics including:
  - System availability
  - Average login response times
  - Unique user logins
  - Fixed and Float AR license usage
  - Incident, Change and Work Order activity
  - Active email volumes

Metrics are presented in easy-to-read colorful graphics. Here's a sample view of the dashboard:

Figure 22: sample view of the dashboard

ServiceNow Sample Performance Report

States and Cardholders can use the support portal to obtain information about the real availability of all their instances. Real availability is the percentage of production time that an instance is up and available for use.

In viewing the below example one can see production and non-production instances are shown. If the Organization has multiple production instances, then the production real availability percentage is shown for each production instance.

For more information about an individual outage, point to a yellow or red "chiclet." The incident number, start time, end time, and impact are listed. The impact is the outage duration calculated as outage start time minus outage end time.

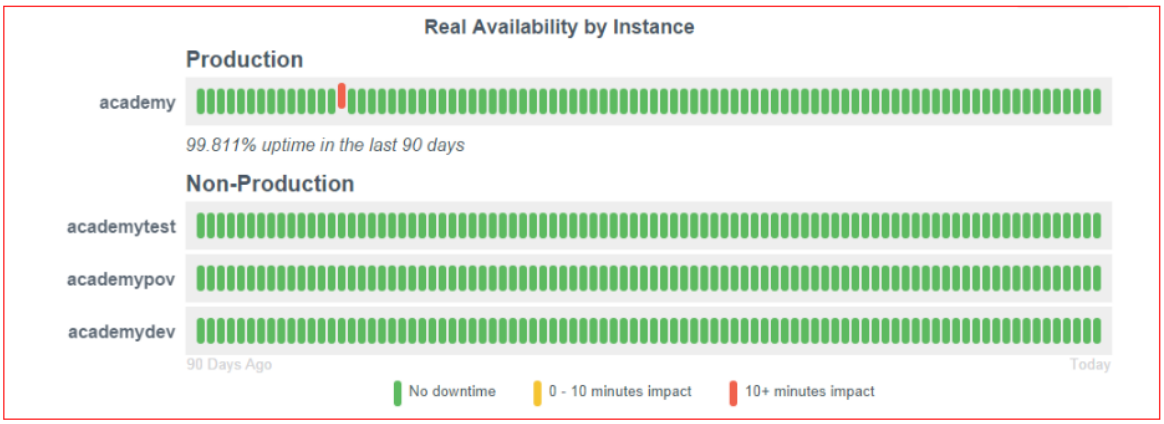

Figure 23: Real Availability by Instance

To drill down on which days may have been impacted by an outage when you're looking at the chicklets, click on "more details" for an instance; you'll be taken to a new page where you can see up to 12 months of available information in an easy-to-read and mobile friendly calendar.

By default, time ranges have been set for 3, 6, and 9 months (90, 120, and 180 days) to make those easy to select, as those have been some of our most frequent requests.

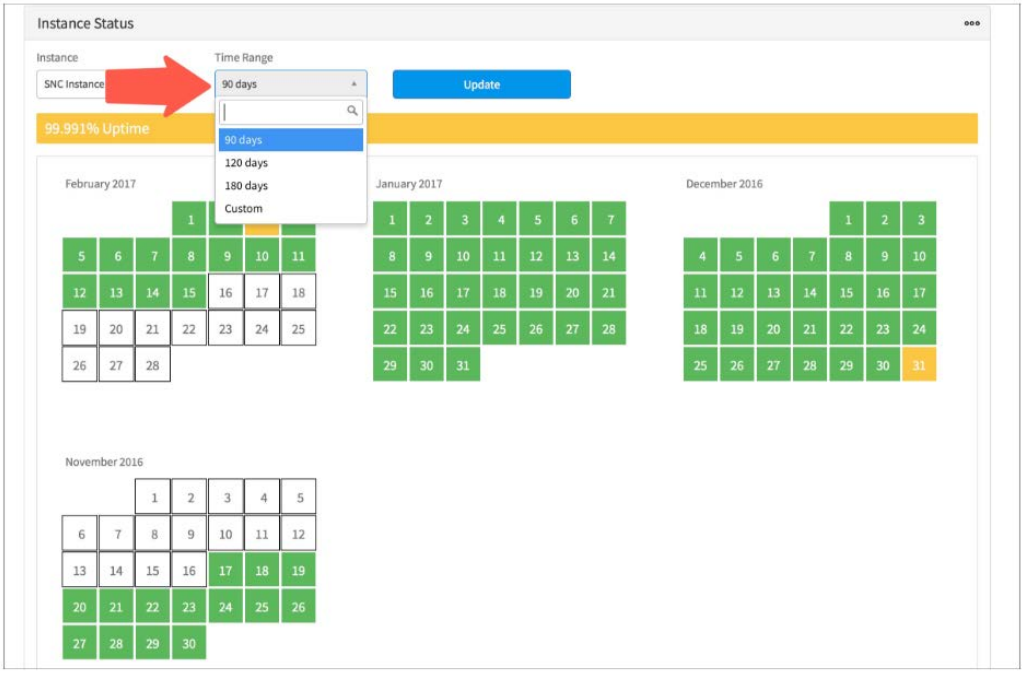

Figure 24: Instance Status

Capgemini iPaaS

The Capgemini Enterprise iPaaS provides tooling for monitoring and alerting of both the platform and services hosted on the platform. There are a number of standard real-time dashboards and retrospective reports available, additionally the platform provides the Purchasing Entity the ability to amend the standard dashboards and reports or to create bespoke versions. Below are some dashboard examples illustrating the metrics monitored for our sample application.

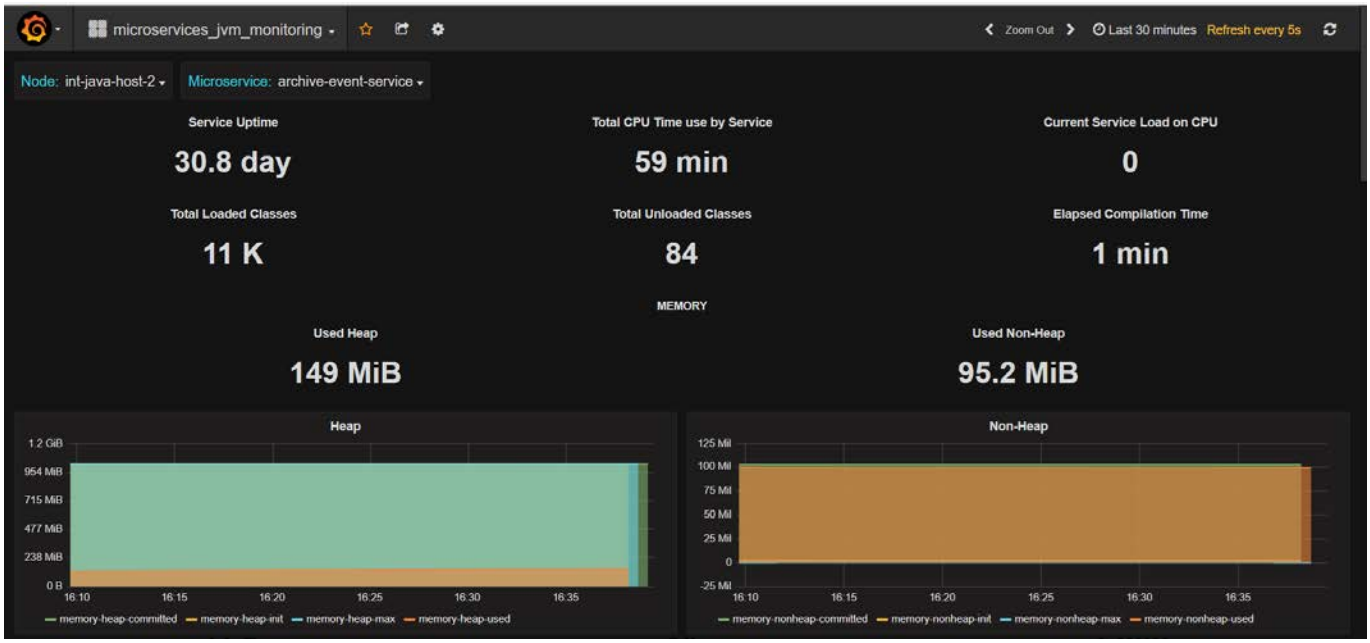

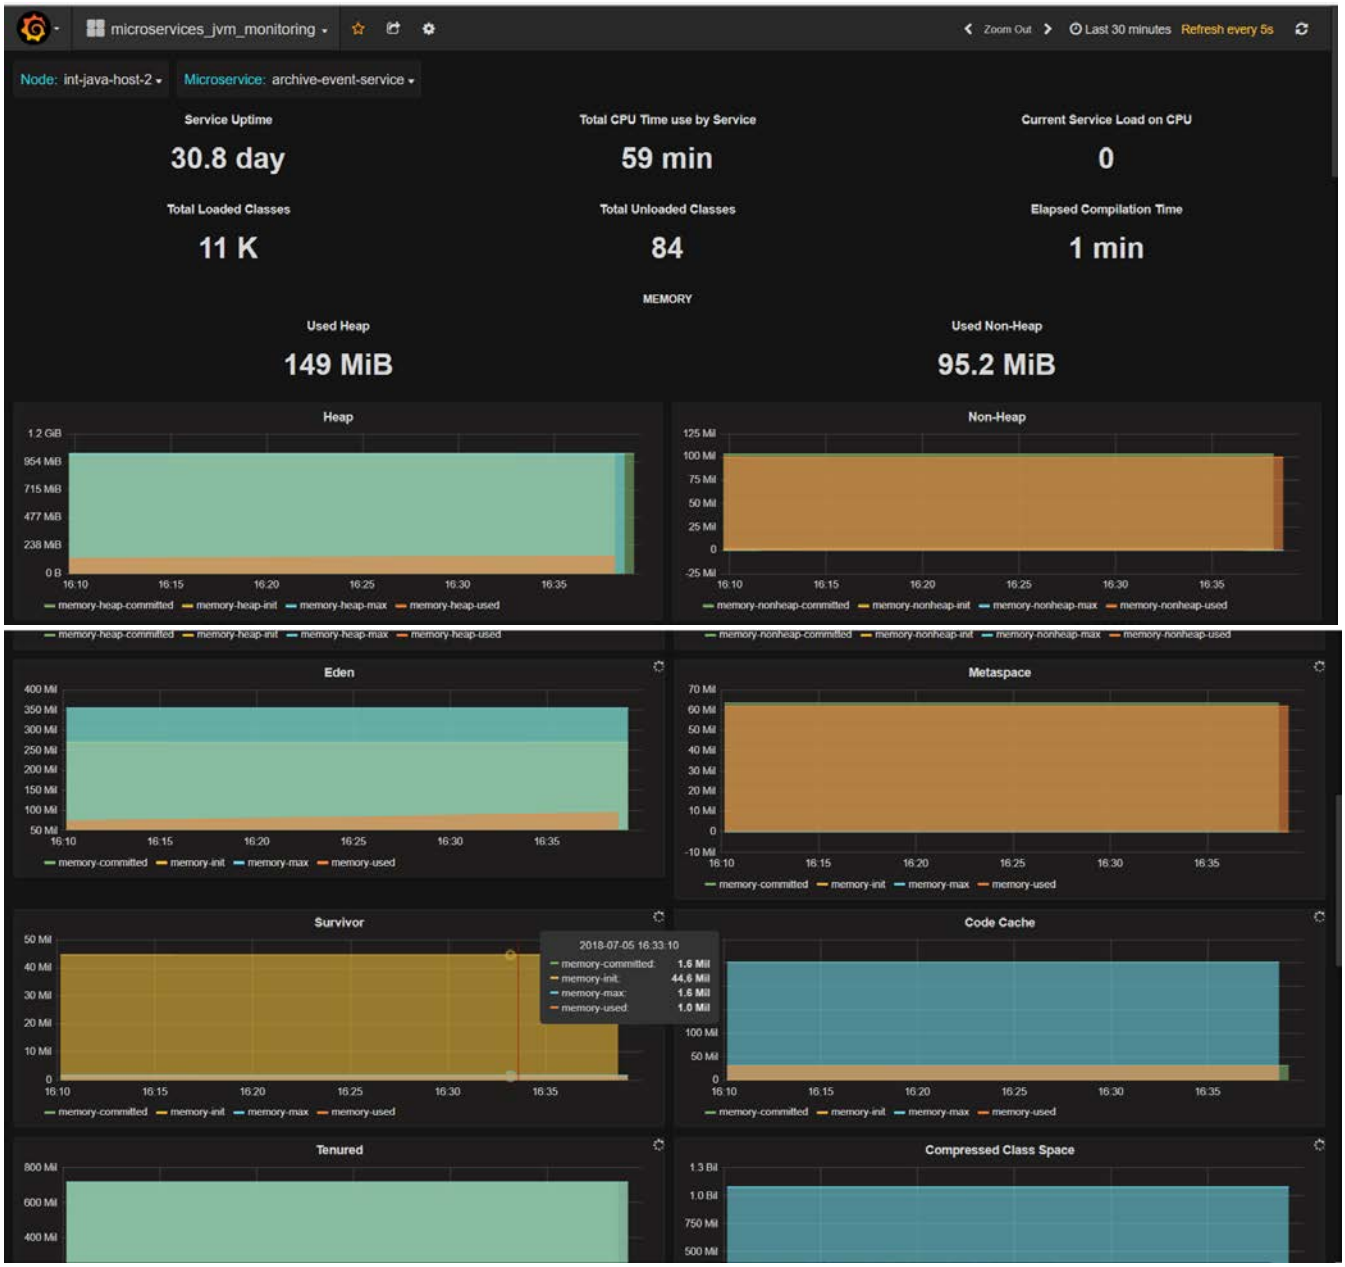

8.12.8 Ability to print historical, statistical and usage reports locally.

Cloud Service Providers can assist customers to establish the processes to gather usage reports and metrics.

Our Cloud Service Providers provide services to generate usage reports, statistics, metrics, and other data for analysis. This data is accessible through their respective web portals through a variety of tools, and by direct download of data. Access is through the Purchasing Entities' web browsers placing the reports, tools, and data on Users' local PCs for further analysis and printing.

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

The services that are provided by Capgemini's Cloud Service Providers include an on-demand deployment service as described in Section 8.1.3. That service along with their other provided services are supported on a 24x7x365 basis in that the Purchasing Entity may at any time log an issue reporting a degradation of service quality. The Cloud Service Provider of the affected service will assess the criticality of the incident against their published criteria and assign the appropriate severity. The

assigned severity will dictate when and how quickly the Provider will work to restore service quality. Further service support details are provided

Cloud Service Providers-IaaS

All Capgemini IaaS Cloud Service Providers offer 24x7x365 support, named individuals, escalation lists, contact details will be documented in the Statement of Work at the time the service details are determined.

Cloud Service Providers-SaaS

BMC and ServiceNow both support their services 24x7x365 with access to the Providers' 24x7 Service Desks. Details of how to access Provider support are contained in the Service Statements issued by the providers at the time services are acquired. The following paragraph is Indicative wording within the Customer Support section of the Service Statement.

"Business Hours"

Customer Support is available 24 hours a day, 7 days a week, including all holidays.

Access Contacts

Customer may make contact using one of the following means:

- Support Portal at https://servicedesk.serviceprovider.com/. Customer may get login access to this self-service portal by contacting its Service Provider administrator.
- Phone using one of the numbers at http://serviceprovider.com/support/contact-support.html.

Cloud Service Providers-PaaS

Capgemini's iPaaS service is supported by 24x7x365. Support requests are communicated via email, logged and worked with the assigned priority.

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

Cloud Service Providers-IaaS

Capgemini can automate the configuration for the scale up and scale down of IaaS resources on the Cloud Service Providers platforms, these services are available on demand 24x7x365.

Cloud Service Providers-SaaS

BMC and ServiceNow provide their software as a service, hosted on the private cloud. The benefit of subscribing to either of these services is that the Purchasing Entities do not have to think about, plan for, or deal with the underlying infrastructure. As Purchasing Entities place greater demands on the service, BMC and ServiceNow will provide that service performance remains will continue to without any action from the Purchasing Entities. The expansion and contraction of the underlying infrastructure are one of the many benefits of subscribing to SaaS.

Cloud Service Providers-PaaS

The Capgemini Enterprise iPaaS includes Operational Management tooling to allow the purchaser to deploy integration and microservice payloads onto the platform and to provision additional compute capacity. Requesting additional compute capacity can be performed 24x7x365.

8.13 (E) CLOUD SECURITY ALLIANCE

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

a. Completion of a CSA STAR Self-Assessment. (3 points)
b. Completion of Exhibits 1 and 2 to Attachment B. (3 points)
c. Completion of a CSA STAR Attestation, Certification, or Assessment. (4 points)
d. Completion CSA STAR Continuous Monitoring. (5 points)

Capgemini has completed the CSA STAR Registry, as per the confirmation image below:

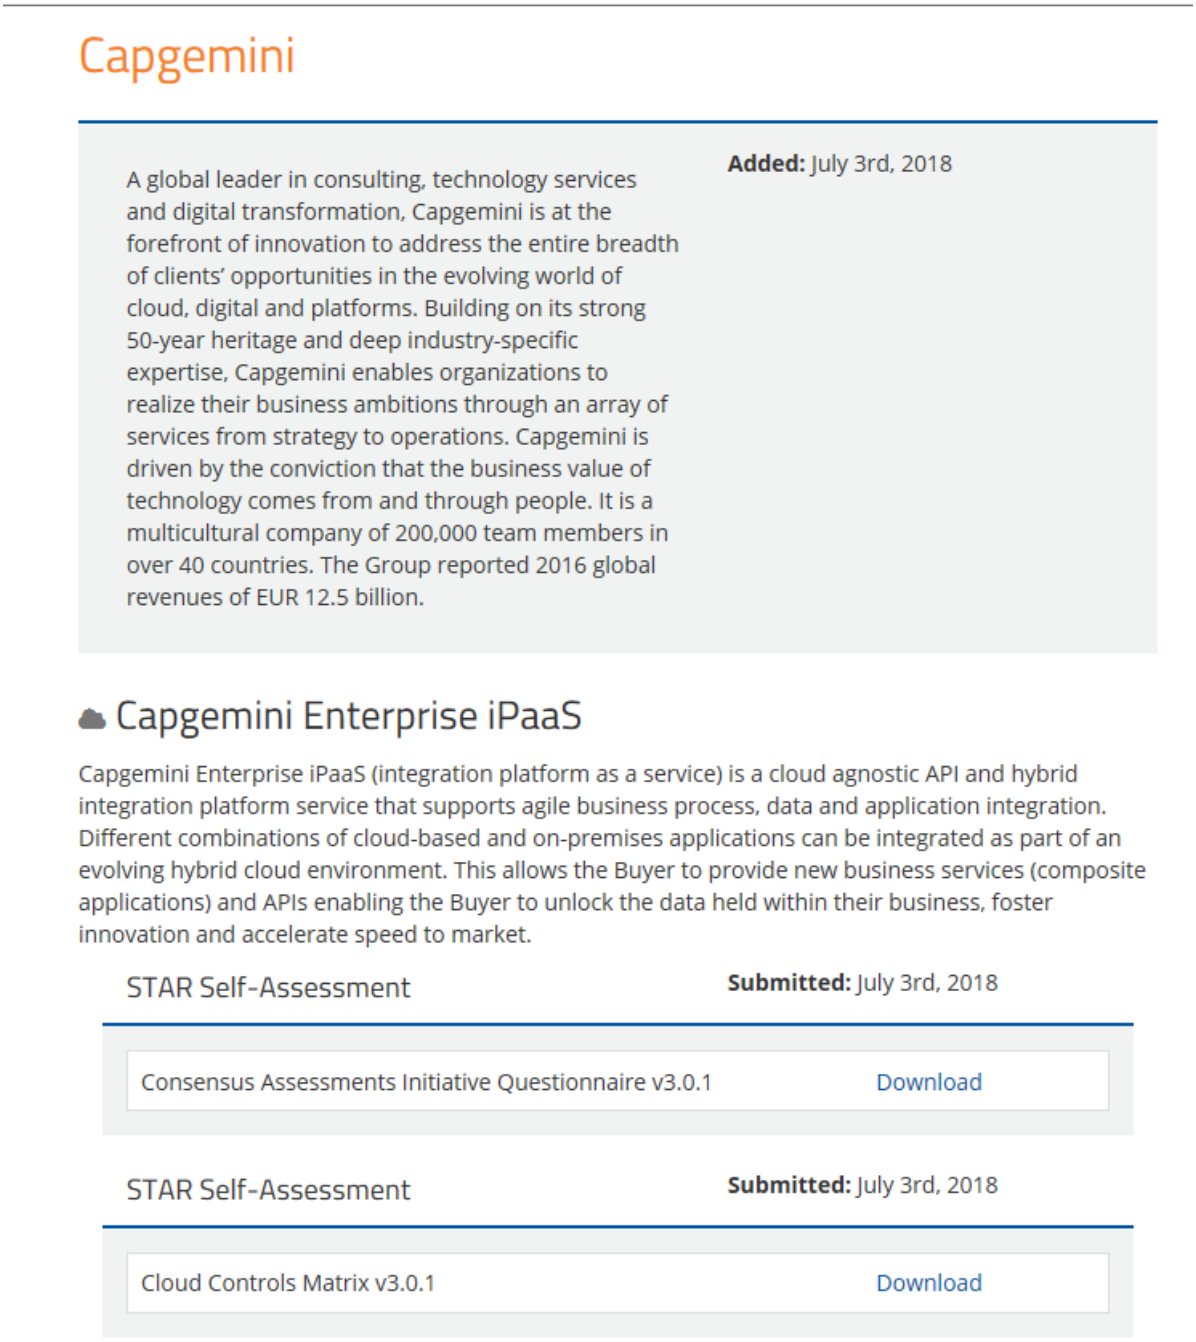

Figure 25: CSA STAR Registry

Capgemini has provided the following CAIQ and CCM forms as separate attachments for Amazon, Microsoft, ServiceNow, BMC, Capgemini and Virtustream IaaS, PaaS, and SaaS Offerings.

Please refer to separate attachments by the following names.

- AWS – CSA Assessment Report
- AWS – Exhibit 1 to Attachment B CAIQ
- Azure – CSA Assessment Report
- Azure – Exhibit 1 to Attachment B CAIQ
- Azure – Exhibit 2 to Attachment B CCM
- BMC – Exhibit 1 to Attachment B CAIQ
- Capgemini iPaaS – Exhibit 1 to Attachment B CAIQ
- Capgemini iPaaS – Exhibit 2 to Attachment B CCM
- ServiceNow – Exhibit 1 to Attachment B CAIQ
- Virtustream – Exhibit 1 to Attachment B CAIQ
- Virtustream – Exhibit 2 to Attachment B CCM

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.

Capgemini can through an agreed upon statement of work define and stand up the Standard Operating Procedures (SOPs) necessary to manage and escalate emergency requests, rush services and non-emergency incidents and requests for AWS and Azure, and work to schedule and address those incidents providing for a higher level of resolution by applying the following processes:

- Continuous monitoring of IT infrastructure
- Registering of events and incidents
- Follow-up and monitoring of events notified through the cloud service providers monitoring services
- Analysis of events to detect the trouble origin
- Escalate and resolve critical incidents
- Report generation
- Creation of monitored resources log data

Capgemini can include a fully integrated Change Management process that is aligned with customers transition milestones and risk factors based on severity and impact of the changes to be implemented, these are categorized in the diagram below. This approach is designed to handle standard changes and complex changes applying the governance standards for all changes.

In any organization, there are two ways in which an Incident can originate – via an End User or an alert generated by the monitoring tools. In this manner, Capgemini will help a Purchasing Entity manage all the Cloud Service Provider incidents raised by end users or monitoring systems, as depicted in the table below.

RISK FACTORS BASED ON THE MAGNITUDE OF CHANGE		
Low	**MEDIUM**	**HIGH**
IT Operational Management Impact IT Staff Impact Business User Impact	Business and IT Commitment to Change	**Type of Transition** **Degree of Process Change** **Geographic Impact**

RISK FACTORS BASED ON READINESS FOR CHANGE		
Low	**MEDIUM**	**HIGH**
Readiness of Leadership Impact of Pat Change Efforts Understanding of OCM Aligned to the Business Case	Business and IT Commitment to Change	**Impact of Competing Initiatives**

Figure 26: Cloud Service Provider Incidents

The expedient timeframe will be determined at beginning of the project with the Purchasing Entity to determine ample change windows for standard and complex changes. It is crucial to understand the landscape of stakeholder interests and organizational capabilities in relation to the anticipated change, thereby reducing risk throughout the process.

Capgemini accounts for Change management on every project. Change Management is a process that needs to be anticipated in the early stages of any cloud migration or cloud transformation engagement. In the case of an emergency or rush service implementation, change management processes will require an evaluation of the impact the suggested changes will have on the organization.

Capgemini's Change Management processes focus on assisting the customer's IT organization in adopting a new way of working, along with gaining acceptance when changes occur across the organization. Change Management contributes to success by maximizing the traction and trajectory of the customer's cloud transformation projects and associated engagements.

Ultimately the success of any project rests on the speed and degree to which it is adopted. Capgemini recognizes each organization has its own unique character and challenges, and it's this character that's assessed in the customer's environment. Effective organizational change planning takes into consideration all aspects of change. Below you will see the key components of our Change Management Strategy designed specifically for each customer's needs. The Change Management Team will orchestrate multiple and integrated workstreams to enable Organizational and Individual change adoption through Awareness, Understanding, Ability, and Ownership.

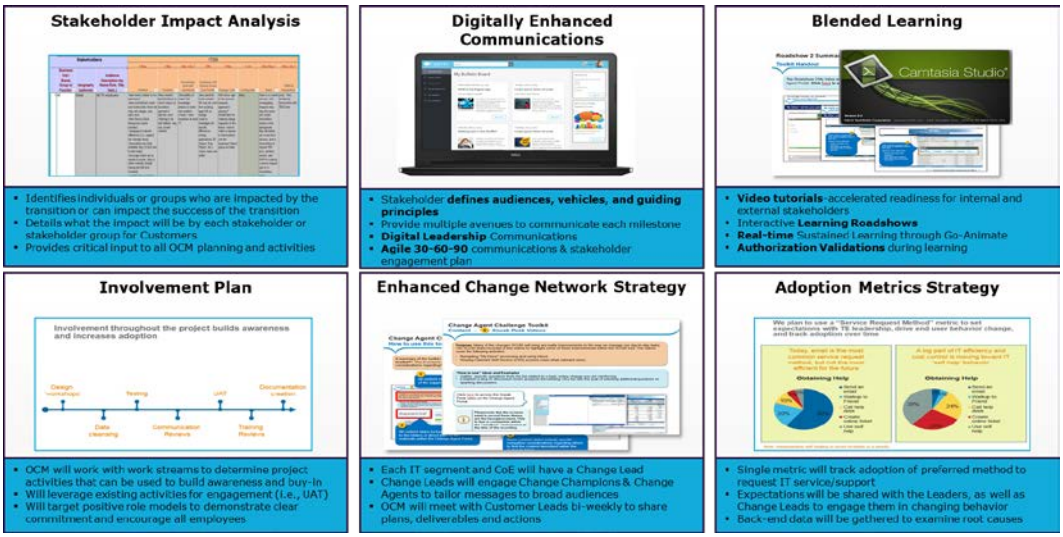

Figure 27: Change Management

Depending on the level of change, a fit for purpose Change Management plan will be designed and executed leveraging our leading practices and processes to takes a customer-centric approach that will reduce risk, provide proper controlled releases to the cloud production environment.

We can meet most Purchasing Entity requests to get to deployment, even “emergency” or “rush” requests. There are a few aspects to every request that needs to be understood by all parties before even the most urgent requests can be successfully undertaken.

11. The full set of requirements must be communicated-what does success look like?
12. Full knowledge of the starting point -what do we have to work with going into the project?
    i. The knowledge/input necessary
    ii. The approvals necessary
    iii. The funding necessary
13. Do we have enough time? Yes, this is a rush job, but even so, if the deadline cannot be met, should we start down the intended path or take a different route to the end goal?

If those points are all understood and agreed, Capgemini would staff the implementation as quickly as possible and then execute.

When the request is for a service that is not yet in production, the Capgemini Change Management process will not apply until the Deployment Phase of the project.

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

Provisioning cloud services for Purchasing Entities with whom we have an executed a Statement of Work does not require any lead time as the cloud services can be self-provisioned through the AWS or Azure management portals.

For a Purchasing Entity for whom we are provisioning the initial solution, the typical lead time from the signing of a Statement of Work to provisioning infrastructure is 5 days. It is then an additional 5 days to provide the cloud services credentials, so the Purchasing Entities can access the AWS, Azure or Virtustream cloud environments.

After those initial onboarding steps are complete, further provisioning of cloud services will not require any lead time as the cloud services can be self-provisioned through the AWS or Azure management portals. Virtustream provisioning given its private in nature is done through service requests that must be approved outside any self-provisioned portals.

The BMC Remedy on Demand and ServiceNow offers Capgemini are making are reselling of subscription-based access to cloud-hosted software as a service. The effort required on Capgemini’s part to complete these transactions is much less than the effort on the part of the Purchasing Entities who have to gather and educate stakeholders, align and document requirements, assess the available choices against those requirements, make a decision, and then secure funding.

Typically, Capgemini is asked to assist at the time a Purchasing Entity is working to articulate requirements. We respond by assigning an Architect to help with the analysis of documentation of requirements. Capgemini has several experienced Architects on staff and having one assigned within 5 days is our normal turnaround time.

The advantage for the Purchasing Entity to involve Capgemini during requirements gathering phase of their project is to start the required interaction with the SaaS providers at the same time requirements are being finalized. Following this timing effectively results in the first instance being provisioned within 24 hours of the commercial arrangements being finalized.

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.

Capgemini Cloud Service Providers will assist the Purchasing Entities with the disposal of the data at the end of the retention period; data will be deleted from the Cloud Service Providers storage pools. Backups of data can be transmitted to the Purchasing Entities if necessary.

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

AWS, Azure and Virtustream (IaaS, PaaS, SaaS)

Capgemini Cloud Service Providers will plan the DR recovery plan and runbook to meet the requirement to recover and restore data according to customers RPO's and RTO's requirements. The runbook will be configured and implemented in the Cloud Service Provider environment. The DR strategy can be defined by configuring data replication recovery time objectives (RTO) and recovery point objectives (RPO) within organizational limits.

For example, Azure Site Recovery helps provide business continuity by keeping business apps and workloads running during outages. Site Recovery replicates workloads running on physical and virtual machines (VMs) from a primary site to a secondary location. When an outage occurs at your primary site, you fail over to a secondary location, and access apps from there. After the primary location is running again, you can fail back to it. Site Recovery provides continuous replication for Azure VMs and VMware VMs, with replication frequency as low as 30 seconds. Similar DR strategy and replication can be configured in AWS environment.

BMC Remedy on Demand and ServiceNow (SaaS)

There are no known risks in the ability of BMC or ServiceNow to identify, initiate, and perform a disaster recovery event in response to any situation preventing the ability of the Provider to meet the Purchasing Entity's arranged Recovery Time Objective.

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.

Capgemini Cloud Service Providers have extensive experience architecting and implementing various levels of Business Continuity Management (BCM) designing and implementing complete Business Continuity and/or Disaster Recovery solutions, providing continuity of service and/or service recovery. Capgemini leverages leading practices to configure a Disaster Recovery plan and infrastructure in the event a failure occurs. The Disaster Recovery design is implemented in AWS or Azure cloud environment, leveraging their availability sites and data centers located throughout multiple regions in the United States.

The following table depicts the data center regions to date.

Data Center Regions by Provider		
AWS	**Azure**	**Virtustream**
US East (N. Virginia) US East (Ohio) US West (N. California) US West (Oregon)	Central US East US 2 East US North Central US South Central US	US East US West US Central

Data Center Regions by Provider		
AWS	Azure	Virtustream
	West US 2 West Central US West US	

Capgemini Cloud Service Providers through the engagement will configure the environment necessary to establish a successful failover in the event of an outage. Capgemini Cloud Service Providers will customize the RTO's and RPO'S to conform to the Purchasing Entities failover timelines and policies. Capgemini Cloud Service Providers will also monitor and manage the on-premise and failover sites for proper replication of infrastructure services, applications, and databases.

The solution diagrams below depict an enterprise architecture for on-premise infrastructure, servers, a line of business web applications and databases, each design to failover to the Cloud Service Provider infrastructure site in an automated, orchestrated manner.

Virtustream Disaster Recovery as a Service

The Virtustream private cloud offering leverages its US-based delivery centers within a Disaster Recovery as a Service (DRaaS) model to provide multiple levels of DR capabilities while allowing clients to scale the service as the criticality of their information demands or as the size of their infrastructure expands.

The DRaaS service eliminates the need for predefined CapEx expenditures and supports service scalability as a client's demands expand and evolve over time. Purchasing Entities select the tier of service which best aligns with their current demands and risk profiles. As these characteristics change over time, the DRaaS model will expand or contract to address the evolution of the client environment. This flexible model allows clients to pay as you grow and to leverage resources only as required by the current consumption levels.

The diagram below defines the Virtustream Tiered DR model.

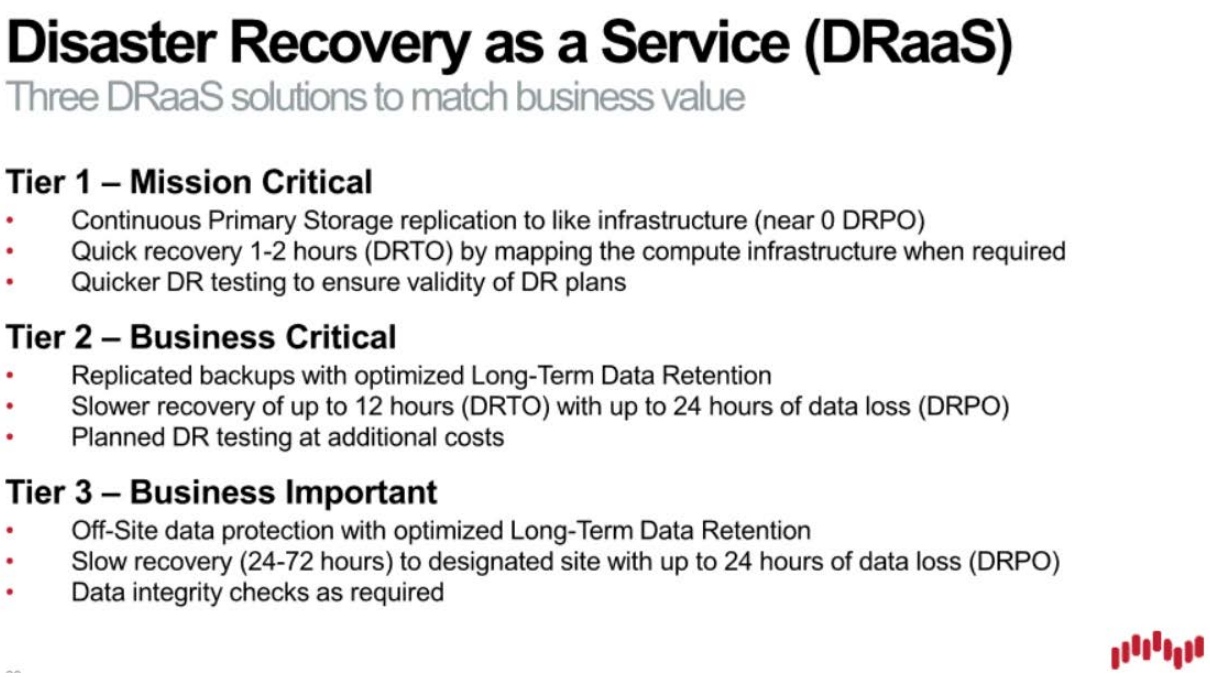

Figure 28: Virtustream Tiered DR model

The diagram below outlines the Virtustream DR benefits.

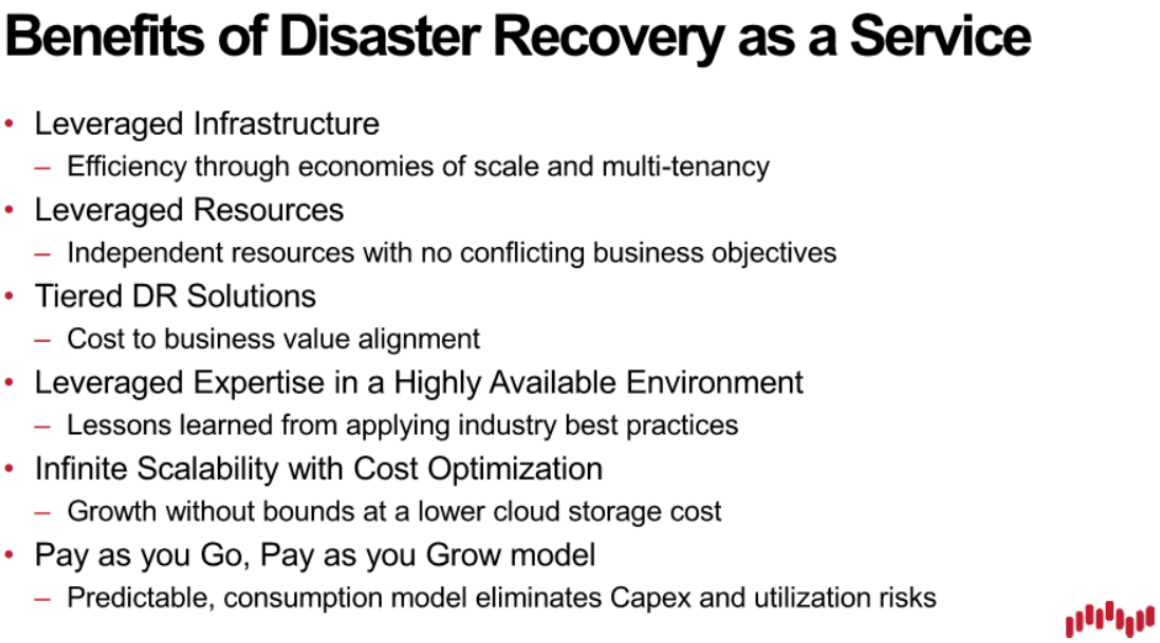

Figure 29: Virtustream DR benefits

Azure Disaster Recovery

The following diagram depicts a typical DR architecture designed to incorporate services available from Azure.

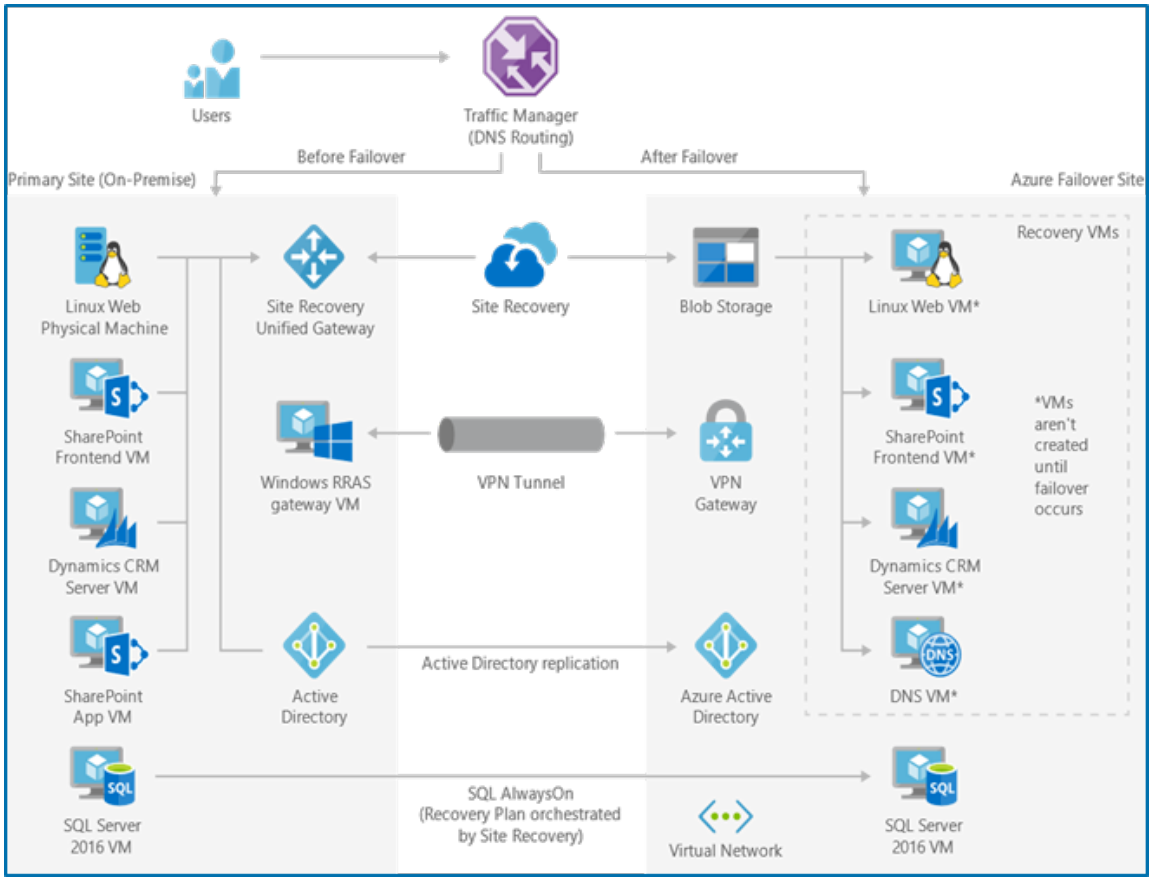

Figure 30: Typical DR Architecture

AWS Disaster Recovery

The following diagram depicts a typical DR architecture designed to incorporate services available from AWS.

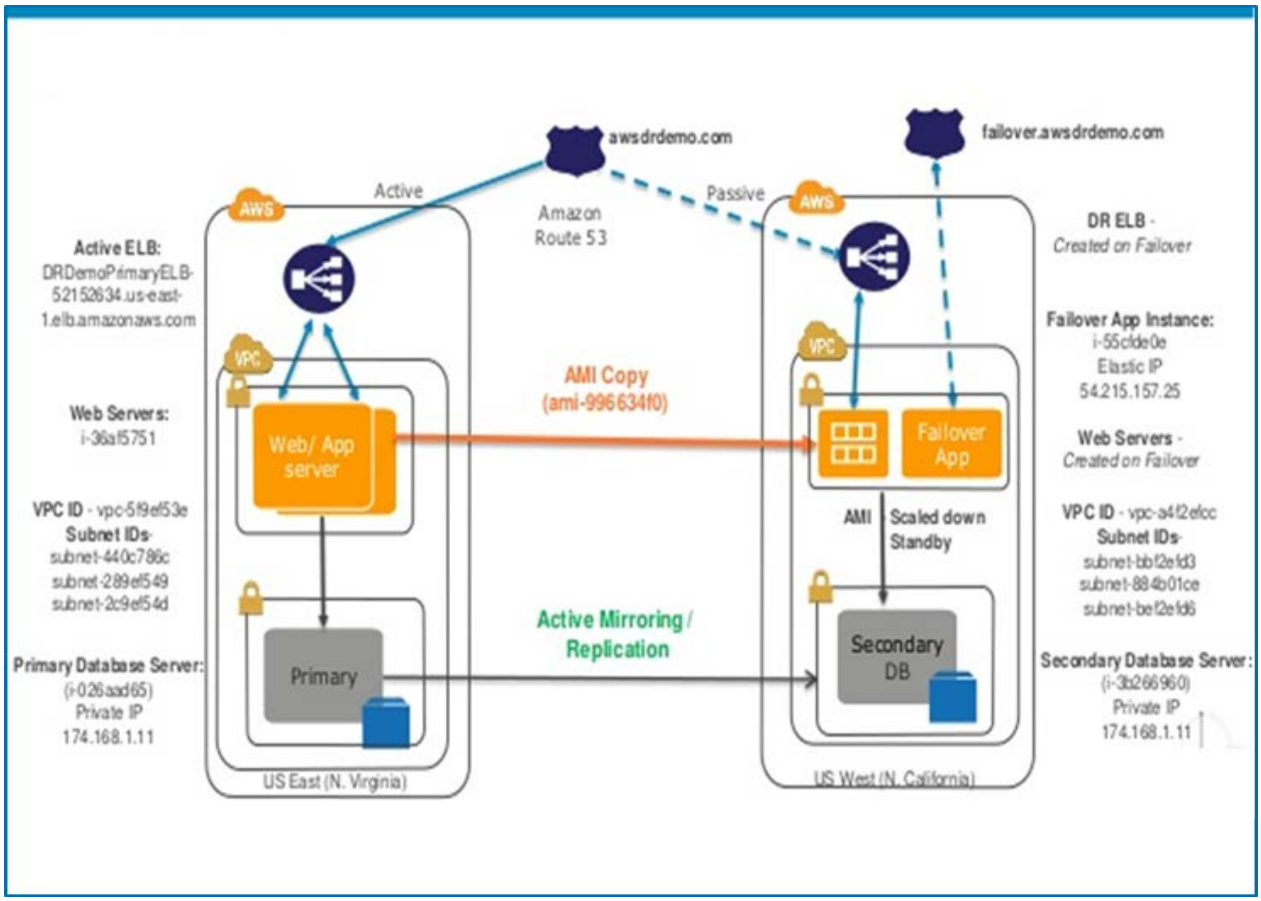

Figure 31: Typical DR Architecture

BMC Disaster Recovery

BMC has developed a SaaS Information Technology Contingency plan to manage operations in the unlikely event of a situation requiring such measures. The following details of the contingency plan explain how contingency operations are managed,

Purpose

Information systems are vital to the mission of BMC and its business functions. It is therefore critical that the services provided by BMC are able to operate effectively without excessive interruption. BMC's Information Technology Contingency plan (ITCP) establishes comprehensive procedures to recover BMC Remedy OnDemand services quickly and effectively following a service disruption.

BMC's ITCP establishes procedures and mechanisms that obviate the need to resort to performing operational functions using manual methods. If manual methods are the only alternative, every effort will be made to continue operating and support functions and processes manually. In order to maintain a normal level of efficiency, it is important to decrease real-time process engineering by documenting notification and activation guidelines and procedures, recovery guidelines and procedures, and reconstitution guidelines and procedures prior to the occurrence of a service disruption. Upon a disaster declaration, appropriate personnel is apprised of current conditions and damage assessment begins. As part of the recovery phase, appropriate personnel takes a course of action to recover the BMC OnDemand components at a site other than the one that experienced the disruption. Moving into

the reconstitution phase, actions are taken to restore system processing capabilities to normal operations.

BMC contracts with Equinix for data center infrastructure and services. BMC's Remedy On Demand services is operated out of their contracted two processing sites in Chicago, IL and in Santa Clara, CA with each city containing a data center pairing. The Chicago Data Center pair and the Santa Clara pairs provide the active/passive arrangement BMC uses in the extremely unlikely event that one of the data centers in either pair experiences an outage that necessitates transferring operations to the paired data center.

Data is also replicated between Chicago and Santa Clara to guard against an outage event that impacts the greater metropolitan area in either city. Following is a representative map of BMC's global data center locations.

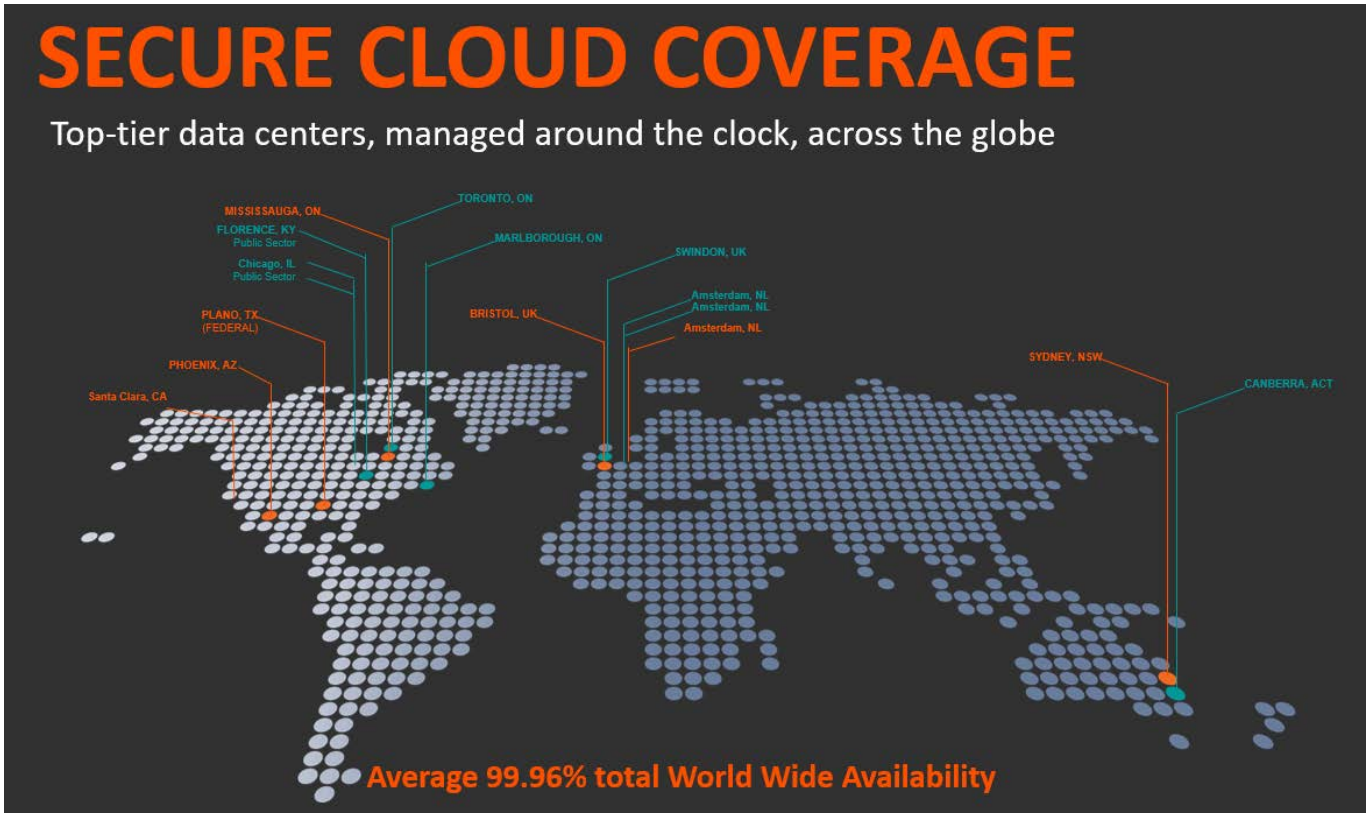

Figure 32: MC's global data center locations

Compliance

BMC's ITCP supports the requirements for the Federal Risk and Authorization Management Program (FedRAMP). The ITCP denotes interim measures to recover services following an unprecedented emergency or system disruption. Interim measures include the relocation of production systems and services to an alternate site. Unless otherwise agreed in advance, alternate sites will always reside within the same country as the primary site.

Scope

In accordance with Federal Information Processing Standards (FIPS) 199, BMC follows guidelines on determining the potential impact to organizational operations and assets, and individuals through a formula that examines three security objectives: confidentiality, integrity, and availability. The procedures in the ITCP have been developed for a moderate-impact system and are designed to recover BMC OnDemand services within arranged RTO targets.

ServiceNow Disaster Recovery

ServiceNow's production cloud environment is architected to host customer instances in region-specific, geographically dispersed data center pairs that operate in an active-active mode. In the US those data center pairs are located in Santa Clara, CA, and Manassas, VA. Instance data is replicated in near real-time between the data center pairs. In the event of an operational fault, failure, outage or attack, customer traffic can be quickly rerouted using our Advanced High Availability (AHA) capability so that you maintain access to your instances and data.

Figure 33: ServiceNow Disaster Recovery Centers

The AHA process is comprised of eight main steps and it is invoked through using an instance of Now Platform (using an implementation of ServiceNow to orchestrate the Nonstop Cloud) when one of two conditions is met:

14. In the event of a service disruption, the ServiceNow operations team determines whether a failover is required.
15. For scheduled maintenance activity, the ServiceNow operations team determines if an AHA transfer should be performed.

While Advanced High Availability is the primary means to recover data and restore service in the case of a disruption, in certain cases it is desirable to use ServiceNow's more traditional data backup and recovery mechanism. It works in concert with AHA and acts as a secondary recovery mechanism.

Backups of the two production databases and the single sub-production database are taken every day for all instances. The backup cycle consists of four weekly full backups and the past six days of daily differential backups that provide 28 days of backups. All backups are written to disk, no tapes are used and no backups are sent off-site. All the controls that apply to live customer data also apply to backups. If data is encrypted in the live database, then it will also be encrypted in the backups. Regular, automated tests are run to provide for the quality of backups. Any failures are reported for remediation within ServiceNow.

8.16 (E) HOSTING AND PROVISIONING

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

Capgemini's Cloud Service Providers according to the Purchasing Entity's direction and agreed upon solution architecture. Capgemini's Cloud Service Providers will leverage internal tools and processes to automate the provisioning of cloud resources. The Resource Management API's and PowerShell scripting will be e utilized to automate repetitive tasks in the provisioning processes. Bash and Shell scripting will be utilized to provision Linux cloud resources. Direct access to the Azure and AWS REST API'S will also be utilized in the provisioning of cloud services on each of these cloud platforms. Virtustream will be provisioning their server environment through their Management Portal, which provides user-friendly access to automation, orchestration, and in-depth monitoring and reporting.

8.16.2 Provide tool sets at a minimum for:

16. Deploying new servers (determining a configuration for both stand-alone or part of an existing server farm, etc.)
17. Creating and storing server images for future multiple deployments
18. Securing additional storage space
19. 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).

AWS and Azure (IaaS, PaaS, SaaS)

20. Deploying New Servers – Purchasing Entity will leverage the Azure Portal and the AWS web portal to store scripts and create new servers, add them to existing farms or standalone configurations.
21. Creating and storing server images - Cloud Providers will create storage repositories to store server images, these images will be used for automated server build processes in the future.
22. Securing additional storage space - Additional storage can be allocated to accommodate the growing number of images.
23. Monitoring tools - For monitoring capabilities, the AWS and Azure web portals will be utilized, proper delegated access to individuals can be configured for segregation of duties and monitoring purposes. Monitoring agents can be installed on public hosted servers or hybrid servers to gather metrics, performance counters and system logs, all data can then be sent to a central monitoring portal for further analysis and viewing performance. Sample Monitoring Portals for Azure and AWS are provided below.

Azure Monitoring Portal

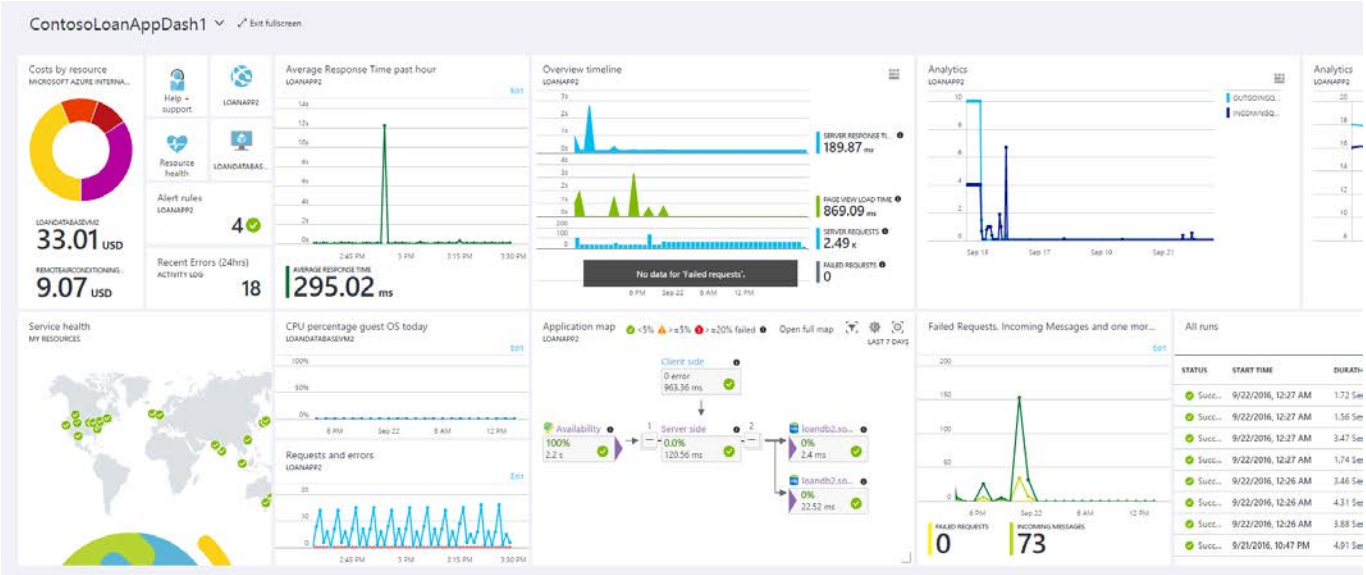

Figure 34: Azure Monitoring Portal

AWS Monitoring Portal

Figure 35: AWS Monitoring Portal

Virtustream (IaaS)

24. Deploying New Servers – Capgemini will leverage the xStream Management Portal to manage the automation and provisioning controls to create new servers, add them to existing farms or standalone configurations.
25. Creating and storing server images - Capgemini will create storage repositories to store server images, these images will be used for automated server build processes in the future.
26. Securing additional storage space - Additional storage can be allocated to accommodate the growing number of images.
27. Monitoring tools - For monitoring capabilities, Virtustream takes a holistic and comprehensive approach to system management, evaluation, and utilization. The portal view contains a variety of solution elements, each focused on a specific element of the client's environment.

Figure 36: Deploy Management Professionals

8.17 (E) TRIAL AND TESTING PERIODS (PRE- AND POST- PURCHASE)

8.17.1 Describe your testing and training periods that your offer for your service offerings.

Capgemini offers value-added services that include the implementation of IaaS hosting services (AWS, Azure, and Virtustream) via the Statement of Work's negotiated with Purchasing Entities. Capgemini will plan for ample time in the Statements of Work to conduct the necessary testing for any workloads that will be migrated to the cloud. Capgemini will also include time in the Statement of Work for the necessary training and knowledge transfer prior to transitioning cloud services to production.

Capgemini offers value-added services that include the implementation of BMC Remedy on Demand and ServiceNow. Within the context of those value-added services included in the Statement of Work, Capgemini conducts formal User Acceptance Testing and Training.

The User Acceptance Testing demonstrates the linkage between documented business requirements and the delivery of agreed functionality. Capgemini authors test scripts that tie the requirements to the expected outcome of a given test case. By observing and signing off on successful completion of the test case, the User Acceptance Tester is acknowledging the delivered work is fit for purpose.

Capgemini can deliver training in several different formats from classroom-based instructor-led training to computer-based training modules delivered through individual Purchasing Entity learning management system, complete with quizzes. The typical approach is to conduct virtual training sessions that encompass both PowerPoint presentations and a live demonstration, with the session being recorded for future use; these activities can be included in the agreed upon Statement of Work.

8.17.2 Describe how you intend to provide a test and/or proof of concept environment for evaluation that verifies your ability to meet mandatory requirements.

Capgemini, through a negotiated Statement of Work, can coordinate the stand-up of the test and proof of concept environments provided by the Cloud Service Providers.

Capgemini through a negotiated Statement of Work can work with the Purchasing Entities to determine the environments comply with any mandatory requirements.

8.17.3 Offeror must describe what training and support it provides at no additional cost.

Capgemini can include time in the agreed upon the statement of work for the necessary training, knowledge transfer and support for Capgemini cloud solutions and iPaaS implementations as a part of transitioning to production. Support will be available post-production for an agreed limited amount of time as indicated in the agreed upon the statement of work. Cloud Service Provider offerings such as Amazon, Azure, Virtustream, ServiceNow, and BMC provide limited web-based training without additional service fees.

Capgemini along with the Cloud Service providers promotes a self-learning approach through access to the iPaaS User Portal. Capgemini's learning portal contains detailed platform documentation, together with self-learning tutorials, 'How to guides' and examples. The images below provide an example of some of these 'How to guides' as a sample from Capgemini's portal, Amazon, Azure, BMC, and ServiceNow examples can be found on their product portals online.

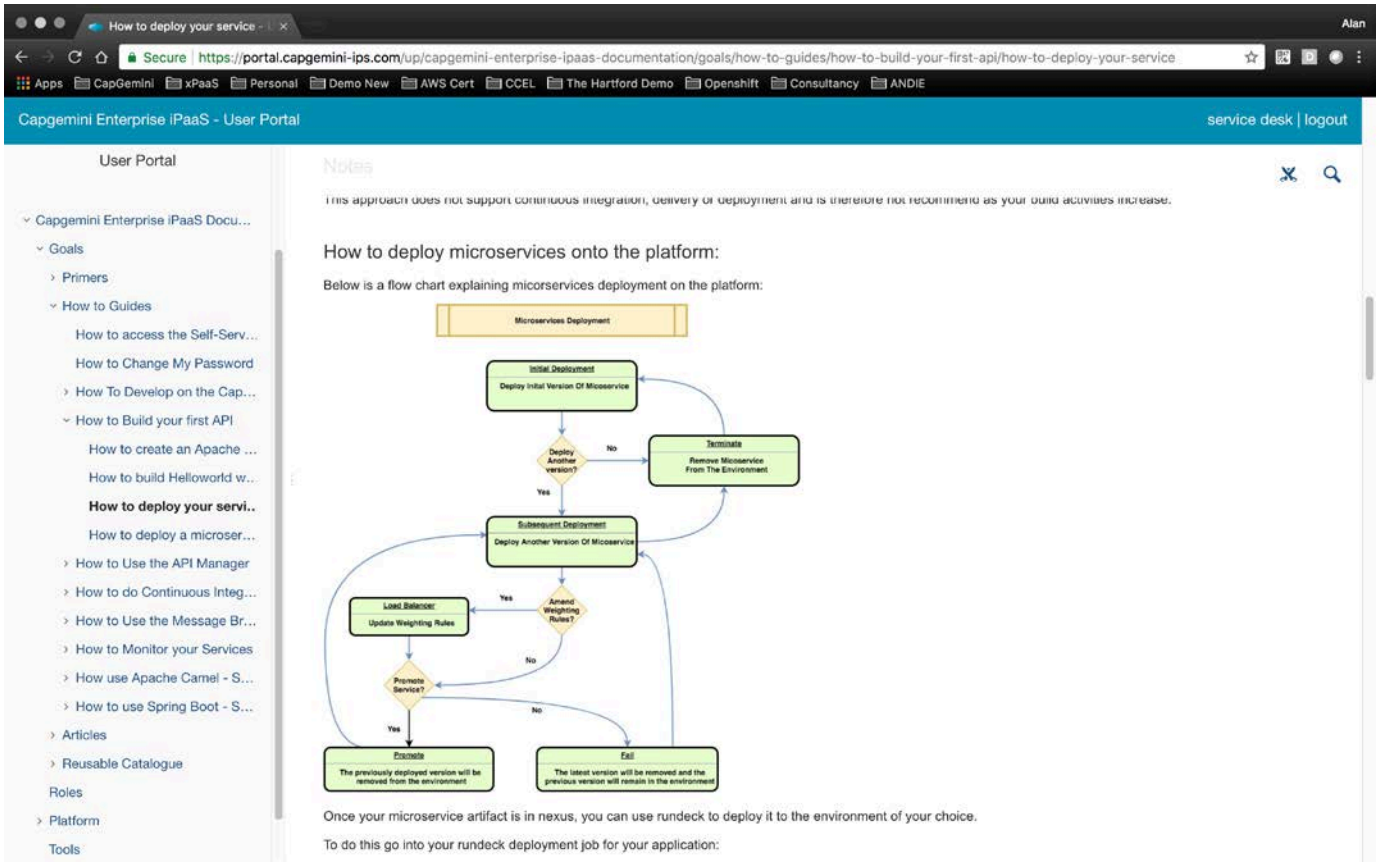

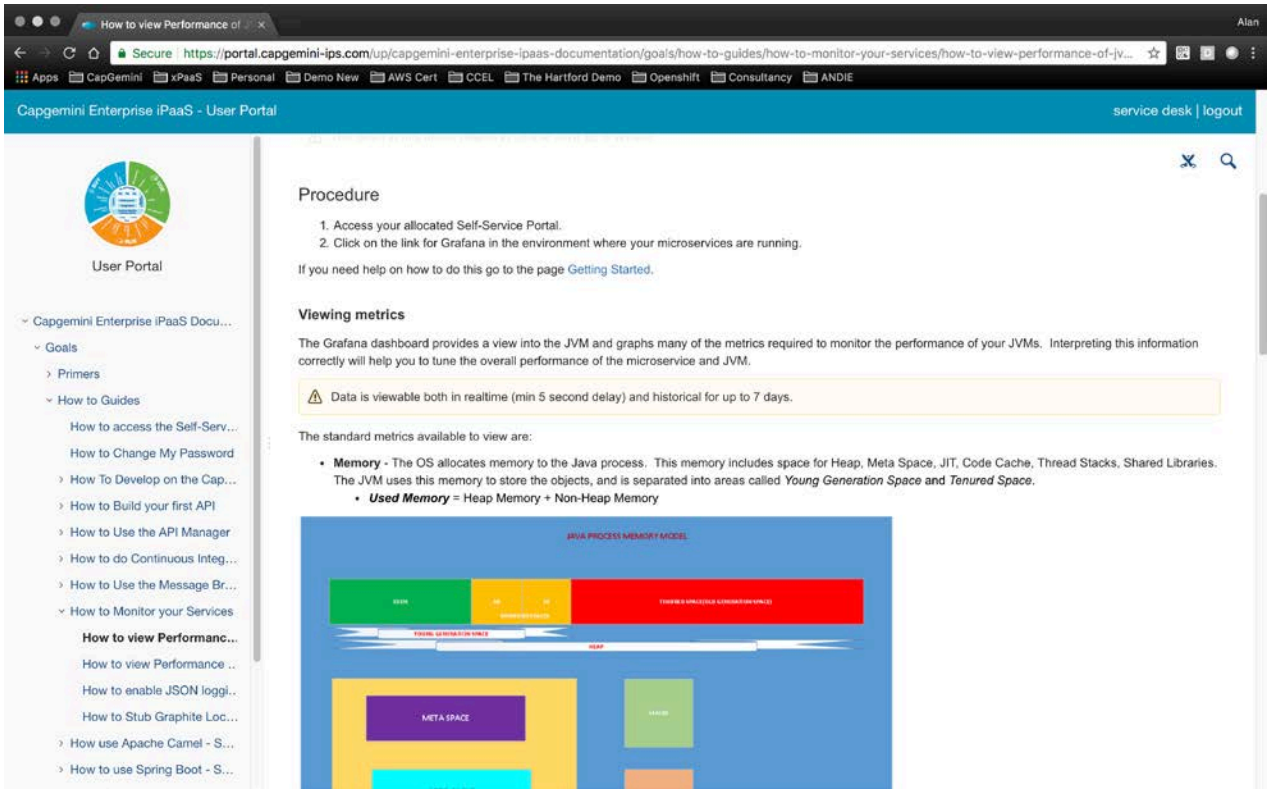

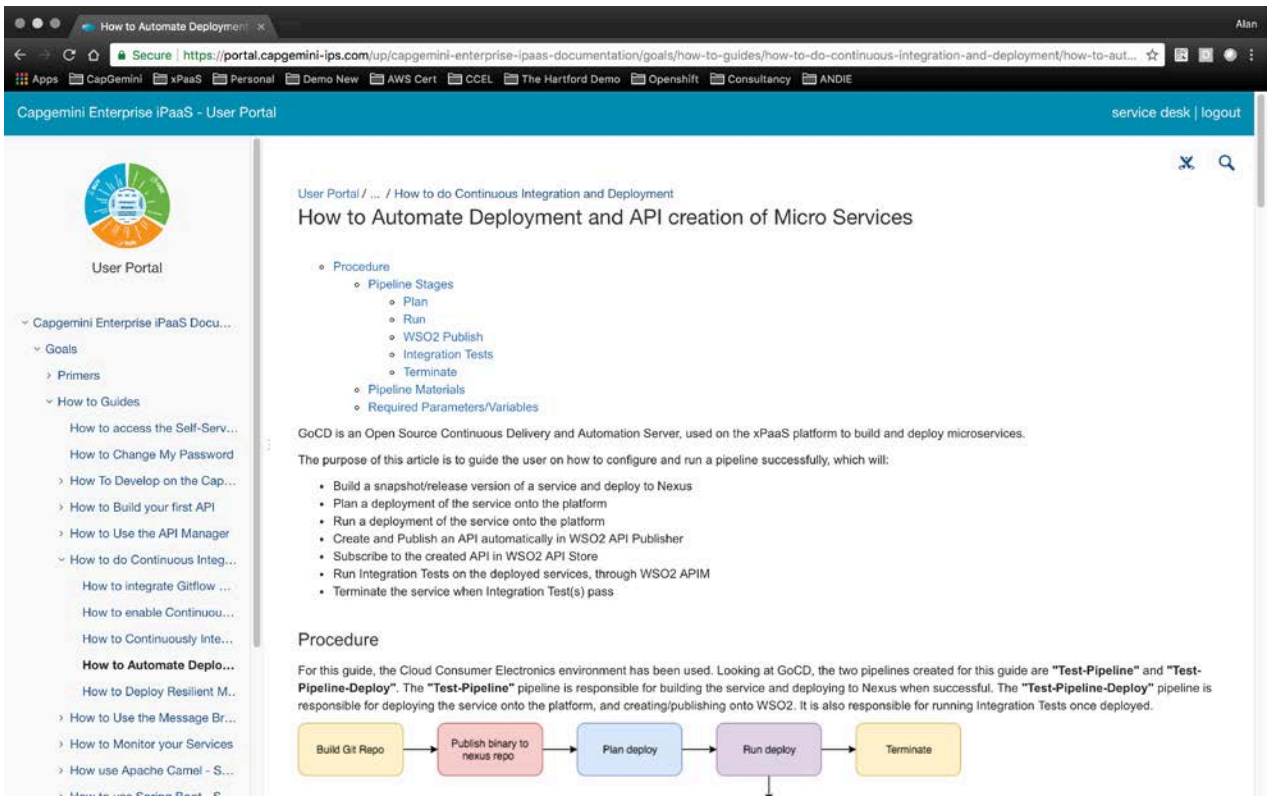

8.18 (E) INTEGRATION AND CUSTOMIZATION

8.18.1 Describe how the Solutions you provide can be integrated with other complementary applications, and if you offer a standards-based interface to enable additional integrations.

The Cloud solutions offered by Capgemini can integrate with other cloud technologies and applications utilizing REST API's for integration to ticketing systems like ServiceNow, Remedy, and other platforms.

Capgemini can engage with the Purchasing Entities to determine the type of integration required to arrive at the full value of their cloud solution spend.

Capgemini Enterprise iPaaS (PaaS)

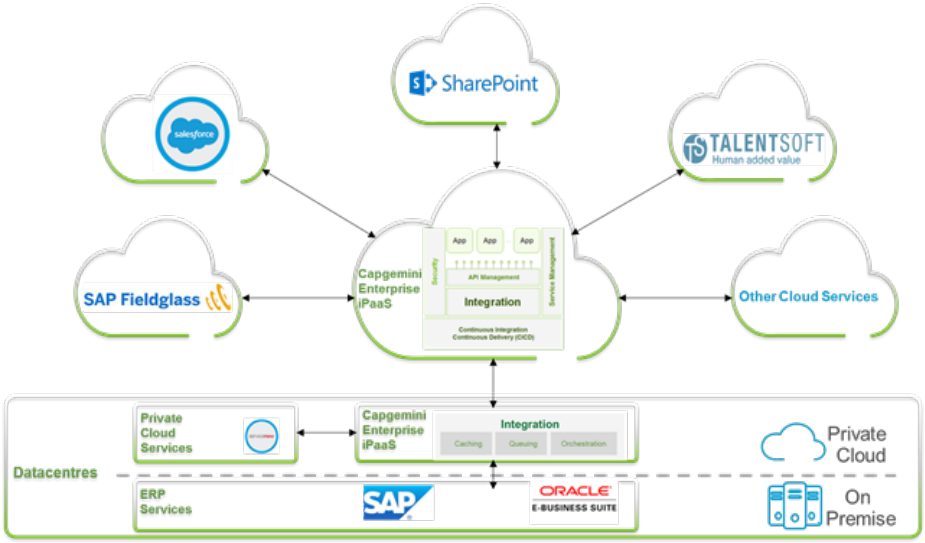

Figure 37: Capgemini Enterprise iPaaS

By essence, the Capgemini Enterprise iPaaS provide integration to a wide variety of other systems. This is achieved through our API and Integration capabilities.

API Capability

Capgemini's iPaaS platform provides an API design tool (Swagger editor) to assist with the creation of APIs. Alternatively, API definitions can be independently created within a Swagger definition and deployed to the platform through our Continuous Integration Continuous Deployment (CICD) tools.

The API design tool provides a step-by-step design process for the creation of an API. This combined with the example API provided within the platform and the user accelerators within our portal (containing concepts, architectural views, how-to guides, examples, governance models) allows for rapid and effective API creation.

Our solution also incorporates a resilient deployment of all of the components that comprise the API management and gateway functionality. These components will be deployed into multiple zones within the AWS region to provide a cloud-based presence for API management, portal, and gateways.

Integration capability

Integrations can be developed in a number of different methods, languages, and tools, however again we provide a sample application and user accelerators within the Capgemini Enterprise iPaaS Portal (containing concepts, architectural views, how-to guides, examples, governance models) to allow for rapid and effective Integration creation.

The Capgemini Enterprise iPaaS incorporates an integration platform which utilizes the Open Source Apache Camel integration framework and RabbitMQ broker software.

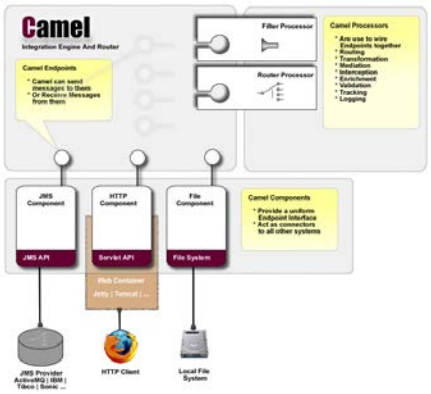

Apache Camel has an extensive list of connectors, which can be incorporated into popular frameworks like Spring, allowing for Java-based routing and mediation rules to be defined. Combined with the Advanced Message Queuing Protocol (AMQP) capabilities of RabbitMQ (message orientation, queuing, routing, reliability and security), micro-services based integrations can be quickly achieved.

Capgemini has developed platform definitions for the majority of the 65 Enterprise Integration Patterns, which allows for rapid deployment.

The Capgemini Enterprise iPaaS utilizes Continuous Integration/Continuous Deployment (CICD) tooling, combined with Cluster Management to provide resilient integration deployments.

8.18.2 Describe the ways to customize and personalize the Solutions you provide to meet the needs of specific Purchasing Entities.

Capgemini is a world-renowned IT Systems Integrator and Development company with the expertise necessary to develop native IaaS, PaaS, iPaaS, and SaaS cloud solutions for customers based on web technologies, enterprise applications, and cloud platforms. Capgemini will engage with the Purchasing Entities to determine the kind of integration or customization needed to meet requirements. For example, solutions can be developed to enable automated provisioning, self-service catalogs, monitoring and alerting, chargeback, financial reporting, and operational reporting that can be used for consumption analysis of cloud services.

In Section 8.20 Value-Added, Capgemini elaborates on the many IaaS, PaaS, and SaaS deployment, advisory, and configuration services that enable a Purchasing Entity to customize and personalize offerings from the Cloud Service Providers.

Capgemini Enterprise iPaaS (PaaS)

The Capgemini Enterprise iPaaS is provided as a managed service, therefore the platform provided is standard to all Purchasing Entities. However, within the platform there are a number of ways in which it can be customized including;

- Platform Size – Can be adjusted on a pay as you use model
- API Portal – Can be branded with corporate colors and logos
- CI/CD Pipelines – Can be created as required
- API and Integrations – Can be deployed as required

8.19 (E) MARKETING PLAN

Describe you how you intend to market your Solutions to NASPO ValuePoint and Participating Entities.

Capgemini Marketing Plan

Capgemini has built a marketing plan to advertise of the value of the services Capgemini offers through the NASPO ValuePoint Cloud Contract. Our messaging will include that this consortium contract will help states leverage the economies of scale in working with Capgemini to provide a wide range of offerings that will serve them wherever they are in their cloud transformation lifecycle. Capgemini understands that we need to bring a thorough and creative marketing plan to the state government

market that will not only carry the message but track the success of our annual marketing plan with NASPO and their Purchasing Entities.

As clients become increasingly savvy and less receptive to unsolicited marketing contact, it is becoming challenging to extract full value from traditional marketing campaigns. Together with the emergence of digital channels, these new realities are causing marketers to rethink their approach to connecting with clients. Understanding this new landscape, Capgemini plans to engage our customers through tactics such as collaboration and co-creation via personalized dialogue and multi-dimensional segmentation based on rational and emotional behavior drivers. We have experience leveraging the assets and methodologies described below through multiple government cooperative contracts. In US state government, Capgemini has been responsible for the marketing to grow the customer base on behalf of the Shared Services programs for both the Texas Department of Information Resources since 2012 and the Georgia Technology Authority since 2014. We will leverage our reputation as a trusted service provider to the states, as well as many of the same methodologies we use for DIR and GTA and other large consortium contracts we describe in our scope of experience section to market the NASPO Cloud Contract to state government.

Our marketing plan consists of the following assets and methodologies:

- Our People
- Annual Plan
- Marketing Channels
- Marketing Campaigns
- Lead Capture and Conversion

Our People

Capgemini has an experienced Sales Teams, cloud subject matter experts, and Creative Services Teams that work together to quickly create professional collateral for events and marketing campaigns. Our Corporate Marketing Team creates messaging around the offerings that we have listed in our catalog for social media and our corporate website.

Our experienced sales team is led by our Executive VP who has been actively engaged in the state government market with Capgemini since 2011. Our team regularly attends events and individual customer meetings. We draw on the expertise of hundreds of solution architects and collaborate with our clients on the best way to transform their legacy environment to an agile cloud environment that will securely serve the citizens of the state. Capgemini understands that many State CIOs are driving the transformation to a secure cloud environment for state government to augment the state's data centers through a shared services model. Capgemini has focused our efforts on coaching state CIOs on how to transform their shared services organizations into an agile, secure, cloud-enabled digital ecosystem for their agencies. We will continue developing these relationships and with the award of the NASPO contract, expand our reach out to additional agency CIOs.

Our Creative Services team has expertise in creating multiple types of collateral for our sales team that include case studies, email campaigns, tradeshow and brochure collateral. We intend to leverage our broader Capgemini North American marketing team and the campaigns and blogs created for our global public and private Capgemini customers that align with the cloud services we have proposed in the NASPO contract. Additionally, we will leverage our strong social media collateral developed by our account based creative services team along with our Capgemini Corporate Marketing team to evangelize our NASPO service catalog out to state government through our marketing channels.

Here are three examples of the quality of social media marketing that we will leverage out to our state government clients:

The way we assemble and build services with AWS is like Lego bricks, says AWS VP Cloud Architecture Strategy Adrian Cockcroft in the latest Cloud Choice podcast.
https://goo.gl/8wQ42C

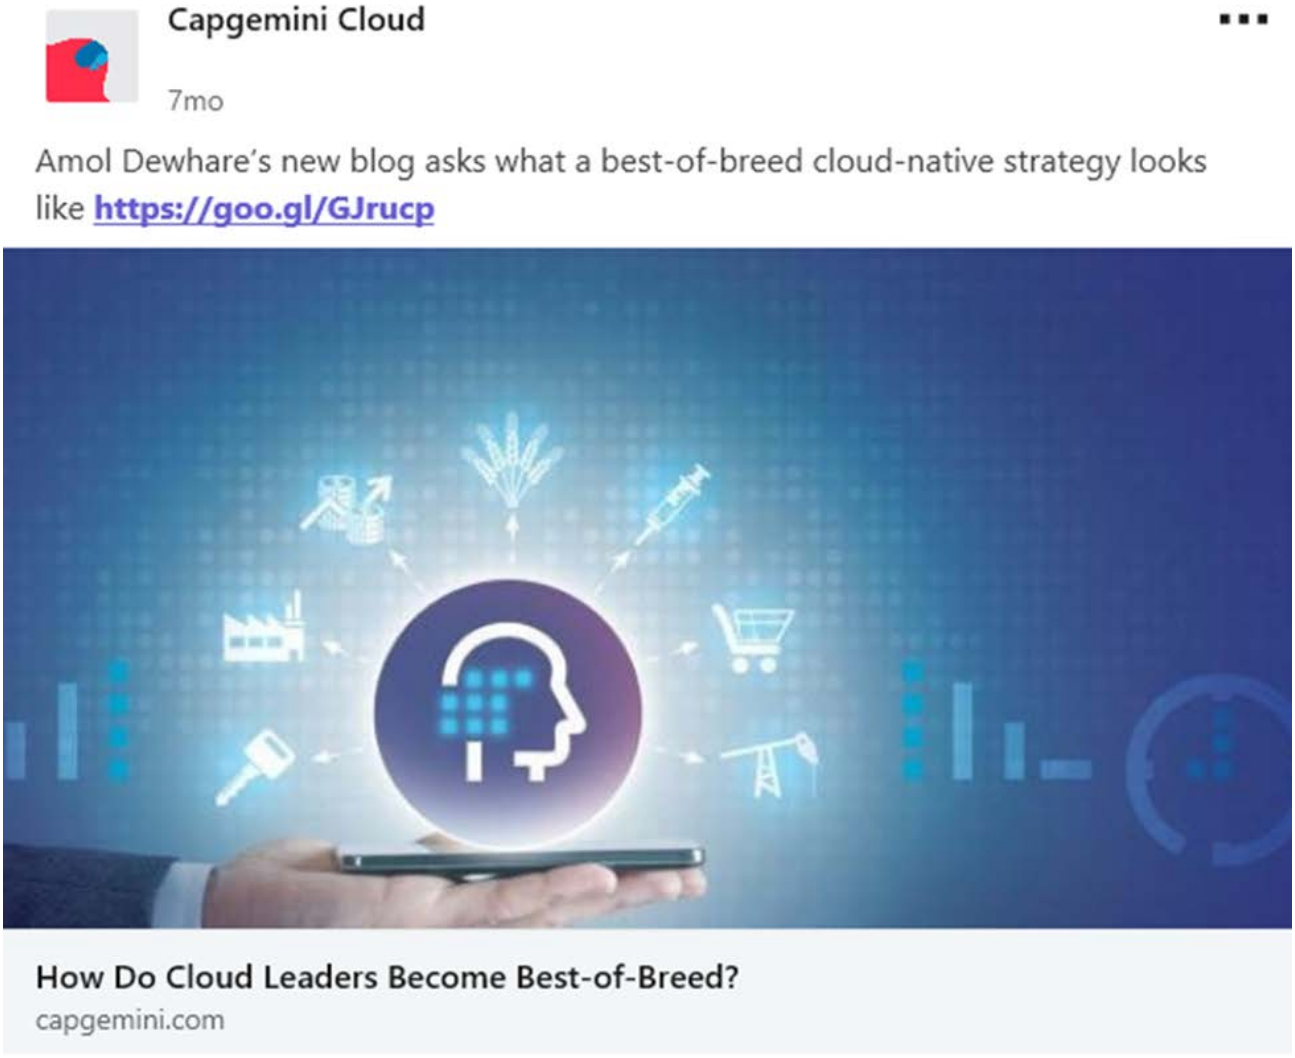
Capgemini Cloud
7mo
Amol Dewhare's new blog asks what a best-of-breed cloud-native strategy looks like https://goo.gl/GJrucp
How Do Cloud Leaders Become Best-of-Breed?
capgemini.com


Cloud-native podcast series
Your guide in five episodes
In these episodes we discuss all aspects of cloud-native transformation: Why do companies need to innovate, why are they choosing cloud native, and what are the implications for CIOs? How can enterprises build the right teams, skills and culture for continuous delivery, and what does the future look like? Capgemini experts Dan O'Riordan and Rene Claudio interview thought leaders from Pivotal, Forrester, Cloud Foundry Foundation, Uber and AWS.
Capgemini
The Cloud Ch...
SOUNDCLOUD
5
TRACKS

Annual Plan

The key to a successful marketing plan is to have a clear, repeatable and actionable annual plan. Our annual plan consists of three phases of subsequent marketing activities.

Annual Plan	
Identify the customers	▪ Review the success of the previous year market segmentation (geographic, vertical, etc.) ▪ Identify key customers, highlight success stories, identify cross-sell/upsell opportunities for the following year
Define the Plan	▪ Analyze market trends and customer base to align potential services with the target segmentation to promote campaigns ▪ Identify new services/bundles for promotion ▪ Define campaign cadence for the calendar year
Review the previous year results	▪ Conduct review of all campaigns from the previous year ▪ Identify processes for standardization ▪ Highlight Success factors

Marketing Activities	
Overall Vision	▪ The overall vision, strategy, and direction for campaigns
Service Priority	▪ A prioritized list of services and capabilities highlighted for that given year
Marketing and Awareness Campaigns	▪ Plan multiple annual NASPO Cloud Contract Awareness Campaigns ▪ The plan which campaigns will be launched when
Channel Strategies	▪ Determine which campaigns will utilize which channels
Measurement	▪ Metrics from Salesforce will be used to measure the success of each campaign and the type of marketing vehicle

Marketing Channels

Part of our annual plan is to determine the proper marketing channels for our clients, as there many channels that can be used to connect with state agencies. To maximize the effectiveness of communications, we have found that multiple communications should be used to market new services. These marketing channels may be used in isolation, or as part of a campaign to progress stakeholders along the commitment curve. The following is a summary of channels we intend to use to market Capgemini NASPO Cloud Services:

- Press Releases
- Social Media
- Landing page on Capgemini public facing website in our public-sector section dedicated to NASPO contract that will capture traffic and direct traffic to the NASPO Cloud Contract site
- State Government and Partner Conferences
- Email Campaigns
- Alliance Partners
- Webinars

- Podcasts and Vidcasts
- Product and Service descriptions and other collateral

Capgemini sponsors both national and local government events. Capgemini has been an active member and sponsor of state focused technology events including but not limited:

- National Association of State Technical Directors
- National Association of State CIOs
- State Specific Digital Summits
- SLED CIO Roundtables

We look forward to adding the following to our annual state government event circuit:

- NASPO ValuePoint Events

In addition to state government focused conferences, Capgemini is a proud sponsor of many of our partner conferences as well. We intend to bring NASPO Cloud Contract collateral to our alliance partner conferences that we sponsor (Microsoft, AWS, Virtustream, ServiceNow, and BMC) to carry the message that these services can be easily procured through the NASPO Valupoint Cloud Contract. Examples of partner trade show that we sponsor which are client facing trade shows where we will provide NASPO Cloud Contract collateral to state Purchasing Entities that attend:

- AWS Reinvent
- Microsoft Ignite
- ServiceNow Knowledge18
- Dell World for Virtustream
- Exchange for BMC Remedy

We also attend the Microsoft inspire partner facing event in which we do account planning and strategy sessions with the Microsoft sales force. We have dedicated alliance managers working closely to build the SLED market. Capgemini is looking forward to showcasing our NASPO Cloud Contract status and portfolio to the state government clients and the partner account teams at these important trade shows on an annual basis.

Annual Marketing and NASPO ValuePoint Cloud Contract Awareness Campaign Strategy

As part of our marketing effort, Capgemini will create awareness campaigns to state government clients as part of our annual marketing plan and activities. Our holistic approach starts with understanding what customers really want and then designing a customer experience and campaigns that meet those customer needs. We know from our experiences in a variety of public and private industries that customers care about more than just cost. Customers tend to weigh several factors including:

- Ease of access
- Customer service
- Information and support availability
- Quality of services being offered
- Overall experience in working with the service provider
- Trust in the service provider
- Service provider reputation and ability to deliver on promises
- Alignment between service offerings and customer goals/vision

With these customer needs in mind, and our annual plan structure, we leverage the various marketing channels described above to conduct campaigns. Each campaign has iterative planning and execution pieces that revolve around content gathering / creating, creative strategy, and delivery. As with every initiative, each campaign closes with the retiring of collateral and an analytical review of performance.

At the end of each year, the campaign plan will be reviewed. Success will be metrics based. Based on results of the previous year, our team will determine the plan of action for the following year. As part of the planning phase, we would propose an annual campaign review with NASPO to gain feedback on our methodology and leverage any broader insights that NASPO can provide to help us better serve state government.

Lead Capture and Conversion

Our internal Capgemini Salesforce CRM solution allows us to track the creation of leads within each campaign, manage those leads to the point of creating new opportunities, and then track ongoing touch points with contacts we have at client accounts. Our platform includes Pardot – Marketing Automation, Sales Cloud, and Sales Lightening. We have also integrated various other tools such as DiscoverOrg (for data hygiene) and LinkedIn Sales Navigator into Salesforce, to help us populate richer data in the system that we will leverage to help market our portfolio in the NASPO Cloud Contract out to the agencies.

Our system meets the needs of inside sales, outside sales, marketing, sales operations, and executives and can support any applicable strategy that we evolve over the 8 remaining years of the NASPO contract. Through Salesforce we can manage our NASPO specific campaigns, using standardized templates and processes that provide consistency in reporting campaign metrics. We also create dashboards and reports that measure the success of the campaigns that span channels like events and digital marketing.

Capgemini will be able to share metrics with NASPO to help track the value of Capgemini in the State Government Market. We will leverage Salesforce for the following marketing functions:

- Measure the success of marketing channels leveraged for campaigns
- Capture and manage leads and convert them to opportunities and contacts within accounts
- A clear understanding of our pipeline
- Provide summarized data, reports, and dashboards

Capgemini's holistic marketing plan has the people, the proven methodology, and the market-leading tools needed to effectively market the NASPO contract for cloud services to state government. We are excited to carry the message to the states that there is an easy contracting method that will allow them to quickly procure the services they need to become a truly agile IT services provider.

8.20 (E) RELATED VALUE-ADDED SERVICES TO CLOUD SOLUTIONS

Describe the value-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 in offering activities with initial setup, training, and access to the services.

The following table depicts the high-level categories of the fee-based, value-added professional services Capgemini will provide in conjunction with the NIST compliant Cloud offerings. Please see the **SK18008 – Capgemini - Detailed Product Offering Document** for individual cloud service and product offers, summarized in this section.

Value-Added Services	IaaS	PaaS	SaaS
Cloud Value-Added Services	▪ Capgemini Cloud Advisory Services ▪ Capgemini IaaS Deployment Services ▪ Capgemini Cybersecurity Services	▪ Capgemini iPaaS Services	▪ Capgemini Service Management Advisory Services ▪ Capgemini Configuration Services ▪ Capgemini SaaS Deployment Services

IaaS

Capgemini Cloud Advisory Services

Covering multiple facets of digital transformation, along with solid references from global enterprises in all industry segments, combined with a technology-agnostic, partnership-oriented approach that gives clients full access to the best solutions and guidance the industry has to offer, Capgemini can meet the Purchasing Entities at any stage in their cloud transformation lifecycle. We provide support in developing a strategy, workload assessment, and adoption change management services.

Cloud Strategy

An assessment which will engage with all areas of the business to determine the following.

- The extent to which the Purchasing Entity or the agencies the Purchasing Entity serves has already adopted 'Cloud', be that Public or Private or both.
- What are the challenges with the existing service and where can Cloud adoption help and/or potentially hinder?
- How is Cloud perceived by these entities? Is there potentially a requirement to educate as to what is possible, what is probable and what the associated risks are based on Capgemini's experience with other clients.
- Evolution versus Revolution – where on the spectrum does the Purchasing Entity sit. What are the drivers and appetite for change and what may impede progress? Is the adoption of the cloud driven by a need to remove the risk of legacy hardware, to consolidate workloads onto a more cost-effective platform or is the business looking more to innovate by embracing cloud-native technologies such as Mobile, IOT, and Big Data. In short, what type and rate of change would be most palatable/appropriate for the business.
- And post-adoption, what assistance the client may need to support and/or manage the technologies ongoing both from a functional and cost perspective.

The output of this is then assessed against Capgemini's client base, especially those in the public sector with similar challenges. Experience gained, and lessons learned from other engagements will inform the optimal approach and set realistic expectations. The strategy will enable the business to adopt Cloud technologies in a manner that is measured, pragmatic and demonstrably beneficial to the business.

Cloud Workload Assessment

Part of the Capgemini Cloud Choice Service Portfolio – “Managing Cloud Mass Migrations”, focuses on managing client cloud migration projects. Capgemini offers functions and teams that manage the entire project – from advisory and strategy to assessment, migration design and planning, and management of the migration itself.

Within the Cloud Workload Assessment service is Capgemini Cloud Migration Assessment Services focusing on cloud assessment and platform architecture, delivering cloud proof of concepts, migration, integration, and testing for IaaS, PaaS, and SaaS Solutions.

Capgemini Cloud Deployment Services

Cloud Migration

Capgemini Cloud Migration services facilitate the planning and quality for the Purchasing Entities cloud migrations, to provide business buy-in and stakeholder alignment. The Capgemini Cloud Migration services are performed through the following phases:

28. Assessment and strategy phase typically carried out to develop the transformation roadmap of initiatives, public cloud, private or hybrid cloud selection identified, architecture definition, Process establishment (SDLC), Security & regulatory guidelines
29. Detailed advisory phase, deep dive assessment and migration preparation phase, for cloud migration programs, including identification of cloud migration candidates (fast track and complex), retirement and legacy candidates. Enterprise & target architecture definition and design, including security standards.
30. Planning, management, and execution phase, of cloud migration programs, including execution of all pre-requisites (target environments, upgrades, patches, transformation) as well as the migrations themselves using multiple delivery units.

Data Center Transformation Services

The Capgemini Data Center Transformation Offering consists of the following services:

31. Data Center Advisory Services focusing on strategy and advisory services – to establish a reputation and lay the foundation for subsequent successful phases
32. Data Center Transformation / Optimization services – upgrades, modernization, automation programs (Client DC to Client DC or Cloud)
33. Data Center Relocation / Consolidation / Migration services – moving or consolidating multiple data centers to new locations, including cloud (Client DC to Client DC or Cloud)

Virtustream Private Cloud Deployment

Private Cloud – enable the on-demand provisioning of enterprise-grade workloads on standard technological platforms

- General purpose applications on x86 platforms
- SAP, including HANA (on SAP RUN, and DCS)
- Oracle RUN
- Big Data Lake aaS (Insights aaS)
- A technological Cloud Platform roadmap leveraging the market leading practices and enabling Hybrid
- Data encryption at rest safeguarded and Improved SLA for cloud services
- Integration in the customer ticketing system
- Services can be provisioned on-demand with Pay as You Go model

Application Migration Services for Rehosting in the Public Cloud – Simple, Medium, High, and Very High Complexity

This migration service provides for basic "lift and shift" operations of simple complexity applications. We define migration complexity based on weighted criteria spanning ten (10) different evaluation components. Rehosting activities within the migration process provide a like to like copy of the solution environment from the source to the target locations. The operating system, application architecture and most configuration components of the targeted solution will remain the same, with changes limited to the configuring the solution to run in the new environment.

Application Migration Services for Redeploying Solutions in the Public Cloud – Simple, Medium, High, and Very High Complexity

This migration service provides for solution re-platforming of simple complexity applications. We define migration complexity based on weighted criteria spanning ten (10) different evaluation components. Redeployment activities within the migration process include the creation of a new hosting environment, installation of a fresh OS instance, and re-deployment of the application into the new environment. The application architecture of the targeted solution will remain the same, with changes limited to the configuring the solution to run in the new environment.

Capgemini Cybersecurity Services

Capgemini offers a complete range of cybersecurity services to guide and secure the digital transformation of companies and Government organizations. Our more than 4000 professionals support you in defining and implementing your cyber security and cloud transformation strategies. We protect your data, Information technology. Industrial systems and the Internet of Things (IoT). We have the resources to enable secure cloud transformation, strengthen your defenses, optimize your investments and control your risks in cloud environments. We provide cybersecurity services to Infrastructures, Applications, Endpoints, user, and data in the cloud, and our Research & Development team that specializes in malware analysis and forensics help our customers to stay ahead of advanced and sophisticated threats. Capgemini has ethical hackers, 10 Security Operation Centers (SOC) around the world, a licensed Information Technology Security Evaluation Facility, and are a global leader in the field of testing.

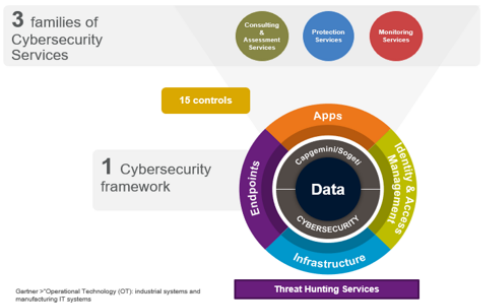

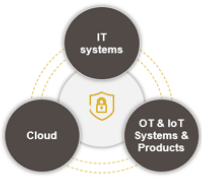

Securing the digital world against cyber-attacks and malicious internal behavior through our end-to-end portfolio covering IT, industrial systems and IoT products. Our services are based on four areas of activity: we advise, protect, monitor, and hunt.

Advise (Consulting & Assessment Service): We confirm that our clients' cybersecurity strategy is fit for purpose and in line with both their appetite for risk and their budget. From cybersecurity maturity and health assessments to roadmaps, risk assessments and information asset inventories, including security controls such as pen-tests and audits, our consulting services are designed to help organizations make the right choices about what to prioritize and where to invest.

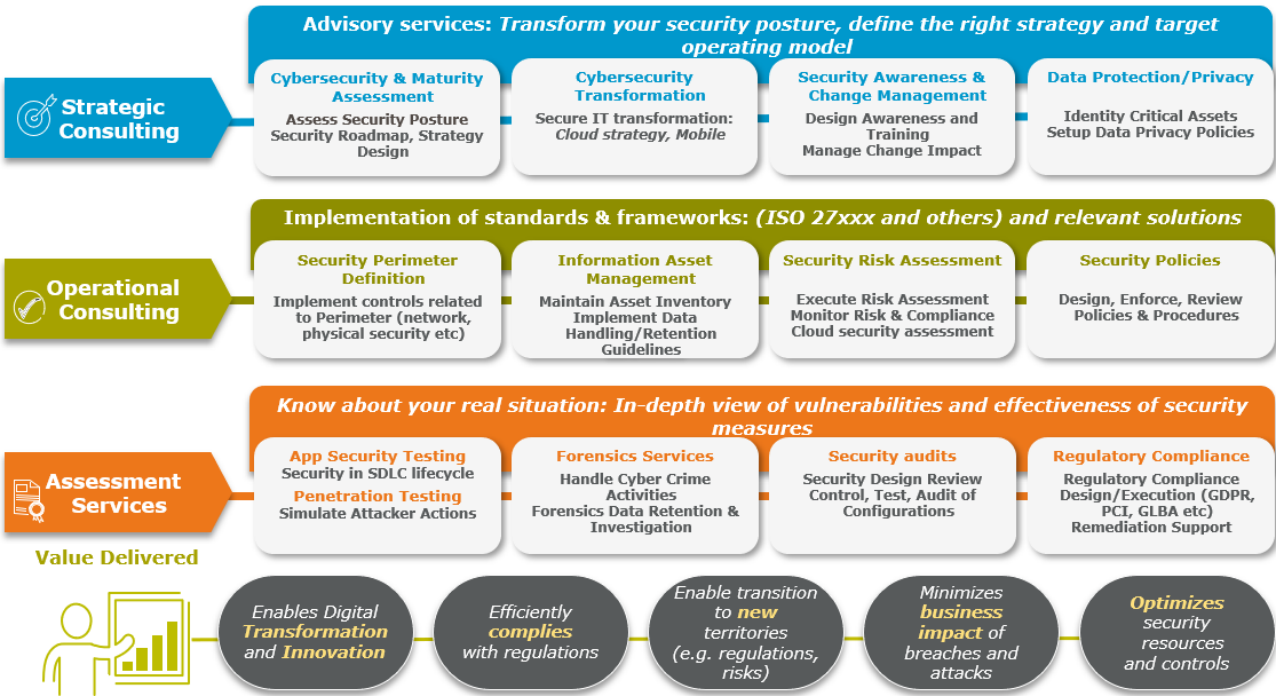

Protect (Protection Services): Powered by the best technology providers and delivered by domain experts, our protection services are designed to keep our clients ahead of the game by securing their Infrastructures, Applications, Endpoints, user, and data in the cloud and on-premise data centers leveraging our Partners and leading technology vendors.

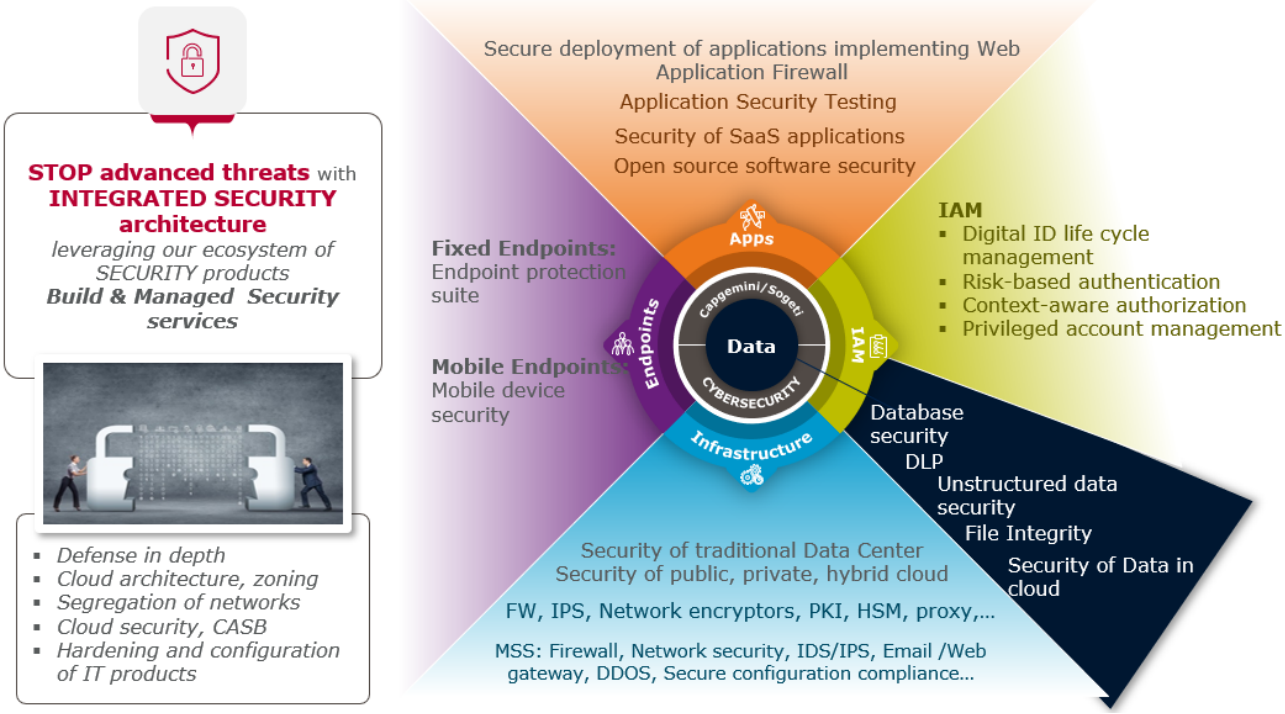

Monitor (Monitoring Services): Organizations gain situational awareness of how their security controls are operating and the threats they face with our security monitoring services. They will detect and react efficiently to cyber-attacks.

PaaS - Capgemini iPaaS Services

Capgemini Enterprise iPaaS

The Capgemini Enterprise iPaaS team offer a set of defined consulting and delivery services, which are defined below. We have then used these defined services throughout the solution phases to express the professional service engagement. Capgemini Enterprise iPaaS comes with three add-on services that can help the Buyer to fast-track the Buyer's business outcomes

Accelerate Service

A Capgemini Enterprise iPaaS team who enable a project team to use the platform and quickly deliver an agreed business outcome that showcases how the Capgemini Enterprise iPaaS can transform a business.

- This comprised a small team of expert resources
- A team to work directly with the Purchasing Entries' integration development team to build the first set of integrations and assist to speed up the understanding of how to use the Capgemini Enterprise iPaaS.
- Allows the Purchasing Entries' project team to use the platform and quickly deliver business benefit.

Connect Service

A Capgemini Enterprise iPaaS team who will enable the project team to connect the Capgemini Enterprise iPaaS with your existing business and technical processes, configure the Capgemini Enterprise iPaaS in-line with a project's requirements, implement an operating model compatible with the existing service management approach and deploy the Capgemini Enterprise iPaaS to a private cloud.

- This comprised a small team of expert resources
- A team to work with the Purchasing Entries' infrastructure team to connect the iPaaS with existing business and technical processes, configuring it to the project requirements, implementing an operating model compatible with the existing service management approach, and deploying to a private cloud.
- Allows the Purchasing Entries' project team to use the quick connect to the required systems and processes.

Expert Service

A Capgemini Enterprise iPaaS team made up of expert resources, tailored for a specific need. This may be either, supporting the project team in acquiring the necessary knowledge and skills to become effective in the use of the Capgemini Enterprise iPaaS, or working with the project team to identify and elaborate opportunities for their business that will exploit the Capgemini Enterprise iPaaS.

- This comprises an expert resource
- Under the Expert service, the Offeror can provide expert resources to support the Purchasing Entries' project team in acquiring the knowledge and skills to maximize value from the iPaaS or work with the Purchasing Entries' project team to assist in identifying and elaborating opportunities to use the iPaaS to deliver desired business outcomes. As part of the above services, the Supplier can provide people with expertise in the following:

- API and integration development and testing
- Capgemini Enterprise iPaaS architecture
- Automation of networking and infrastructure deployment
- Secure integration of cloud and on-premises services
- Agile methods
- DevOps
- Requirements analysis

- Allows the Purchasing Entries' project team to use the platform and quickly deliver business benefit.

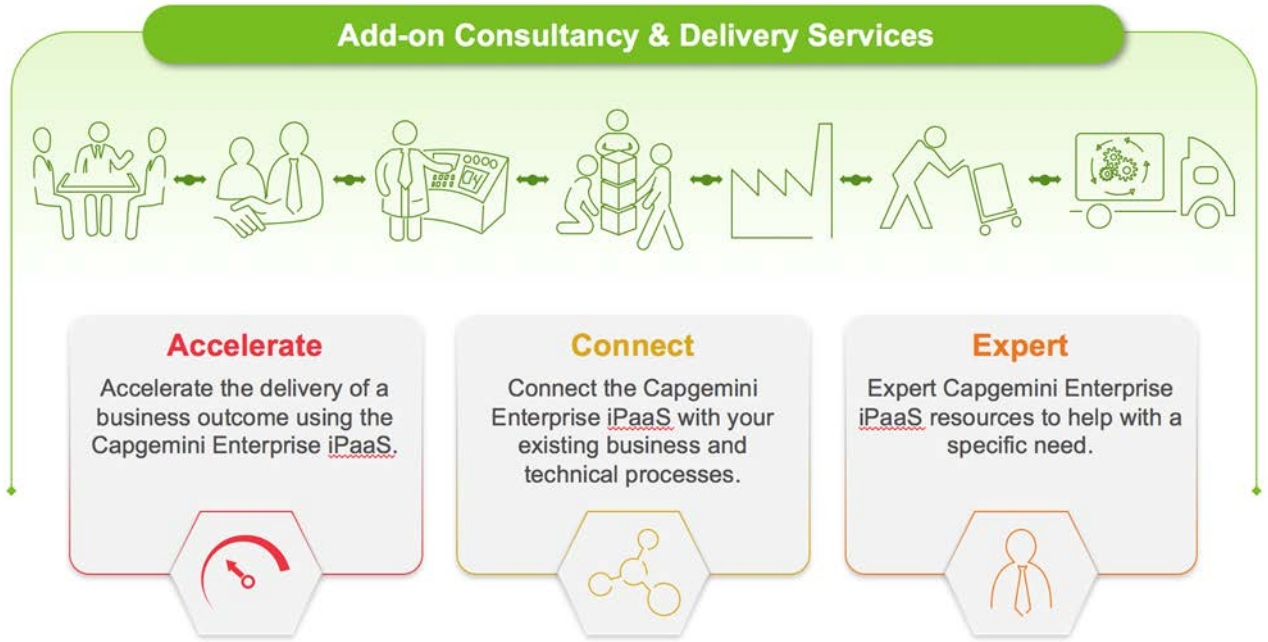

Figure 38: Add-On Consultancy & Delivery Services

SaaS

Capgemini's Software as a Service (SaaS) catalog selections provide consultation, design, implementation, and post-deployment coaching services to support our Cloud-based offerings. Capgemini SaaS value-added services enable our clients to accelerate their software selection process while equipping them to maintain an objective and rigorous evaluation process throughout, mitigating the risk of making a wrong decision. Incorporating an effective service management tool solution is an integral component to implementing the Service Integration and Management (SIAM) framework of technology and processes.

Capgemini's definition of SIAM:

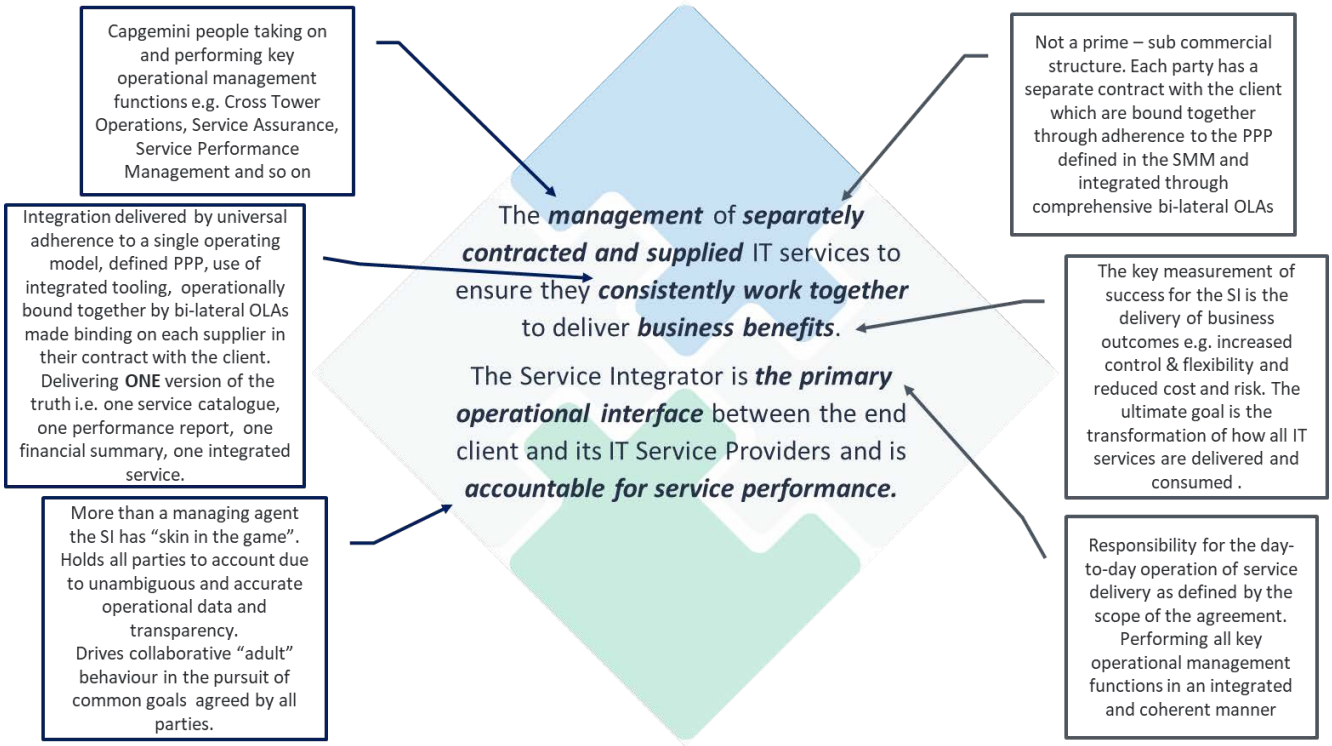

Figure 39-Capgemini's SIAM Definition

Capgemini Service Management Advisory Services

Capgemini Service Management Advisory Services Service Asset and Configuration Management

Prior to migrating your assets to the cloud, Capgemini recommends understanding the assets that comprise your ecosystem. It is a critical step that each Purchasing Entity must make when they are thinking of a cloud transformation project – "what assets do I have, and where are they?" Establishing your starting point is the first, essential step in any effort at improving your situation. Performing a Service, Asset, and Configuration Management (SACM) Assessment is the step Capgemini recommends before proposing any other SACM services.

Capgemini can effectively execute a SACM Assessment to determine the IT environment health by defining the scope, deliverables, governance, and reporting of IT assets.

To successfully conduct a SACM Assessment, the following activities must be executed:

- Look at the current state of the IT Asset Management practices and the related processes including purchasing, financial and IT Service Management processes
- Identifies the sources (tools, physical and virtual locations) that will provide IT Asset data for future phases
- Reviews the people and cultural aspects of the environment for practices that increase overall asset program risk
- Technology enablement focusing on the tools that provide the discovery data and integration with consuming systems as to whether they provide the correct information. The asset repository tool itself, and follow-on reporting is reviewed for suitable functionality.
- Hardware and Software Asset Inventory to assess the level of accuracy and completeness.

Capgemini's SACM Assessment provides a report of findings, recommended improvements, potential cost savings, and a proposal for achieving sustainability of the IT asset estate.

Capgemini Service Management Advisory Services SIAM Assessment

Many CIOs across the country are looking at moving to a shared services model by implementing a SIAM framework. Conducting a Service Management Advisory Services SIAM assessment is a smart first step in that move. Engaging Capgemini to perform this assessment will help the CIO understand their Organization's people, process, and technology readiness to support build a cloud-enabled shared services platform. A platform that will deliver cloud services back to their organization through a digital, cloud-based marketplace.

Capgemini's SIAM Assessment is performed by executing diagnostic exercises to analyze the state of maturity and identify gaps against the target SIAM framework depicted in the following figure.

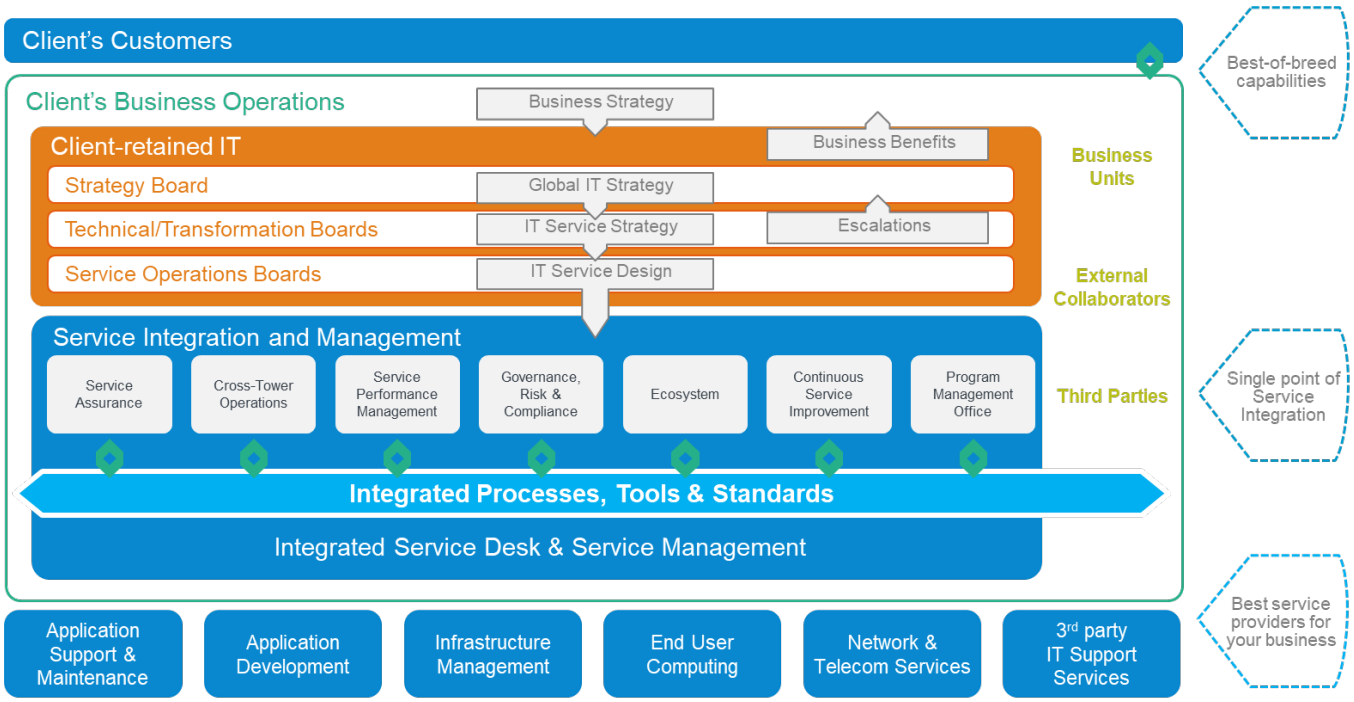

Figure 40-Capgemini's SIAM Framework

Effective implementation requires execution of the following activities:

- Discover = discovery of current state, business drivers, priority issues, desired outcomes, assessment objectives and approach. Information is captured to conduct a gap analysis.
- Assess = analyze operational and organizational effectiveness, ecosystem delivery capability, gap analysis, change impact and readiness, and quick fixes.
- Recommend = execute transformation and actionable plans, change roadmaps, operating model change, stakeholder plan, and governance model.

Capgemini's SIAM Assessment provides the CIO with guidance, recommendations for improvement, a concrete action list, and a high-level transformation plan.

Capgemini Service Management Advisory Services SIAM Coaching

After Capgemini has helped the Purchasing Entity design and implement a cloud-based, shared services SIAM framework, you may be in need of additional coaching/training to operationalize your shared services model. You need a way to reinforce the daily, weekly, and monthly cadence of SIAM activities until they are adopted as the accepted standard.

Capgemini understands your situation and offers a SIAM Coaching service designed to place this specialized skill set within your organization. Our SIAM Coaches will both educate your retained staff on SIAM methods and reinforce the cadence of SIAM activities. Our SIAM Coaches will work with your teams for the agreed period of time in the statement of work and assist with instruction, guidance, reinterpretation, and organizational change management. Our coaching approach reduces the risk of not attaining your strategic IT goals and leveraging your shared services model.

Capgemini SIAM coaches work in several capacities across your IT Organization to transform understanding, change behavior, and deliver a steady level of support.

The Capgemini SIAM coach's will:

- Teach SIAM methodologies, processes, and activities.
- Help create an environment focused on delivery and facilitating continuous improvement.
- Assist with process adoption and identify improvements required.
- Showcase enabling tools and techniques.
- Engage with stakeholders at all levels of the organization.
- Provide for key metrics tie SIAM activities to Business outcomes and are accurately reported.
- Equip Customer with the ability to coach future Providers.

SIAM Coaching provides the direction, guidance, and reporting your team needs to confidently adopt new processes and procedures with reporting evidence that their efforts are making a difference.

Capgemini Configuration Services

Capgemini Configuration Services Cloud Provisioning, Monitoring, and Automation

Purchasing Entity organizations are looking for leading IT Service Providers who can expand IT services to address increasingly complex IT requirements. IT Providers must be able to manage and broker IaaS and SaaS cloud functions to enrich and control the Customer Experience.

Capgemini's Cloud Provisioning and Automation services integrate your existing Service Catalog tools with Cloud Providers using our developed service connectors and Application Programming Interfaces (APIs) to seamlessly translate your Business requests into an expansion of your ecosystem and IT capabilities on an unattended basis.

Capgemini's Cloud Provisioning supports business transformation initiatives by delivering a business architecture that reflects decision making, budgeting, and supports Cloud services. Capgemini`s automation service provides design, consultation, business analysis and service management in support of the design and implementation of Cloud-based services; replacing or reducing human effort in these types of processes; helps to enable the delivery of accurate and fast results, providing a cost-effective integration approach.

Capgemini's Cloud Provisioning and Automation services integrate your Service Catalog and Cloud Platforms by:

- Automatically provisioning Cloud services with or without approval workflow.
- Streamlining the order-to-activation process for your LOB and Customers.
- Provisioning services from different cloud instances to present to a single marketplace.
- Automatically updating your CMDB with what was provisioned.

Upon successful provisioning and automation integration of the specific catalog items with the Cloud Provider services, Capgemini will turn over all documentation and implementation notes to your IT Support Team for ongoing support.

Capgemini Configuration Services Performance Management and Reporting

Purchasing Entities moving to a cloud enabled Service Management tool may need to have support to create the right level of reporting for different stakeholders. The services performance and reporting team may also be new to a consumption based reporting data which is a key different from traditional data center hosting versus newer cloud billing and performance reporting.

In our Value-Added Services Offering of Service Performance Management & Reporting, the Purchasing Entity will get a unified view of how their cloud enabled multi-provider infrastructure is a performing. Engaging Capgemini to implement this function provides the Purchasing Entity's Service Management organization with processes, technology, and training to deliver flexible, integrated performance reports and insightful analysis to enable management of cross-tower operations, contracted service levels, and suppliers' management.

The Service Performance & Reporting (SP&R) Team performs the following functions:

- Receive service reports from all suppliers and consolidate the data to produce and publish daily, weekly and monthly operational performance reports and dashboards
- Provide performance trending and analysis across service towers

The SP&R Team will manage Suppliers through service level reports and contractual compliance visibility:

- Publishes performance of all service providers against SLA targets
- Participates in monthly Supplier Management sessions
- Defines and maintains the linkage from KPIs to SLAs
- Defines and maintains cross-tower Service Reports and dashboards
- Continuous Service Improvement opportunity identification and support

Capgemini's implemented Service Performance Management & Reporting function coordinates the drive towards sustainable service improvements with a focus on outcomes, supported by better performance management culture, systems and reporting.

Capgemini Configuration Services Cloud Cost Management & Chargeback

Transformation programs are comprised of many moving parts. Two drivers to move to a cloud-based ecosystem are flexibility and cost optimization. Managing costs and services can be quite challenging for organizations that have not had experience with a hybrid cloud environment. Capgemini understands the extreme importance for an efficient level of Cost Management. Financial focus is achieved by first understanding what challenges you have today, what immediate needs are required, what future requirements should be considered, and what is the ultimate outcome desired.

Capgemini will collaborate with you to create, design, and implement IT cost chargeback tools and processes that allow you to allocate IT costs to the Purchasing Entity's consuming organization.

- Define resource allocation and resource units
- Determine unit costs
- Determine resource unit consumption by organizational unit
- Implement effective and accurate invoicing, minimizing inaccurate invoices

All billing artifacts are pulled together to provide a clear allocation IT costs and who is consuming them.

Engaging Capgemini to implement your IT Cost Chargeback function allows you to calculate, analyze, allocate, and invoice IT costs back to the consuming Organization.

Capgemini Configuration Services SACM Sustainability Model

In a hybrid environment, it is critical for Purchasing Entities to have a single source of truth for assets in their ecosystem, both physical and virtual. Capgemini had developed the Service Asset and Configuration Management (SACM) Sustainability Model to establish and maintain a single source of truth.

The Sustainability Model has been created to help guide Asset Managers on how best to sustain the quality and accuracy of the data in the CMDB. This Model requires all components to be implemented and followed and is not effective if not implemented in its full format.

The SACM Sustainability Model is a set of interrelated processes, tool enhancements, reports, and leading practices, including recommended corrective actions. The scope of Sustainability Model covers both SACM tasks as well as interfaces with other services that impact or benefit from SACM.

The key components of the Sustainability Model include:

- Working interfaces with key processes – Change Management and Procurement
- Sustainability Reporting
- BAU Trackers for SACM Team
- Waterfall Methodology and Reporting
- SACM Documentation and Work Instructions
- Common Standards and Data Model
- An Audit Plan

Engage Capgemini to implement and turn over to your SACM Team the SACM Sustainability Model and realize the benefits of:

- Reliable data available for use in Incident and Change Management
- Identification of CIs critical to IT Service continuity, information on dependency and configuration
- Reliable billing validation
- Confirmed return of leased assets
- Sustainable model providing for continued service improvement
- KPI Matrix and Dashboard

Capgemini SaaS Deployment Services

Deploying Service Integration and Management (SIAM)

Moving between Service Management Ticketing Tools is a major endeavor. Whether the move is to upgrade versions, re-tool for increased functionality, cost-cutting measures, or to implement Service Integration and Management (SIAM), Capgemini has the resources to perform this work for you.

Using with our mature project methodology, our experienced SIAM Project team can deploy the cloud-hosted Service Management Systems listed below. As our Project Team works through the research, design, configure, test, and deploy project phases, they will:

- Collect requirements and data
- Analyze and normalize foundation and master data,
- Design and build the configuration
- Test the configuration, including a User Acceptance Test with your Business Users
- Deploy to production

With this project methodology, we can implement any of the below SaaS/Module combinations:

You receive the system configured, populated with foundation and master data, and tested to meet your business requirements. Included in the delivery of the system, is the delivery of our world class SIAM Processes and Work Instructions tailored to your specific requirements and ready for consumption by your teams.

We complete the implementation with Process Training on a "Train the Trainer" basis to teach your teams how to conduct the daily activities laid out by the processes, using your new system.

Capgemini SaaS Deployment Services for Ivanti Heat

- Incident/Problem/Change Management
- Self Service/Service Catalog
- Knowledge Management
- Configuration Management Database/Configuration Management
- Service Level Management.

Capgemini SaaS Deployment Services for Cherwell

- Incident Management
- Problem Management
- Change Management
- Self-service Portal
- Knowledge Management
- Configuration Management
- Database Management
- Request Management
- Asset Management
- Service Level Management
- Survey

Capgemini SaaS Deployment Services for ServiceNow

- Incident/Problem/Change Management
- Request and Service Catalog
- Asset and Configuration Management Database
- Knowledge Management
- Service Level Management
- Surveys

Capgemini SaaS Deployment Services for BMC Remedy

- MyIT Service Catalog
- Incident Management
- Problem Management
- Change Management

- Self-service Portal
- Knowledge Management
- Configuration Management
- Database Management
- Work Order Management
- Asset Management
- Service Level Management
- Survey

Capgemini SaaS Deployment Services for BMC Discovery

Undocumented hardware not only exposes your company to financial risks but to security risks as well; servers that do show up on the roles are seldom accounted for or patched. Fixing the situation is very difficult when most asset discovery tools or services require you to know about the devices before the tools can start tracking the CIs.

BMC Discovery is an agentless discovery tool that will find the objects in your environment whether you know about them or not. Discovery will work across your public, private, and facility-based environments to combine detailed device information, installed software titles, and which devices are communicating with one another

Combining all this data, Discovery provides insightful analysis of your ecosystem enabling an organization to resolve several key business blockers through:

- Correcting CMDB inaccuracies
- Updating the CMDB automatically
- Accurately showing Configuration Drift
- Exposing Shadow IT
- Answering security audits and providing other control mechanisms

Capgemini's Discovery professionals have the experience to engineer and execute a deployment across your organization that will integrate Discovery with your ITSM CMDB, update your ITSM CMDB on an unattended basis, and schedule reports to run as needed. They will conclude the project to put Discovery data into your Team's hands by conducting knowledge transfer sessions with your Asset Management Team, going over the prepared documentation

8.21 (E) SUPPORTING INFRASTRUCTURE

8.22.1 Describe what infrastructure is required by the Purchasing Entity to support your Solutions or deployment models.

- The typical infrastructure required for a Purchasing Entity are typically, OSI Stack – High Available Network Layers 1-3:
- Network connectivity, logical or physical, routers, switches, cabling at the Purchasing Entity site or Virtual data Center at the Cloud Service Providers.
- IP addressing and naming systems, and/or servers at the Purchasing Entity site or Virtual data Center at the Cloud Service Providers.
- FW rules logical and/or firewall appliances at the Purchasing Entity site or Virtual data Center at the Cloud Service Providers.

- Load balancer systems at the Purchasing Entity site or Virtual data Center at the Cloud Service Providers.

Capgemini, through an agreed-upon Statement of Work, will set up and configure the connectivity to Cloud Service Providers according to the Purchasing Entities network security requirements, reference the diagrams below for a configuration of Cloud Service Provider connectivity.

AWS Direct Connect (IaaS, PaaS, SaaS)

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an Amazon Virtual Private Cloud (VPC) using private IP space while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

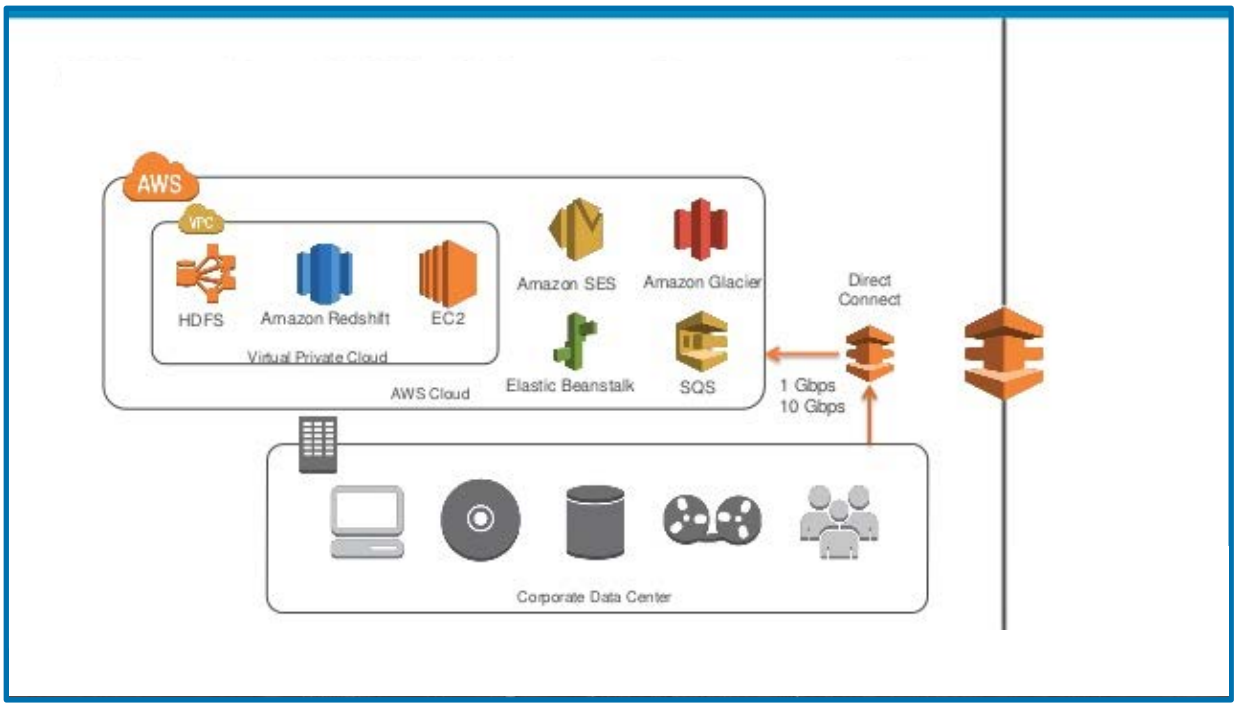

Figure 41: AWS Direct Connect

Azure ExpressRoute (IaaS, PaaS, SaaS)

Azure ExpressRoute is used to create private connections between Azure datacenters and infrastructure on your premises or in a colocation environment. ExpressRoute connections don't go over the public Internet, and they offer more reliability, faster speeds, and lower latencies than typical Internet connections. In some cases, using ExpressRoute connections to transfer data between on-premises systems and Azure can give you significant cost benefits.

With ExpressRoute, establish connections to Azure at an ExpressRoute location, such as an Exchange provider facility, or directly connect to Azure from your existing WAN network, such as a multiprotocol label switching (MPLS) VPN, provided by a network service provider.

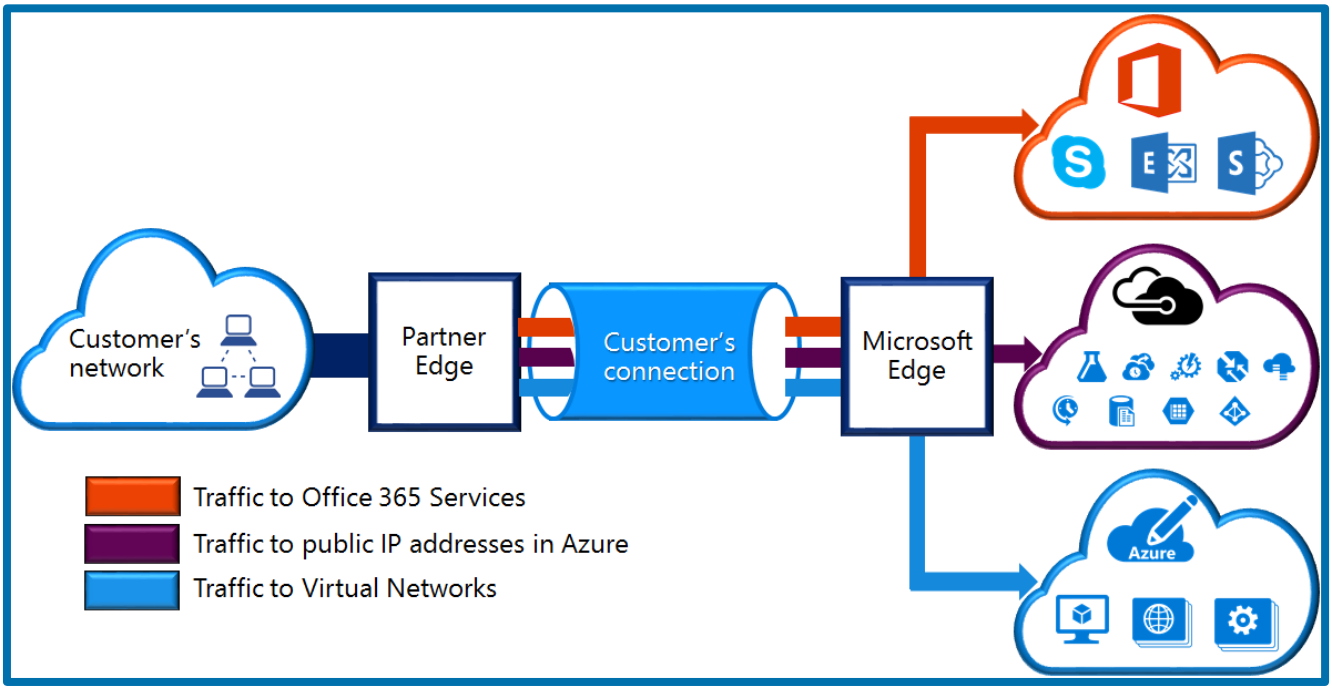

Figure 42: Azure ExpressRoute

Virtustream Network Connectivity (IaaS)

Virtustream architects their solution as an extension to a customer's current network infrastructure. Purchasing Entities will have the capability to directly connect to Virtustream from their existing WAN network, such as a multiprotocol label switching (MPLS) VPN, provided by a network service provider. The Virtustream environment is seen as a separate node within the client's ecosystem, with private IP addresses available to isolate servers, applications and other solution components.

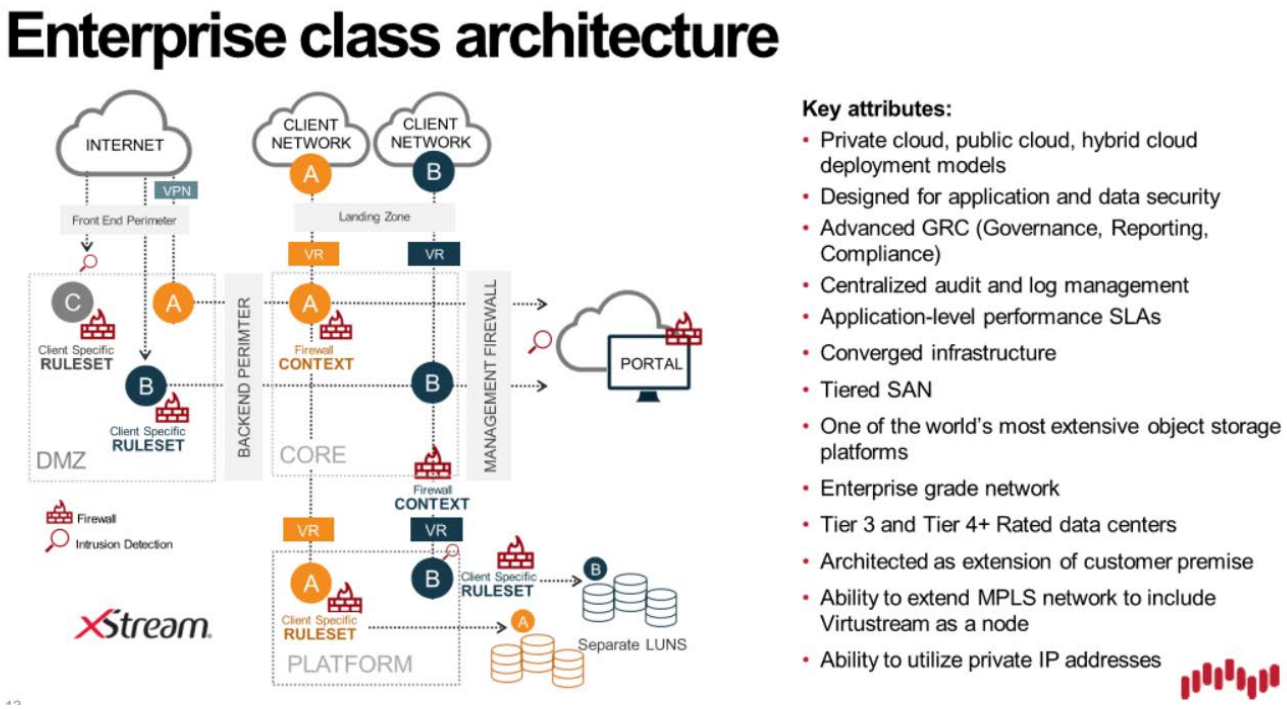

Figure 43: Virtustream Enterprise Class Architecture

BMC and ServiceNow (SaaS)

The BMC Remedy on Demand and the ServiceNow SaaS solutions do not require additional infrastructure to configure or use the services as they are offered.

The value of both solutions would be greatly enhanced by building integrations with your other systems. These integrations typically involve making a server to server data connection between applications residing within your trusted network and the SaaS services residing in the cloud. Both providers can facilitate this connection in a secure manner without requiring a VPN and the associated complexity, cost, and time penalties associated with VPN architectures.

The connections are accomplished by installing integration software on a server in your network that can communicate with your applications and can connect to the internet for access to the cloud. A very simple view of this connection is depicted in the following figure.

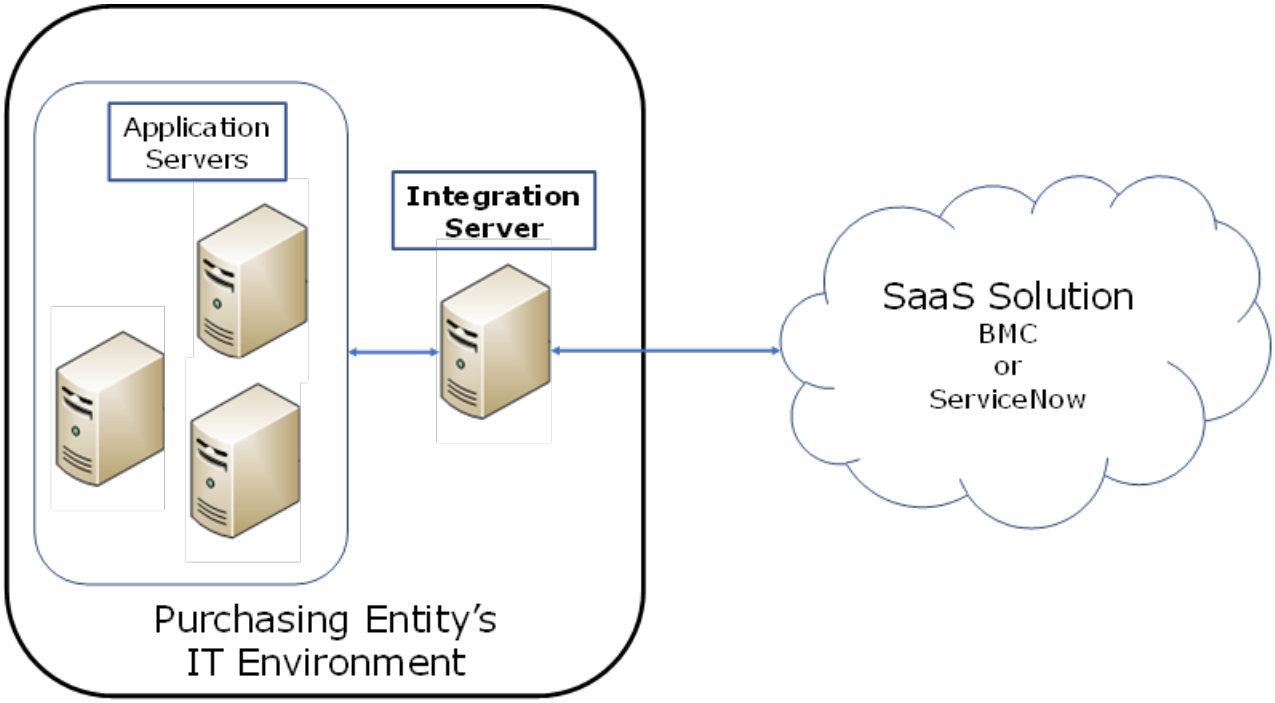

Figure 44: BMC and ServiceNow Connectivity

The system requirements for the integration server will vary depending on the number and nature of the integrations implemented. Minimum system requirements are determined by which SaaS solution is deployed.

SaaS Solution	Operating System	Java version	CPU	Memory	Storage
BMC Client Gateway	Windows Or Linux	JSDK 8, 64 bit	2	4GB	40GB
ServiceNow Mid Server	Windows Or Linux	JRE version 1.8.	1 quad-core	5GB	40GB

8.22.2 If required, who will be responsible for the installation of new infrastructure and who will incur those costs?

- Capgemini can be responsible to build out the cloud infrastructure and connectivity as part of a negotiated Statement of Work.
- There are recurring monthly costs with establishing the direct network connectivity, which will be the responsibility of the Purchasing Entity.

- The monthly payments for the direct network connectivity are made to Capgemini's Cloud Service Providers or Purchasing Entity Network providers depending on negotiated SOW.
- If the customer purchased direct network connectivity services directly from Cloud Service Providers, then payment is made by the purchasing agent directly to the Cloud Service Provider.

9. Appendices

9.1 Appendix A – Role Descriptions

The following Role Descriptions support the projects for consulting, design, and implementation that the Purchasing Entities will leverage in a negotiated statement of work with Capgemini. Capgemini will utilize a mixture of roles to perform the various tasks, subprojects, create and review deliverables necessary to execute the value-added services in section 8.20 RELATED VALUE-ADDED SERVICES TO CLOUD SOLUTIONS. These roles are also priced in Attachment F Cost Proposal Form under the value add section.

Role Name	Role Description
Service Delivery Manager	The Service Delivery Managers reporting to the Capgemini's Relationship Manager are responsible for the day to day delivery of work streams. These individuals escalate appropriate issues to the Service Delivery Executive and manage service issues and changes in accordance with the adopted change management procedures. Specifically, the Service Delivery Managers manage the services delivered onsite, onshore, and offshore service locations. Manage the Capgemini personnel providing service under this agreement Manage service level performance and implement improvement initiatives where appropriate On-boarding and off-boarding of Capgemini project team members Manage issue resolution including escalation to the Service Delivery Executive on an "as needed" basis Perform quality assurance on projects and operation activities Work with delivery teams to introduce innovation into the service delivery process and assist the implementation of agreed to innovative ideas Participate in the Executive Committee Lead the Scope Steering Committee for Capgemini. Bring team members with necessary subject matter expertise to participate in activities, analysis, and meetings.
Program Manager	Program managers confirm master plans and schedules are followed, developing solutions to program challenges, and directing others for successful completion of the project on time and on budget. Manages people and/or services delivered from a unit via Program Office to the client. Making sure that the services are provided to an appropriate quality to meet contractual obligations and that targets are met. Co-ordinates delivery of a portfolio of services across technologies Meets the defined SLA/OLA's Ramp-up or run-down of services and resources, deploying transition managers where necessary Provide for correct and appropriate charging of the program budget allocations As a respected and trusted advisor, supports the customer's initiatives by deploying standards, providing for timely delivery of projects Thorough understanding of all contractual obligations and confirms they are met.
Project Manager	The Project Manager provides the central focus for the project with overall responsibility for the delivery of the project. He is responsible for delivering a quality end product that meets the requirements of the client. He controls the project and is fully committed to and accountable to the Project Executive Sponsor for its success. The Project Manager confirms that (i) the project meets the agreed requirements.

Role Name	Role Description
	Delivers products on time and to budget to the agreed level of quality and (ii) the delivered system works in accordance with the requirements specification. Sets up and manages the administration of the project, including support, documentation and project records Assembles the project team and making sure that they are all aware of their responsibilities. The project manager establishes a strong working relationship with the client and its customers. The Project Manager establishes requirements, documenting requirement and managing changes to requirements throughout the entire project. The Project Manager Develops the Overview Project schedule, supporting plans/schedules, and project management products, obtaining agreement and keeping them up to date throughout the project. Defines, schedules, controls and adjusts tasks for the project. Maintains tight financial control, both for man-days and costs, so that the project is completed to the agreed budget and consistent with quality. The Project Managers manages the project risks, including the development of contingency plans. Monitors and report progress to the Project Executive Sponsor/Steering Committee. The Project Manager establishes and operates change control in accordance with Project Management procedures. Organizes the implementation and handover. Provides effective communication within the project team, with the Project Executive Sponsor/Steering Committee, with users, and with any other interested parties. The Project manager along with the project executive sponsor organizes and carry's out the Project Evaluation Review and participation in the Post Implementation Review.
Transition Manager	The Transition Manager will have primary responsibility for managing the team and all activities related to transferring delivery responsibilities from Capgemini to Purchasing Entities. The Transition Manager will interact with Purchasing Entities personnel and Capgemini executives sponsoring the transition effort. Develops the detailed transition plan encompassing transition activities for the Help Desk, Deskside support, Network Support, Server Support, ITIL Process Implementation, Reporting and the Governance Committee Manages the transition work plan Assigns and manage transition team members Leads transition team meetings and prepare weekly status reports Identifies transition issues and discuss these with Purchasing Entities and Capgemini leadership Prepares a Standards and Procedures manual specifically defining how service will be delivered to Purchasing Entities
Project Analyst	The primary responsibilities of a Project Analyst are performing, analyzing and providing project analysis and support to the entire project team reporting directly to the Project Manager. Creates, managing and disbursing reports related to the project. Maintains project assets, communications and related database(s). Evaluates and monitors the overall project. Reviews and reports the project's budget and finances. Routinely performs complete or component analysis. Notifies the entire project team about abnormalities or variances.
Team Manager/Leader	Team Manager is responsible for the overall project direction and day-to-day activities of the Technical Team Members. Creates and develops performance report, delivery method, the scope of work, and general duties records. Helps in managing customer demands to provide for maximum satisfaction, and to maintain quality over quantity. Defines customer requirements and assist in creating comprehensive technical

Role Name	Role Description
	documents Supervises activities between internal and external resources, and facilitate smooth workflow for service delivery Evaluates project data for accuracy, and takes the lead in setting project targets and priorities Resolves disputes between team members and management and assist in addressing issues amongst team members to avoid unacceptable behaviors Gives training and mentorship to team members to make them better on the job Reviews customer technical demands and instructions and assist internal team and customers to identify effective methods for delivering technical solutions
Operations Manager	Plays an interface role between Service Delivery Manager and the Delivery Team– that provide the different services to the client. In this role, the Operations Manager is accountable for delivering the run services towards a client specific project deliverable. Produces periodic reports to update company management on the progress of operations Confirms compliance with all company policies and procedures when delivering customer projects Liaises with other departmental heads to plan and implement action plans for improved technical operations Maintains an up-to-date knowledge of technical processes, functions, and requirements for customer projects
Senior Delivery Architect	Senior Delivery Architects define and provide for a comprehensive and coherent view across Business, Information, Systems, and Technology, not just to design IT systems but to deliver Business Change which may also be supported and enabled by IT. The Senior Delivery Architect Director works with the CIO/CTO of an organization to align business and technical requirements for projects.
Infra Architect	Infrastructure Architects design and implement information systems to support the enterprise infrastructure of an organization. They provide that all systems are working at optimal levels and support the development of new technologies and system requirements. Infrastructure Architects generally lead and direct a team and report directly to top management. Infrastructure Architects leads the design and review processes for new systems. They develop and document the proposed technical design for the integration and implementation of any new software, working across the IT department.
Cloud Lead Architect	The Cloud Lead Architect is one of the senior project roles. Provides that the project builds will meet the technical Strategy and the associated benefits. The Cloud Lead Architect is responsible for designing and implementing enterprise cloud infrastructure and platforms required for cloud computing. Analyzes system requirements and confirms that systems will be securely integrated with current applications. Has a deep understanding of system development in cloud environments, including Software as Service (SaaS), Platform as Service (PaaS), or Infrastructure as a Service (IaaS).
Cloud Network Architect	The Cloud Network Architect is one of the senior project roles. Provides that the project builds will meet the technical Strategy and the associated benefits. The Cloud Network Architect is responsible for designing and implementing enterprise cloud network infrastructure and platforms required for cloud computing. Analyzes network requirements and provides that networks will be securely integrated with current workloads and cloud-native applications. Has a deep understanding of networking and cloud environments, including Software as Service (SaaS), Platform as

Role Name	Role Description
	Service (PaaS), or Infrastructure as a Service (IaaS).
Cloud Infrastructure Lead	The Cloud Infrastructure Lead is one of the senior project roles. Working with the Cloud Lead Architect provides that the project builds will meet the technical Strategy and the associated benefits. The Cloud Infrastructure Lead is responsible for the day-to-day activities and designing and implementing enterprise cloud infrastructure and platforms required for cloud computing. Analyzes system requirements and provides that systems will be securely integrated with current applications. Has a deep understanding of system development in cloud environments, including Software as Service (SaaS), Platform as Service (PaaS), or Infrastructure as a Service (IaaS).
Cloud Administrator	The Cloud Administrator is responsible for cloud day-to-day operations. The Cloud Administrator configures and fine-tunes cloud infrastructure systems. Installs and configures virtual cloud instances. Support cloud servers including security configurations, patching, and troubleshooting. Establishes Virtual Private Networks (VPNs) to customer environments. Develops scripts for automating client/server functions Monitor automated systems recovery solutions
Cloud Infrastructure Engineer	The Cloud Infrastructure Engineer is a key role to implement Cloud Infrastructure Technologies. Working with the Cloud Infrastructure Architect provides that the project builds will meet the technical Strategy and the associated benefits. The Cloud Infrastructure Engineer is responsible for the day-to-day operations and maintaining and designing enterprise cloud infrastructure and platforms required for cloud computing. Configures system requirements and provides that systems will be securely integrated with current applications. Has a deep understanding of infrastructure systems and cloud environments, including Software as Service (SaaS), Platform as Service (PaaS), or Infrastructure as a Service (IaaS).
Cloud Network Engineer	The Cloud Network Engineer is a key role to implement Cloud Infrastructure technologies. Works with the Cloud Network Architect to provide that the project builds will meet the technical Strategy and the associated benefits. The Cloud Network Engineer is responsible for the day-to-day operations and maintaining and designing cloud network infrastructure and platforms required for cloud computing. Configures cloud network requirements and provides that networks will be securely integrated with current workloads and cloud-native applications. Has an understanding of networking and cloud environments, including Software as Service (SaaS), Platform as Service (PaaS), or Infrastructure as a Service (IaaS).
SIAM Architect	The SIAM Architect is the lead role for a given SIAM Statements of Work and is responsible for the successful outcome of the SIAM project. This role requires at least five years of experience delivering SIAM solutions, proficient with SIAM tools, processes, and integration points. They must be able to lead a multi-disciplinary team in a fast-paced, dynamic environment.
SIAM Workstream Lead	The SIAM Workstream Lead is the lead for one of several SIAM Workstreams that participate in a SIAM Implementation; Tools, Incident/Problem Management, Request Management & Service Catalog, Change Management, Service Asset and Configuration Management, Service Performance & Reporting, and Chargeback. The Workstream Lead is responsible for the successful outcome of their stream and for the necessary integrations with other streams. This role has at least 3 years' experience in their respective stream having participated on several previous implementations. This role directs the SIAM Process Analysts.
SIAM Infrastructure Lead	The SIAM Infrastructure Lead is responsible for the successful deployment of the toolset they have been assigned to deploy, such as the BMC or ServiceNow toolsets. This role has at least three years of experience in their chosen tooling. This role directs the SIAM

Role Name	Role Description
	Administrators.
SIAM Administrator	The SIAM Administrator role is responsible for developing and configuring the SIAM tools being deployed. They receive direction and requirements from the SIAM Infrastructure Lead.
SIAM Process Analyst	The SIAM Process Analyst role is responsible for working directly with Clients to gather business requirements, and rendering those business requirements into Process and Configuration functional requirements. This role also writes documentation, updates test scripts, and organizes foundation data or normalizes asset data depending on their area of assignment.
Database Architect	Database Architect is responsible for the design, structure, and maintenance of data, organized in a relational database. The Database Architect provides that the accuracy and accessibility of data relevant to the customer's data requirements for a project. Provides data structure and management for customer projects. Leveraging advanced skills with data-oriented computer languages such as SQL and XML.
Database Administrator	The Database Administrator (DBA) is responsible for managing and coordinating all database activities. The DBA responsibilities include database design, user coordination, backup, recovery, overall performance, and database security. DBA regularly performs routine tests and modifications to provide that a database is performing and running correctly. The DBA troubleshoots the programs and hardware. Based on the findings, repairs or changes can be made to fix the problem. The DBA routinely discusses and coordinates security measures with the Database Architect and Management Team.
Cybersecurity Architect	Cybersecurity Architect plays an important role to enable business through Information technology by making secure security is considered in every step of digital transformation and operations. On a high-level Cybersecurity Architect is responsible to Design, build and implement enterprise-class security systems, Align standards, frameworks and security with overall business and technology strategy, Identify and communicate current and emerging security threats, Design security architecture elements to mitigate threats as they emerge, Create solutions that balance business requirements with information and cyber security requirements, Identify security design gaps in existing and proposed architectures and recommend changes or enhancements, Train users in implementation or conversion of systems. Cybersecurity Architect will industry standard certifications such as CISSP, CRISC, TOGAF, ISSAP, ISSEP, including other technical certification and will have Fifteen (15) or more years of experience in Cybersecurity including, Security architecture for cloud and on-premise environments, demonstrating solutions delivery, principles and emerging technologies - Designing and implementing security solutions. This includes continuous monitoring and making improvements to those solutions, working with an information security team, Consulting and engineering in the development and design of security leading practices and implementation of solid security principles across the organization, to meet business goals along with customer and regulatory requirements, Security considerations of cloud computing: They include data breaches, broken authentication, hacking, account hijacking, malicious insiders, third parties, APTs, data loss and DoS attacks, Identity and access management (IAM) – the framework of security policies and technologies that limit and track the access of those in an organization to sensitive technology resources, Relevant National Institute of Standards and Technology (NIST) standards. A system that is not in compliance with the standards set by NIST, along with ISO27001, COBIT and COSO (below), will lack both compliance and adequate security architecture, ISO27001 – specifications for a framework of policies and procedures that include all legal, physical and technical controls involved in an organization's risk management, Control Objectives for Information and Related Technologies (COBIT), Committee of Sponsoring Organizations (COSO) of the Treadway Commission, a joint initiative to combat corporate fraud, Windows, UNIX, Mainframe,

Role Name	Role Description
	PoS, IOT and ICT systems. Additionally, Cybersecurity Architect will have, Exceptional communication skills with diverse audiences - Strong critical thinking and analytical skills, Strong leadership, project and team-building skills, including the ability to lead teams and drive projects and initiatives in multiple departments, Demonstrated ability to identify risks associated with business processes, operations, information security programs and technology projects, The ability to be the enterprise security subject matter expert who can explain technical topics to those without a technical background.
Senior Cybersecurity Consultant	The Senior Security Consultant will be responsible for developing, architecting and implementing security solutions at an enterprise level. This includes establishing a formal information security framework, strategic roadmaps, as well as managing/overseeing projects to implement solutions to support the corporate security and technology strategy and confirms solutions work with a defense in depth program, as well as application and technology portfolio.
Cybersecurity Consultant	The Security Consultant will be responsible for implementing security architecture designs for security solutions at an enterprise level. This includes implementing a security framework, standards, managing day to day project activities to support the corporate security and technology strategy, confirms that solutions work with a defense in depth program, as well as application and technology portfolio.
Cybersecurity Analyst	The Cybersecurity Analyst collects and analyses intelligence regarding cyber threats and vulnerabilities, and direct and coordinate the response to such threats and vulnerabilities. Analyses of threats and vulnerabilities to determine their impact upon the IT systems. Identify the necessary actions required to proactively mitigate the risk posed by the threats and vulnerabilities. The Cybersecurity Analyst works in conjunction with the Cybersecurity Architect.
iPaaS Architect	The iPaaS Architect is a Technical Architect with industry experience of API Management, Integration and Microservices and specific experience of the architecture and operation of the Capgemini Enterprise iPaaS. This Architect will as the link between our Client, Project Teams, and other technical resources. The role is to technically oversee the assignment to deliver the desired outcome.
iPaaS Delivery Manager	The iPaaS Delivery Manager is a Delivery Manager with industry knowledge of DevOps and Agile methodologies and a track record of successful project delivery. This resource will have specific knowledge of the Capgemini Enterprise iPaaS and knowledge of our leading practices and ways of working. The iPaaS Delivery Manager is responsible for designing and implementing enterprise cloud iPaaS platform solutions required for cloud computing integration. Analyses business and system requirements confirms that systems will be securely integrated with current applications and platforms. Has a deep understanding of system development in cloud environments, including Software as Service (SaaS), Platform as Service (PaaS), or Infrastructure as a Service (IaaS).
iPaaS Platform Engineer	The iPaaS Platform Engineer is a specialist technical resource with industry knowledge of DevOps ways of working and Agile methodologies. This resource has specific knowledge of the Capgemini Enterprise iPaaS and the code, scripts, and tools utilised to deliver the platform.
iPaaS Integration Developer	The iPaaS Integration Developer is a specialist technical resource with industry knowledge of Integration development, DevOps ways of working and Agile methodologies. This resource has specific knowledge of the Capgemini Enterprise iPaaS and the leading practices, integration patterns and tools utilized to deliver APIs and Integrations.
Training	The Training Coordinator responsibilities include communicating with managers to

Role Name	Role Description
Coordinator	identify training needs and mapping out development plans for teams and individuals. Training Coordinators are responsible for managing, designing, developing, coordinating and conducting all training programs.
Training Analyst	The Training Analyst serves as a specialist in the planning and execution of instructional and/or research assignments. This position conducts needs assessments and development of measurement instruments for instructional assignments.
Training Lead	The Training Lead is responsible for improving the productivity of the organization's employees. This position assesses property-wide developmental needs to drive training initiatives and identifies and arranges suitable training solutions for employees
Help Desk 1st Line IT Service Desk	The Help Desk 1st Line IT Service Desk's primary responsibility is the management of the Problem Management process in the scope of the Service Support ITIL processes. The goal of the Help Desk 1st Line IT Service Desk is to prevent, remove or minimize the adverse impact of Incidents and Problems on the business that is caused by errors within the IT Infrastructure. Seek and eliminate recurrence of incidents, proactively identify hidden risk and improvement opportunities. In order to achieve this goal, Help Desk 1st Line IT Service Desk seeks to get to the root cause of Incidents and then initiate actions to improve or correct the situation. Undertakes pro-active investigations and preventive activities mitigation even potential IT business related risk. Carries out root cause analysis and preventative management as required, so that all necessary parties are informed and involved in the process. To build and maintain effective working relationships across tiers (Capgemini, Client, 3rd Party and Partner contacts) so that business and technical knowledge are applied effectively, in order to achieve the best possible levels of service quality and availability.
Help Desk Resolver	The Help Desk Resolver's primary responsibility is the management of the Problem Management process in the scope of the Service Support ITIL processes. The goal of Help Desk Resolver is to prevent, remove or minimize the adverse impact of Incidents and Problems on the business that is caused by errors within the IT Infrastructure. Seek and eliminate recurrence of incidents, proactively identify hidden risk and improvement opportunities. In order to achieve this goal, Help Desk Resolver seeks to get to the root cause of Incidents and then initiate actions to improve or correct the situation. Undertakes pro-active investigations and preventive activities mitigation even potential IT business related risk. Carries out root cause analysis and preventative management as required, making sure that that all necessary parties are informed and involved in the process. To build and maintain effective working relationships across tiers (Capgemini, Client, 3rd Party and Partner contacts) so that business and technical knowledge are applied effectively, in order to achieve the best possible levels of service quality and availability
Help Desk Reporting Analyst	The Help Desk Reporting Analyst operates and provides reporting platform provisioning pre-defined and standard reports set for defined stakeholders. Understands and is able to interpret reporting requirements in the context of the available technical solutions; Contributes towards the production of high-level and detailed design documents; Executes and deploys defined solution following standards and instructions; Provides that the reporting output produced is of high quality.
Help Desk Knowledge Manager	The role of the Help Desk Knowledge Manager is to undertake the management of Major Incidents, making sure that they are managed and communicated within the scope of Service Level Agreement, securing that all engaged parties perform up to high standards and follow agrees on the process. The role confirms that the output from Major Incident Management process is of high quality and provides management with the right level of business intelligence. The primary goal of Help Desk Knowledge Manager is, as a matter of urgency, to minimize or remove the adverse impact of

Role Name	Role Description
	Incidents on the business that is caused by errors within the IT Infrastructure, process failures, engaged staff competencies. The ultimate objective is to restore IT services as soon as possible. The role of Help Desk Knowledge Manager is embedded into Service Management Team and therefore works closely with Incident Management, Problem Management, Change Management, Service Desk Capgemini Service Management Leads, Service Delivery Managers and at times directly with the Customer. The role closely aligns with Capgemini's ITIL leading practices.
Help Desk/OCM Manager	The Help Desk/OCM Manager Undertakes management of changes, so that they are logged, progressed, updated/ authorized and actioned/resolved within the scope of the Service Level Agreement, so that all necessary parties are informed and involved in the process. Prepares reports on statistical information.

Amazon CloudFront Service Level Agreement

Last Updated June 1, 2013

This Amazon CloudFront Service Level Agreement ("SLA") is a policy governing the use of Amazon CloudFront under the terms of the Amazon Web Services Customer Agreement (the "AWS Agreement") between Amazon Web Services, Inc. and its affiliates ("AWS", "us" or "we") and users of AWS' services ("you"). This SLA applies separately to each account using Amazon CloudFront. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement

Service Commitment

AWS will use commercially reasonable efforts to make Amazon CloudFront available with a Monthly Uptime Percentage (defined below) of at least 99.9% during any monthly billing cycle (the "Service Commitment"). In the event Amazon CloudFront does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Error Rate" means: (i) the total number of internal server errors returned by Amazon CloudFront divided by (ii) the total number of requests during that five minute period. We will calculate the Error Rate for each Amazon CloudFront account as a percentage for each five minute period in the monthly billing cycle. The calculation of the number of internal server errors will not include errors that arise directly or indirectly as a result of any of the Amazon CloudFront SLA Exclusions (as defined below).

- "Monthly Uptime Percentage" is calculated by subtracting from 100% the average of the Error Rates from each five minute period in the monthly billing cycle.

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for Amazon CloudFront for the billing cycle in which the error occurred in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99% but less than 99.9%	10%
less than 99%	25%

We will apply any Service Credits only against future Amazon CloudFront payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the error occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar ($1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability, non-performance, or other failure by us to provide Amazon CloudFront is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you must submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each incident of non-zero Error Rates that you are claiming; and

iii. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the Monthly Uptime Percentage applicable to the month of such request is confirmed by us and is less than 99.9%, then we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon CloudFront SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon CloudFront, or any other Amazon CloudFront performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon CloudFront; (iii) that result from any actions or inactions of you or any third party; (iv) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); (v) arising from our suspension and termination of your right to use Amazon CloudFront in accordance with the AWS Agreement; (vi) that result from exceeding usage limits stated in the Amazon CloudFront documentation; or (vii) that result from use of an origin server other than Amazon S3 (collectively, the “Amazon CloudFront SLA Exclusions”). If availability is impacted by factors other than those used in our calculation of the Error Rate, then we may issue a Service Credit considering such factors at our discretion.

Amazon Compute Service Level Agreement

Last Updated February 12, 2018

This Amazon Compute Service Level Agreement (this "SLA") is a policy governing the use of the Included Products and Services (listed below) by you or the entity you represent ("you") under the terms of the AWS Customer Agreement (the "AWS Agreement") between Amazon Web Services, Inc. and its affiliates ("AWS", "us" or "we") and you. This SLA applies separately to each account using the Included Products and Services. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Included Products and Services

- Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Fargate for Amazon ECS (Amazon Fargate)

Service Commitment

AWS will use commercially reasonable efforts to make the Included Products and Services each available with a Monthly Uptime Percentage (defined below) of at least 99.99%, in each case during any monthly billing cycle (the "Service Commitment"). In the event any of the Included Products and Services do 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 any of the Included Products and Services, as applicable, was in the state of "Region Unavailable." Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon Compute Services SLA Exclusion (defined below).
- "Availability Zone" and "AZ" mean an isolated location within a region identified by a letter identifier following the region code (e.g., us-west-1a).

- “Region Unavailable” and “Region Unavailability” mean:
  - For Regions with only one AZ, when that AZ and one AZ in any other Region, in which you are running an instance or task (one or more containers), as applicable, are concurrently “Unavailable” to you.
  - For all other Regions, when more than one AZ within the same Region, in which you are running an instance or task (one or more containers), as applicable, are concurrently “Unavailable” to you.
- “Unavailable” and “Unavailability” mean:
  - For Amazon EC2, Amazon ECS, or Amazon Fargate, when all of your running instances or running tasks, as applicable, 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 99.99% 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.

Amazon DynamoDB Service Level Agreement

Last Updated June 19, 2018

This Amazon DynamoDB Service Level Agreement ("SLA") is a policy governing the use of Amazon DynamoDB ("DynamoDB") under the terms of the AWS Customer Agreement or other agreement with us governing your use of our Services (the "Agreement"). This SLA applies separately to each account using DynamoDB. Unless otherwise provided herein, this SLA is subject to the terms of the Agreement and capitalized terms will have the meaning specified in the Agreement. We reserve the right to change the terms of this SLA in accordance with the Agreement.

Service Commitment

AWS will use commercially reasonable efforts to make DynamoDB available with a Monthly Uptime Percentage (defined below) for each AWS region, during any monthly billing cycle, of (a) at least 99.999% if the Global Table SLA (defined below) applies, or (b) at least 99.99% if the Standard SLA (defined below) applies (the "Service Commitment"). In the event DynamoDB does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.
- The "Global Table SLA" is the Service Commitment that applies if all of your DynamoDB tables in the applicable AWS region are part of global tables as described on the AWS Site ("Global Tables") throughout the applicable monthly billing cycle, and you make reasonable attempts to failover in the event of an availability issue in a single AWS region.
- The "Standard SLA" is the Service Commitment that applies if any of your DynamoDB tables in the applicable AWS region are not part of Global Tables, or if the Global Table SLA would otherwise apply but you do not make reasonable attempts to failover in the event of an availability issue in a single AWS region.
- "Monthly Uptime Percentage" for a given AWS region is calculated as the average of the Availability for all 5-minute intervals in a monthly billing cycle. Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any DynamoDB SLA Exclusion (defined below).

- "Availability" is calculated for each 5-minute interval as the percentage of Requests (defined below) processed by DynamoDB that do not fail with Errors (defined below). If you did not make any Requests in a given 5-minute interval, that interval is assumed to be 100% available.
  - The Global Table SLA Availability calculation considers all Requests for all of your DynamoDB tables in the AWS region and also their corresponding replica tables in other AWS regions.
  - The Standard SLA Availability calculation considers all Requests for all of your DynamoDB tables in the applicable AWS region.
- A "Request" is a customer-initiated action of a type specifically listed as being supported by DynamoDB in the DynamoDB API Reference Documentation on the AWS Site. For the avoidance of doubt, Requests do not include actions listed under other products or services (e.g., Amazon DynamoDB Accelerator, Amazon DynamoDB Streams).
- An "Error" is any Request that returns a 500 or 503 error code, as described in DynamoDB Common Errors on the AWS Site.

Service Credits

Service Credits are calculated as a percentage of the following charges paid by you for DynamoDB for the monthly billing cycle in which the Monthly Uptime Percentage for a given AWS region fell within the ranges set forth in the table below: (a) if the Global Table SLA applies, the total charges paid by you for DynamoDB for the AWS region plus the charges described in the "Global Tables" section of the Amazon DynamoDB Pricing page of the AWS Site paid by you for corresponding replica tables in other AWS regions, or (b) if the Standard SLA applies, the total charges paid by you for DynamoDB in the applicable AWS region.

	Monthly Uptime Percentage	Service Credit Percentage
Global Table SLA	Less than 99.999% but equal to or greater than 99%	10%
	Less than 99%	25%
Standard SLA	Less than 99.99% but equal to or greater than 99%	10%

Less than 99%	25%

We will apply any Service Credits only against future DynamoDB payments otherwise due from you. At our discretion, we may issue the Service Credits 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. Service Credits 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 Agreement, your sole and exclusive remedy for any unavailability or non-performance or other failure by us to provide DynamoDB is the receipt of Service Credits (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive Service Credits, you will need to submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the billing cycle and AWS regions with respect to which you are claiming Service Credits, together with the Monthly Uptime Percentage for that AWS region for the billing cycle and the specific dates, times, and Availabilities for each 5-minute interval with less than 100% Availability in that AWS region throughout the billing cycle;

iii. your Request logs that document the errors for 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 credit request is confirmed by us and is less than the Service Commitment, then we will issue the Service Credits to you within one billing cycle following the month in which the credit request occurred. Your failure to provide the credit request and other information as required above will disqualify you from receiving Service Credits.

DynamoDB SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of DynamoDB, or any other DynamoDB performance issues:

(i) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of DynamoDB;

(ii) that result from any voluntary actions or inactions from you or any third party (e.g. scaling of provisioned capacity, misconfiguring security groups, VPC configurations or credential settings, disabling encryption keys or making the encryption keys inaccessible, etc.);

(iii) that result from you not following the best practices described in the DynamoDB User Guide on the AWS Site;

(iv) that result in additional recovery time due to insufficient read capacity units (RCUs) and write capacity units (WCUs) provisioned for your database workload;

(v) 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);

(vi) that result from any maintenance as provided for pursuant to the Agreement; or

(vii) arising from our suspension and termination of your right to use DynamoDB in accordance with the Agreement (collectively, the "DynamoDB SLA Exclusions").

If availability is impacted by factors other than those explicitly used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

Amazon RDS Service Level Agreement

Create Free Account »

Last Updated March 25, 2016

This Amazon RDS Service Level Agreement ("SLA") is a policy governing the use of the Amazon Relational Database Service ("Amazon RDS") under the terms of the AWS Customer Agreement (the "AWS Agreement") between Amazon Web Services, Inc. and its affiliates ("AWS", "us" or "we") and users of AWS' services ("you"). This SLA applies separately to each account using Amazon RDS. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Service Commitment

AWS will use commercially reasonable efforts to make Multi-AZ instances available with a Monthly Uptime Percentage (defined below) of at least 99.95% during any monthly billing cycle (the "Service Commitment"). In the event Amazon RDS does not meet the Monthly Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Monthly Uptime Percentage" for a given Multi-AZ instance is calculated by subtracting from 100% the percentage of 1 minute periods during the monthly billing cycle in which the Multi-AZ instance was "Unavailable". If you have been running that Multi-AZ instance for only part of the month, your Multi-AZ instance is assumed to be 100% available for the portion of the month that it was not running. Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon RDS SLA Exclusion (defined below).
- "Multi-AZ instance" means an Amazon RDS for MySQL, MariaDB, Oracle or PostgreSQL database instance with the Multi-AZ parameter set to true.
- "Unavailable" means that all connection requests to the running Multi-AZ instance fail during a 1 minute period.
- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible account.

Service Credits

Service Credits are calculated as a percentage of the charges paid by you for the Multi-AZ instances that did not meet the Monthly Uptime Percentage commitment in a billing cycle in accordance with the schedule below.

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0%	25%

We will apply any Service Credits only against future Amazon RDS payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the unavailability occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar ($1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability or non-performance or other failure by us to provide Amazon RDS is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you will need to submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

i. the words "SLA Credit Request" in the subject line;

ii. the dates and times of each Unavailability incident you are claiming;

iii. the DB Instance IDs and the AWS Regions of the affected Multi-AZ instances; and

iv. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the Monthly Uptime Percentage of such request is confirmed by us and is less than the Service Commitment, then we will issue the Service Credit to you within one billing cycle following the month in which the request occurred. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon RDS SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon RDS, or any other Amazon RDS performance issues:

(i) that result from a suspension described in Section 6.1 of the AWS Agreement;

(ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon RDS;

(iii) that result from any voluntary actions or inactions from you or any third party (e.g., rebooting a database instance, scaling compute capacity, not scaling storage when the storage is full, misconfiguring security groups, VPC configurations or credential settings, disabling encryption keys or making the encryption keys inaccessible, etc.);

(iv) that result from instances belonging to the Micro DB instance class or other instance classes which have similar CPU and memory resource limitations;

(v) that result from you not following the basic operational guidelines described in the Amazon RDS User Guide (e.g., overloading a database instance to the point it is inoperable, creating excessively large number of tables that significantly increase the recovery time etc.);

(vi) caused by underlying database engine software that lead to repeated database crashes or an inoperable database instance;

(vii) that result in long recovery time due to insufficient IO capacity for your database workload;

(viii) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); or

(ix) that result from any maintenance as provided for pursuant to the AWS Agreement; or

(x) arising from our suspension and termination of your right to use Amazon RDS in accordance with the AWS Agreement (collectively, the "Amazon RDS SLA Exclusions").

If availability is impacted by factors other than those explicitly used in our Monthly Uptime Percentage calculation, then we may issue a Service Credit considering such factors at our discretion.

Amazon Route 53 Service Level Agreement

Create Free Account »

Last Updated December 1, 2016

This Amazon Route 53 Service Level Agreement (“SLA”) is a policy governing the use of Amazon Route 53 (including Private DNS) under the terms of the Amazon Web Services Customer Agreement (the “AWS Agreement”) between Amazon Web Services, Inc. and its affiliates (“AWS”, “us” or “we”) and users of AWS’ services (“you”). This SLA applies separately to each account using Amazon Route 53. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Service Commitment

AWS will use commercially reasonable efforts to make Amazon Route 53 100% Available (defined below). In the event Amazon Route 53 does not meet the foregoing commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- “100% Available” means that Amazon Route 53 did not fail to respond to your DNS queries during a monthly billing cycle.
- A “Service Credit” is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon Route 53 account.

Service Credits

Service Credits are calculated based on 1 day of Service Credit, which is equal to your average daily Amazon Route 53 query charges for the monthly billing cycle preceding the monthly billing cycle in which the period that Amazon Route 53 was not 100% Available occurred, and are available as follows:

Duration Amazon Route 53 was not 100% Available	Service Credit
5 - 30 minutes	1 day Service Credit

31 minutes - 4 hours	7 days Service Credit
More than 4 hours	30 days Service Credit

We will apply any Service Credits only against future Amazon Route 53 payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the error occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar ($1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability, non-performance, or other failure by us to provide Amazon Route 53 is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you must submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

i. the words “SLA Credit Request” in the subject line;

ii. the dates and times of each period that Amazon Route 53 was not 100% Available that you are claiming; and

iii. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the period that Amazon Route 53 was not 100% Available is confirmed by us, then we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon Route 53 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon Route 53, or any other Amazon Route 53 performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon Route 53; (iii) that result from any actions or inactions of you or any

third party; (iv) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); (v) arising from our suspension and termination of your right to use Amazon Route 53 in accordance with the AWS Agreement; (vi) that result from you exceeding usage limits stated in the Amazon Route 53 documentation; or (vii) that, with respect to public DNS only, result during a period that you were not using all four virtual name servers (for example, ns123.awsdns.com, ns123.awsdns.net, ns123.awsdns.co.uk and ns123.awsdns.org) assigned to your “hosted zone” (collectively, the “Amazon Route 53 SLA Exclusions”). If availability is impacted by factors other than those used in our calculation of 100% Available, then we may issue a Service Credit considering such factors at our discretion.

Amazon S3 Service Level Agreement

Create Free Account »

Last Updated April 4, 2018

This Amazon S3 Service Level Agreement ("SLA") is a policy governing the use of Amazon Simple Storage Service ("Amazon S3") under the terms of the Amazon Web Services Customer Agreement (the "AWS Agreement") between Amazon Web Services, Inc. and its affiliates ("AWS", "us" or "we") and users of AWS' services ("you"). This SLA applies separately to each account using Amazon S3. Unless otherwise provided herein, this SLA is subject to the terms of the AWS Agreement and capitalized terms will have the meaning specified in the AWS Agreement. We reserve the right to change the terms of this SLA in accordance with the AWS Agreement.

Service Commitment

AWS will use commercially reasonable efforts to make Amazon S3 available with the applicable Monthly Uptime Percentage (as defined below) during any monthly billing cycle (the "Service Commitment"). In the event Amazon S3 does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Definitions

- "Error Rate" means: (i) the total number of internal server errors returned by Amazon S3 as error status "InternalError" or "ServiceUnavailable" divided by (ii) the total number of requests for the applicable request type during that five minute period. We will calculate the Error Rate for each Amazon S3 account as a percentage for each five minute period in the monthly billing cycle. The calculation of the number of internal server errors will not include errors that arise directly or indirectly as a result of any of the Amazon S3 SLA Exclusions (as defined below).
- "Monthly Uptime Percentage" is calculated by subtracting from 100% the average of the Error Rates from each five minute period in the monthly billing cycle.
- A "Service Credit" is a dollar credit, calculated as set forth below, that we may credit back to an eligible Amazon S3 account.

Service Credits

Service Credits are calculated as a percentage of the total charges paid by you for Amazon S3 for the billing cycle in which the error occurred in accordance with the schedule below.

For all requests not otherwise specified below:

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99.0% but less than 99.9%	10%
Less than 99.0%	25%

For requests to Amazon S3 Standard – Infrequent Access (Standard-IA) and Amazon S3 One Zone – Infrequent Access (OneZone-IA):

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 98.0% but less than 99.0%	10%
Less than 98.0%	25%

We will apply any Service Credits only against future Amazon S3 payments otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the error occurred. Service Credits will not entitle you to any refund or other payment from AWS. A Service Credit will be applicable and issued only if the credit amount for the applicable monthly billing cycle is greater than one dollar ($1 USD). Service Credits may not be transferred or applied to any other account. Unless otherwise provided in the AWS Agreement, your sole and exclusive remedy for any unavailability, non-performance, or other failure by us to provide Amazon S3 is the receipt of a Service Credit (if eligible) in accordance with the terms of this SLA.

Credit Request and Payment Procedures

To receive a Service Credit, you must submit a claim by opening a case in the AWS Support Center. To be eligible, the credit request must be received by us by the end of the second billing cycle after which the incident occurred and must include:

1. the words “SLA Credit Request” in the subject line;
2. the dates and times of each incident of non-zero Error Rates that you are claiming; and
3. your request logs that document the errors and corroborate your claimed outage (any confidential or sensitive information in these logs should be removed or replaced with asterisks).

If the Monthly Uptime Percentage applicable to the month of such request is confirmed by us and is less than the applicable Service Commitment, then we will issue the Service Credit to you within one billing cycle following the month in which your request is confirmed by us. Your failure to provide the request and other information as required above will disqualify you from receiving a Service Credit.

Amazon S3 SLA Exclusions

The Service Commitment does not apply to any unavailability, suspension or termination of Amazon S3, or any other Amazon S3 performance issues: (i) that result from a suspension described in Section 6.1 of the AWS Agreement; (ii) caused by factors outside of our reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon S3; (iii) that result from any actions or inactions of you or any third party; (iv) that result from your equipment, software or other technology and/or third party equipment, software or other technology (other than third party equipment within our direct control); or (v) arising from our suspension and termination of your right to use Amazon S3 in accordance with the AWS Agreement (collectively, the “Amazon S3 SLA Exclusions”). If availability is impacted by factors other than those used in our calculation of the Error Rate, then we may issue a Service Credit considering such factors at our discretion.

Volume Licensing

Service Level Agreement for Microsoft Online Services June 1, 2018

Microsoft

Table of Contents

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
Azure Container Instances	
Azure Database for MySQL	
Azure Database for PostgreSQL	
Azure DDoS Protection	
Azure Maps API	

Table of Contents / Definitions

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. Unless as otherwise provided in a specific SLA, only one Service Credit is permitted per Service for an Applicable Monthly Period.

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. That results from failures in a single Microsoft Datacenter location, when your network connectivity is explicitly dependent on that location in a non-geo-resilient manner;
4. 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;
5. 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;
6. 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;
7. 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;
8. That result from faulty input, instructions, or arguments (for example, requests to access files that do not exist);
9. That result from your attempts to perform operations that exceed prescribed quotas or that resulted from our throttling of suspected abusive behavior;
10. Due to your use of Service features that are outside of associated Support Windows; or
11. 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

Service Specific Terms

Microsoft Dynamics 365

Microsoft Dynamics 365 for Customer Service

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 Dynamics 365 Business Central

Downtime: Any period of time when end users are unable to login to their instance.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 Dynamics 365 for Finance and Operations (Enterprise edition)

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 Dynamics 365 for Retail

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 access 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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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%

Monthly Uptime Percentage	Service Credit
< 95%	100%

Table of Contents / Definitions

Microsoft Dynamics 365 for Sales Enterprise; Microsoft Dynamics 365 for Sales Professional

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:

$$\frac{User\ Minutes\ -\ Downtime}{User\ Minutes}\ x\ 100$$

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 Dynamics 365 for Talent; Microsoft Dynamics 365 for Talent: Attract; Microsoft Dynamics 365 for Talent: Onboard

Additional Definitions:
"**Active Tenant**" means a tenant with an active high availability production topology in the Management Portal that has an active database that users can log into.

Downtime: Any period of time when end users are unable to read or write any Service data for which they have appropriate permission. Downtime does not include Scheduled Downtime.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -\ Downtime}{User\ Minutes}\ x\ 100$$

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
< 99.5%	25%
< 99%	50%
< 95%	100%

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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. There is no Scheduled Downtime for this service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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

Exchange Online Archiving

Downtime: Any period of time when users are unable to access the email messages stored in their archive. There is no Scheduled Downtime for this service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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. There is no Scheduled Downtime for this service.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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

Microsoft Teams

Downtime: Any period of time when end users are unable to read or post to chat conversations for which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 MyAnalytics

Downtime: Any period of time when users are unable to access the MyAnalytics dashboard.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 Advanced Compliance

Downtime: Any period of time when Customer Lockbox component of Office 365 Advanced Compliance 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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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%

Monthly Uptime Percentage	Service Credit
< 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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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%

[1]Online meeting functionality applicable only to Skype for Business Online Plan 2 Service.

Table of Contents / Definitions

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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

This SLA applies to any eligible call placed by any voice service user within the subscription (enabled for making any type of call VOIP or PSTN).

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 based on numerous factors that could impact call quality in the networks managed by Microsoft. While the current Poor Call classifier is built primarily on network parameters like RTT (Roundtrip Time), Packet

Loss Rate, Jitter and Packet Loss-Delay Concealment Factors, it is dynamic and continually updated based on new learnings from analysis using millions of Skype and Skype for Business calls and evolution of Devices, Algorithms and end user ratings.

Monthly Good Call Rate: The Monthly Good Call Rate is calculated using the following formula:

$$\frac{Total\ Calls\ -Poor\ Quality\ Calls}{Total\ Calls}\ x\ 100$$

Service Credit:

Monthly Good Call Rate	Service Credit
< 99.9%	25%
< 99%	50%
< 95%	100%

Table of Contents / Definitions

Workplace Analytics

Downtime: Any period of time when users are unable to access the Workplace Analytics website.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 Azure Services

AD Domain Services

Additional Definitions:
"**Managed Domain**" refers to an Active Directory domain that is provisioned and managed by Azure Active Directory Domain Services.
"**Maximum Available Minutes**" is the total number of minutes that a given Managed Domain has been deployed by Customer in Microsoft Azure during a billing month in a given Microsoft Azure subscription.
"**Downtime**" is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Managed Domain is unavailable. A minute is considered unavailable if all requests for domain authentication of user accounts belonging to the Managed Domain, LDAP bind to the root DSE, or DNS lookup of records, made from within the virtual network where the Managed Domain is enabled, either return an Error Code or fail to return a Success Code within 30 seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Active Directory Domain Services:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Analysis Services

Additional Definitions:
"**Server**" means any Azure Analysis Services server.
"**Maximum Available Minutes**" is the total number of minutes that a given Server has been deployed in Microsoft Azure during a billing month in a given Microsoft Azure subscription.
"**Client Operations**" is the set of all documented operations supported by Azure Analysis Services.

Downtime: is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Server is unavailable. A minute is considered unavailable for a given Server if more than 1% of all Client Operations completed during the minute return an Error Code.

Monthly Uptime Percentage: The Monthly Uptime Percentage for a given Server is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

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.

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit for Basic Tier, Standard Tier and Premium Tier deployments scaled within a single region:

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 an API App, Logic App, 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 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Application Insights

Additional Definitions:
"**Application Insights Resource**" is the container in Application Insights that collects, processes and stores the data for a single instrumentation key.
"**Maximum Available Minutes**" is the total number of minutes that Application Insights Resource(s) have been deployed within a Microsoft Azure subscription during a billing month.
"**Data Latency**" is the number of minutes that data received from the instrumentation in Customer's application is delayed from appearing in Application Insights service where the delay is greater than 2 hours.
"**Downtime**" is the total accumulated number of minutes that are part of Maximum Available Minutes that experience Data Latency.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Automation Service – Desired State Configuration (DSC)

Additional Definitions:
"**Deployment Minutes**" is the total number of minutes that a given Automation account has been deployed in Microsoft Azure during a billing month.
"**DSC Agent Service**" is the component of the Automation Service responsible for receiving and responding to pull, registration, and reporting requests from DSC nodes.
"**Maximum Available Minutes**" is the sum of all Deployment Minutes across all Automation accounts deployed in a given Microsoft Azure subscription during a billing month

Downtime: The total accumulated Deployment Minutes, across all Automation accounts deployed in a given Microsoft Azure subscription, during which the DSC Agent Service is unavailable. A minute is considered unavailable for a given Automation account if all continuous pull, registration,

and reporting requests from DSC nodes associated with the Automation account to the DSC Agent Service throughout the minute either result in 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Automation Service – Process Automation

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:

$$\frac{Total\ Jobs - Delayed\ Jobs}{Total\ Jobs}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Advanced Threat Protection

Additional Definitions:
"**Downtime**" is Any period of time when the admin is unable to access the Azure ATP portal.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ - Downtime}{User\ Minutes}\ x\ 100$$

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 Bot Service

Additional Definitions:
"**Azure Bot Service Premium Channel**" is a Bot Framework channel in the premium category.
"**Bot**" is the developer's Internet facing conversational application which is registered with and is configured to send and receive messages from the Azure Bot Service.
"**Bot Framework**" is a platform for building, connecting, testing, and deploying powerful and intelligent bots.
"**Client**" is the end user facing portion of a Bot.
"**Premium Channels API Endpoint**" is a Bot Framework REST API endpoint for Azure Bot Service Premium Channels
"**Total API Requests**" is the total number of requests made by the Bot or the Client to the Premium Channel's API Endpoint in a Microsoft Azure subscription during a billing month.
"**Failed API Requests**" are the total number of requests within Total API Requests that return an Error Code or do not respond within 2 minutes.
"**Monthly Uptime Percentage**" is calculated as Total API Requests less Failed API Requests divided by Total API Requests multiplied by 100.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ API\ Requests\text{-}Failed\ API\ Requests}{Total\ API\ Requests}\ x\ 100$$

The following Service Levels and Service Credits are applicable to Customer's use of the Azure Bot Service Premium Channels.

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Container Instances

Additional Definitions:
"**Connectivity**" is bi-directional network traffic between the Container Group and other IP addresses using TCP or UDP network protocols in which the Container Group is configured for allowed traffic.
"**Container Group**" is a collection of co-located containers that shares the same lifecycle and networking resources.
"**Maximum Available Minutes**" is the total number of minutes that a given Container Group has been deployed by Customer in a Microsoft Azure subscription during a billing month. Maximum Available Minutes is measured from Customer action that results in starting a given Container Group to the time Customer action that results in stopping or deleting a given Container Group.
"**Downtime**" is the total number of minutes within Maximum Available Minutes that have no Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Container Instances:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Cosmos DB

Additional Definitions:
"**Collection**" is a container of JSON documents, and a unit of scale for transactions and queries.
"**Consumed RUs**" is the sum of the Request Units consumed by all the requests which are processed by the Azure Cosmos DB Collection in a given second.

"**Database Account**" is the top-level resource of the Azure Cosmos DB resource model. A Azure Cosmos DB Database Account contains one or more databases.

"**Failed Requests**" are requests within Total Requests that either return an Error Code or fail to return a Success Code within the maximum upper bounds documented in the table below.

"**Failed Read Requests**" are requests within Total Read Requests that either return an Error Code or fail to return a Success Code within the maximum upper bounds documented in the table below.

Operation	Maximum Upper Bound on Processing Latency
All Database Account configuration operations	2 Minutes
Add a new Region	60 Minutes
Manual Failover	5 Minutes
Resource Operations	5 Seconds
Media Operations	60 Seconds

"**Provisioned RUs**" is the total provisioned Request Units for a given Azure Cosmos DB Collection for a given second.
"**Rate Limited Requests**" are requests which are throttled by the Azure Cosmos DB Collection after Consumed RUs have exceeded the Provisioned RUs for a partition in the Collection for a given second.
"**Request Unit (RU)**" is a measure of throughput in Azure Cosmos DB.
"**Resource**" is a set of URI addressable entities associated with a Database Account.

"**Successful Requests**" are Total Requests minus Failed Requests.

"**Total Read Requests**" is the set of all the read requests, including Rate Limited Requests and all the Failed Read Requests, issued against Resources within a one-hour interval within a given Azure subscription during a billing month.

"**Total Requests**" is the set of all requests, including Rate Limited Requests and all Failed Requests, issued against Resources within a one-hour interval within a given Azure subscription during a billing month.

Availability SLA

"**Read Error Rate**" is the total number of Failed Read Requests divided by Total Read Requests, across all Resources in a given Azure subscription, during a given one-hour interval. If the Total Read Requests in a given one-hour interval is zero, the Read Error Rate for that interval is 0%.
"**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%.
"**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.
"**Average Read Error Rate**" for a billing month is the sum of Read Error Rates for each hour in the billing month divided by the total number of hours in the billing month.

Monthly Availability Percentage: For the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. The Monthly Availability Percentage is represented by the following formula:

100% - Average Error Rate

Service Credit:

Monthly Availability Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Monthly Availability Percentage: For the Azure Cosmos DB Service with multiple regions is calculated by subtracting from 100% the Average Read Error Rate for a given Microsoft Azure subscription in a billing month. Monthly Read Availability Percentage is represented by the following formula:

100% - Average Read Error Rate

Service Credit:

Monthly Read Availability Percentage	Service Credit
< 99.999%	10%
< 99%	25%

Throughput SLA

"**Throughput Failed Requests**" are requests which are throttled by the Azure Cosmos DB Collection resulting in an Error Code, before Consumed RUs have exceeded the Provisioned RUs for a partition in the Collection for a given second.
"**Error Rate**" is the total number of Throughput 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%.
"**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.

"**Monthly Throughput Percentage**" for the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. Monthly Throughput Percentage is represented by the following formula:

100% - Average Error Rate

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Consistency SLA

"**K**" is the number of versions of a given document for which the reads lag behind the writes.
"**T**" is a given time interval.
"**Consistency Level**" is the setting for a particular read request that supports consistency guarantees. The following table captures the guarantees associated with the Consistency Levels.

Consistency Level	Consistency Guarantees
Strong	Linearizability
Sessions	Read Your Own Write (within write region) Monotonic Read Consistent Prefix
Bounded Staleness	Read Your Own Write (within write region) Monotonic Read (within a region) Consistent Prefix Staleness Bound < K,T
Consistent Prefix	Consistent Prefix
Eventual	Eventual

"**Consistency Violation Rate**" is Successful Requests that could not be delivered when performing the consistency guarantees specified for the chosen Consistency Level 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 Consistency Violation Rate for that interval is 0%.
"**Average Consistency Violation Rate**" for a billing month is the sum of Consistency Violation Rates for each hour in the billing month divided by the total number of hours in the billing month.

"**Monthly Consistency Attainment Percentage**" for the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Consistency Violation Rate for a given Microsoft Azure subscription in a billing month.

Monthly Consistency Percentage: For the Azure Cosmos DB Service is calculated by subtracting from 100% the Average Consistency Violation Rate for a given Microsoft Azure subscription in a billing month. The Monthly Consistency Percentage is represented by the following formula:

100% - Average Consistency Violation Rate

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Latency SLA

"**Application**" is a Azure Cosmos DB application deployed within a local Azure region using the Azure Cosmos DB client SDK configured with TCP direct connectivity for a given Microsoft Azure subscription in a billing month.
"**N**" is the number of Successful Requests for a given Application performing either document read or document write operations with a payload size less than or equal to 1 KB in a given hour.

"**S**" is the latency-sorted set of Successful Request response times in ascending order for a given Application performing document read or document write operations with a payload size less than or equal to 1 KB in a given hour.
"**Ordinal Rank**" is the 99th percentile using the nearest rank method represented by the following formula:

$$Ordinal\ Rank = \frac{99}{100}\ x\ N$$

"**P99 Latency**" is the value at the Ordinal Rank of S.
"**Excessive Latency Hours**" is the total number of one-hour intervals during which Successful Requests submitted by an Application resulted in a P99 Latency greater than or equal to 10ms for document read or 15ms for document write operations. If the number of Successful Requests in a given one-hour interval is zero, the Excessive Latency Hours for that interval is 0.
"**Average Excessive Latency Rate**" for a billing month is the sum of Excessive Latency Hours divided by the total number of hours in the billing month.

"**Monthly P99 Latency Attainment Percentage**" for a given Azure Cosmos DB Application is calculated by subtracting from 100% the Average Excessive Latency Rate for a given Microsoft Azure subscription in a billing month. Monthly P99 Latency Attainment Percentage is represented by the following formula::

$$100\% - Average\ Excessive\ Latency\ Rate$$

Service Credit:

Monthly P99 Latency Attainment Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Table of Contents / Definitions

Azure Database for MySQL

Additional Definitions:
"**Server**" is any given Azure Database for MySQL server.
"**Maximum Available Minutes**" is the total number of minutes for a given Server deployed by Customer in a Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes during which a Server is unavailable. A minute is considered unavailable if all continuous attempts by Customer to establish a connection to the Server returned an Error Code.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Database for MySQL:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%

Table of Contents / Definitions

Azure Database for PostgreSQL

Additional Definitions:
"**Server**" is any given Azure Database for PostgreSQL server.
"**Maximum Available Minutes**" is the total number of minutes for a given Server deployed by Customer in a Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes during which a Server is unavailable. A minute is considered unavailable if all continuous attempts by Customer to establish a connection to the Server returned an Error Code.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Database for PostgreSQL:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%

Table of Contents / Definitions

Azure DDoS Protection

"**Maximum Available Minutes**" is the total number of minutes DDoS Protection Service is enabled for a given Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes where protected Azure resources were not available. A minute is considered unavailable when DDoS Protection did not mitigate an attack which directly resulted in underlying Azure resources not meeting respective SLA.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure DDoS Protection:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99.95%	25%

Table of Contents / Definitions

Azure Functions

For Function Apps running on App Service Plans we guarantee that the associated Functions compute will be available 99.95% of the time.
No SLA is provided for Functions Apps running under Consumption Plans.

Additional Definitions:
"**Deployment Minutes**" is the total number of minutes that a given Function App is available to be triggered during a billing month. Deployment Minutes are measured based on the total time that the service is available to trigger a function execution and not based on the potential number of Function executions that might be triggered during a given month.
"**Maximum Available Minutes**" is the sum of all Deployment Minutes across all Function Apps deployed by Customer in a given Microsoft Azure subscription during a billing month.
"**Function App**" is an individual Function deployed on an App Service Plan with an associated trigger.
"**Downtime**" The total accumulated Deployment Minutes, across the Function App deployed by a customer in a given Microsoft Azure subscription, during which the Function App is unavailable to be triggered. A minute is considered unavailable for a given Function App when there is no connectivity between the App Service Plan on which the Function App is hosted and Microsoft's Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Table of Contents / Definitions

Azure Load Balancer

Additional Definitions:
"**Load Balanced Endpoint**" is an IP address and associated IP transport port definition.

"**Healthy Virtual Machine**" is a Virtual Machine which returns a Success Code for the health probe sent by the Azure Standard Load Balancer. The Virtual Machine must have Network Security Group rules permitting communication with the load balanced port.
"**Connectivity**" is bi-directional network traffic over supported IP transport protocols that can be sent and received from any IP address configured to allow traffic.
"**Maximum Available Minutes**" is the total number of minutes that a given Azure Standard Load Balancer (serving two or more Healthy Virtual Machines) has been deployed by Customer in a Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes during which the given Azure Standard Load Balancer is unavailable. A minute is considered unavailable if all Healthy Virtual Machines have no Connectivity through the Load Balanced Endpoint. Downtime does not include minutes resulting from SNAT port exhaustions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Load Balancer:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99.9%	25%

Service Level Exceptions: The Basic Load Balancer is not covered by this SLA.

Table of Contents / Definitions

Azure Maps API

Additional Definitions:
"**Total Transaction Attempts**" is the total number of authenticated API requests made by Customer for a given Azure Map API during a billing month in a given Microsoft Azure subscription. "**Failed Transactions**" is the set of all requests within Total Transaction Attempts that result in an Error Code or otherwise do not return a Success Code within 60 seconds after receipt by the Service.
Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ Transaction\ Attempts\ -\ Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Maps API:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Monitor

Additional Definitions:
"**Action Group**" is a collection of actions deployed by Customer in a given Microsoft Azure subscription which defines preferred notification delivery methods.
"**Deployment Minutes**" is the total number of minutes that a given Action Group has been deployed by Customer in Microsoft Azure subscription during a billing month.
"**Maximum Available Minutes**" is the sum of all Deployment Minutes across all Action Groups deployed by Customer in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment Minutes, across all Action Groups, during which the Action Group is unavailable. A minute is considered unavailable for a given Action Group if all continuous attempts to send alerts or perform registration management operations with respect to the Action Group throughout the minute either return an Error Code or do not result in a Success Code within five minutes.

Monthly Uptime Percentage: is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes} \ x\ 100$$

Service Levels and Service Credits:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Monitor Alerts

Additional Definitions:
"**Alert Rule**" is a collection of signal criteria used to generate alerts using monitoring event data already available to Alert Service for analysis.
"**Maximum Available Minutes**" is the total number of minutes which Alert Rule(s) are deployed by Customer in a given Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes during which the Alert Rule is unavailable. A minute is considered unavailable for a given Alert Rule if all continuous attempts to analyze telemetry signals for resources defined within the Alert Rule throughout the minute either return an Error Code or do not result in a Success Code within five minutes from scheduled Alert Rule start time.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes} \ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Monitor Alerts:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Monitor Notification Delivery

Additional Definitions:
"**Action Group**" is a collection of actions that defines preferred notification delivery methods.
"**Maximum Available Minutes**" is the total number of minutes which Action Group(s) are deployed by Customer in a given Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes during which the Action Group is unavailable. A minute is considered unavailable for a given Action Group if all continuous attempts to send alerts or perform registration management operations with respect to the Action Group 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes} \ x\ 100$$

Service Levels and Service Credits are applicable to Customer's use of Azure Monitor Notification Delivery:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Security Center

Additional Definitions:
"**Protected Node**" is a Microsoft Azure resource, counted as a node for billing purposes that is configured for the Azure Security Center Standard Tier

"**Security Monitoring**" is the assessment of a Protected Node resulting in potential findings such as security health status, recommendations, and security alerts, exposed in Azure Security Center.

"**Maximum Available Minutes**" is the total number of minutes during a billing month that a given Protected Node has been deployed and configured for Security Monitoring.

"**Downtime**" is the total accumulated minutes during a billing month for which Security Monitoring information of a given Protected Node is unavailable. A minute is considered unavailable for a given Protected Node if all continuous attempts to retrieve Security Monitoring information throughout the minute result in either an Error Code or do not return a Success Code within two minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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 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: for the Batch Service is calculated by subtracting from 100% the Average Error Rate for a given Microsoft Azure subscription in a billing month. The "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. Monthly Uptime Percentage is represented by the following formula:

$$Monthly\ Uptime\ \% = 100\% - Average\ Error\ Rate$$

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.

"**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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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.

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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

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%

Monthly Uptime Percentage	Service Credit
< 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.
"**Role Instance Connectivity**" is bi-directional network traffic between the role instance and other IP addresses using TCP or UDP network protocols in which the role instance is configured for allowed traffic. The IP addresses can be IP addresses in the same Cloud Service as the virtual machine, IP addresses within the same virtual network as the virtual machine or public, routable IP addresses.
"**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 Customer to the time Customer has 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 Role Instance Connectivity.

Monthly Uptime Percentage: Monthly Uptime Percentage is represented by the following formula:

$$Monthly\ Uptime\ \% = \frac{(Maximum\ Available\ Minutes\text{-}Downtime)}{Maximum\ Available\ Minutes}\ X\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Table of Contents / Definitions

Container Registry

Additional Definitions:
"**Managed Registry**" is any instance of Basic, Standard or Premium Container Registry.
"**Registry Endpoint**" is the host name from which a given Managed Registry is accessed by clients to perform Container Registry related operations.
"**Registry Transactions**" is the set of transaction requests sent from the client to the Registry Endpoint.
"**Maximum Available Minutes**" is the total number of minutes that a given Managed Container Registry has been deployed by Customer in a Microsoft subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes during which Managed Registry is unavailable. A minute is considered unavailable if all continuous attempts to send Registry Transactions receive an Error Code or do not respond within the Maximum Processing Time outlined in the table below.

Transaction Types	Maximum Processing Time
List (Repository, Manifests, Tags)	8 Minutes
Others	1 Minute

"**Monthly Uptime Percentage**" for Managed Container Registry is calculated using the following formula:

$$Monthly\ Uptime\ \% = \frac{(Maximum\ Available\ Minutes - Downtime)}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Data Catalog

Additional Definitions:
"**Deployment Minutes**" is the total number of minutes for which a Data Catalog has been purchased during a billing month.

"**Entries**" means any catalog object registration in the Data Catalog (such as a table, view, measure, cluster or report).
"**Maximum Available Minutes**" is the sum of all Deployment Minutes for the Data Catalog associated with a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment minutes, during which the Data Catalog is unavailable. A minute is considered unavailable for a given Data Catalog if all attempts by administrators to add or remove users to the Data Catalog or all attempts by users to execute API calls to the Data Catalog for registering, searching, or deleting Entries either result in an Error Code or do not return a response within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

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:

$$\frac{Total\ Activity\ Runs - Delayed\ Activty\ Runs}{Total\ Activity\ Runs}\ x\ 100$$

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 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: of the API calls made to the Data Factory Services is calculated as Total Requests less Failed Requests divided by Total Requests in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$Monthly\ Uptime\ \% = \frac{(Total\ Requests - Failed\ Requests)}{Total\ Requests}$$

Service Credit:
The following Service Credits are applicable to Customer's use of API calls within the Data Factory Service

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Data Lake Analytics

Additional Definitions:
"**Total Operations**" is the total number of authenticated operations attempted within a one-hour interval across all Data Lake Analytics accounts in a given Azure subscription during a billing month.
"**Failed Operations**" is the set of all operations within Total Operations that either return an Error Code or fail to return a Success Code within 5 minutes for account creation and deletion and 25 seconds for all other operations with an additional 2 seconds per MB for operations with payload.
"**Error Rate**" is the total number of Failed Operations divided by Total Operations during a given one-hour interval. If the Total Operations in a one-hour interval is zero, the Error Rate for that interval is 0%.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - Average\ Error\ Rate$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Data Lake Store

Additional Definitions:
"**Total Operations**" is the total number of authenticated operations attempted within a one-hour interval across all Data Lake Store accounts in a given Azure subscription during a billing month.
"**Failed Operations**" is the set of all operations within Total Operations that either return an Error Code or fail to return a Success Code within 5 minutes for account creation and deletion, 2 seconds per file for operations on multiple files, 2 seconds per MB for data transfer operations, and 2 seconds for all other operations.
"**Error Rate**" is the total number of Failed Operations divided by Total Operations during a given one-hour interval. If the Total Operations in a one-hour interval is zero, the Error Rate for that interval is 0%.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$100\% - Average\ Error\ Rate$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Event Grid

Additional Definitions:
"**Maximum Available Minutes**" is the total number of minutes that an Event Grid has been deployed by Customer in a Microsoft Azure subscription during a billing month.

"**Downtime**" is the total number of minutes within Maximum Available Minutes across all Event Grids deployed by Customer in a given Microsoft Azure subscription during which Event Grid is unavailable. A minute is considered unavailable for a given Event Grid if all requests to publish a message either return an Error Code or do not result in a Success Code within one minute.

"**Monthly Uptime Percentage**": The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\ -Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
<99.99%	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 ExpressRoute connectivity 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**" is 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**" is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit The following Service Levels and Service Credits are applicable to Customer's use of each Dedicated Circuit within the ExpressRoute Service.

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

IoT hub

Additional Definitions:
"**Deployment Minutes**" is the total number of minutes that a given IoT hub has been deployed in Microsoft Azure during a billing month.
"**Device Identity Operations**" refers to create, read, update, and delete operations performed on the device identity registry of an IoT hub.
"**Maximum Available Minutes**" is the sum of all Deployment Minutes across all IoT hubs deployed in a given Microsoft Azure subscription during a billing month.
"**Message**" refers to any content sent by a deployed IoT hub to a device registered to the IoT hub or received by the IoT hub from a registered device, using any protocol supported by the Service.

Downtime: The total accumulated Deployment Minutes, across all IoT hubs deployed in a given Microsoft Azure subscription, during which the IoT hub is unavailable. A minute is considered unavailable for a given IoT hub if all continuous attempts to send or receive Messages or perform Device Identity Operations on the IoT 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Log Analytics

Additional Definitions:
"**Batch**" means a group of Log Data entries that are either uploaded to the Log Analytics Service or read from storage by the Log Analytics 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 Log Analytics 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 Log Analytics Service during a given billing month.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ Queued\ Batches - Delayed\ Batches}{Total\ Queued\ Batches}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Logic Apps

Additional Definitions:
"**Deployment Minutes**" is the total number of minutes that a given Logic App has been set to running in Microsoft Azure during a billing month. Deployment Minutes is measured from when the Logic App was created or Customer initiated an action that would result in running the Logic App to the time Customer initiated an action that would result in stopping or deleting the Logic App.
"**Maximum Available Minutes**" is the sum of all Deployment Minutes across all Logic Apps deployed by Customer in a given Microsoft Azure subscription during a billing month.

"**Downtime**" The total accumulated Deployment Minutes, across all Logic Apps deployed by Customer in a given Microsoft Azure subscription, during which the Logic App is unavailable. A minute is considered unavailable for a given Logic App when there is no connectivity between the Logic App and Microsoft's Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes - Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Azure Machine Learning Studio – 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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

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 Azure Machine Learning Studio BES and Management API Service. The Free Azure Machine Learning Studio tier is not covered by this SLA.

Table of Contents / Definitions

Azure Machine Learning Studio – 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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

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 Azure Machine Learning Studio RRS and Management API Service. The Free Azure Machine Learning Studio 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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Microsoft Cognitive Services

Additional Definitions:

"**Total Transaction Attempts**" is the total number of authenticated API requests by Customer during a billing month for a given Cognitive Service API. Total Transaction Attempts do not include API requests that return an Error Code that are continuously repeated within a five-minute window after the first Error Code is received.

"**Failed Transactions**" is the set of all requests to the Cognitive Service API within Total Transaction Attempts that return an Error Code . Failed Transaction Attempts do not include 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**" for each API Service is calculated as Total Transaction Attempts less Failed Transactions divided by Total Transaction Attempts in a billing month for a given API subscription. Monthly Uptime Percentage is represented by the following formula:
Monthly Uptime % = (Total Transaction Attempts - Failed Transactions) / Total Transaction Attempts * 100

$$Monthly\ Uptime\ \% = \frac{(Total\ Transaction\ Attempts - Failed\ Transactions)}{Total\ Transaction\ Attemtps}\ x\ 100$$

Service Credit

The following Service Levels and Service Credits are applicable to Cognitive Services APIs:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: No SLA is provided to free tier or offerings in preview.

Table of Contents / Definitions

Microsoft Genomics

Additional Definitions:

"**Maximum Available Minutes**" is the total accumulated minutes for all Microsoft Genomics accounts created by Customer and active during a billing month for a given Microsoft Azure Subscription.

"**Downtime**" is the total number of minutes within Maximum Available Minutes during which Microsoft Genomics is unavailable. A minute is considered unavailable if all continuous attempts to send authenticated Genomics service REST API requests throughout the minute either return an Error Code or do not respond with an acknowledgement within the minute.

"**Monthly Uptime Percentage**" for Microsoft Genomics is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Mobile Engagement

Additional Definitions:
"**Average Error Rate**" for a billing month is the sum of Error Rates for each hour 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:

$$100\% - Average\ Error\ Rate$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	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:

$$\frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}\ x\ 100$$

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

Network Watcher

Additional Definitions:
"**Network Diagnostic Tools**" is a collection of network diagnostic and topology tools.

"**Maximum Diagnostic Checks**" is the total number of diagnostic actions performed by the Network Diagnostic Tool as configured by Customer in a billing month for a given Microsoft Azure subscription.
"**Failed Diagnostic Checks**" is the total number of diagnostic actions within Maximum Diagnostic Checks that returns an Error Code or does not return a response within the Maximum Processing Time documented in the table below.

Diagnostic Tool	Maximum Processing Time
IPFlow Verify NextHop Packet Capture Security Group View Topology	2 minutes
VPN Troubleshoot	10 minutes

"**Monthly Uptime Percentage**" is calculated by using the following formula:

$$\frac{Maximum\ Diagnostic\ Checks - Failed\ Diagnostic\ Checks}{Maximum\ Diagnostic\ Checks}\ x\ 100$$

Service Levels:

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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 RemoteApp Service. The RemoteApp free trial is not covered by this SLA.

Table of Contents / Definitions

SAP HANA on Azure

Additional Definitions:
"**Announced Single Instance Maintenance**" means periods of Downtime related to network, hardware, or Service maintenance or upgrades impacting Single Instances. We will publish notice or notify you at least five (5) days prior to the commencement of such Downtime.

"**High Availability Pair**" refers to two or more identical SAP HANA on Azure large instances deployed in the same region and configured by the customer for system replication at the application layer. Customer must declare the members of a High Availability Pair to Microsoft during the architecture design process.
"**SAP HANA on Azure Connectivity**" is bi-directional network traffic between the SAP HANA on Azure large instance and other IP addresses using TCP or UDP network protocols in which the instance is configured for allowed traffic. The IP addresses must be IP addresses on the Virtual Network of the associated Azure subscription.
"**Single Instance**" is defined as any single Microsoft SAP HANA on Azure Large Instance machine that is not deployed in an High Availability Pair.

Monthly Uptime Calculation and Service Levels for SAP HANA on Azure High Availability Pair

"**Maximum Available Minutes**" is the total accumulated minutes during a billing month for all SAP HANA on Azure instances deployed in the same High Availability Pair. Maximum Available Minutes is measured from when two or more instances in the same High Availability Pair have both been started resultant from an action initiated by Customer to the time Customer has initiated an action that would result in stopping the instances.

"**Downtime**" is the total accumulated minutes that are part of Maximum Available Minutes that have no SAP HANA on Azure Connectivity.

Monthly Uptime Percentage: The Monthly Uptime Percentage for SAP HANA on Azure High Availability Pair is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit for SAP HANA on Azure High Availability Pair:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99.9%	25%

Monthly Uptime Calculation and Service Levels for SAP HANA on Azure Single Instance

"**Maximum Available Minutes**" is the total accumulated minutes for all SAP HANA on Azure Single Instances deployed by Customer during a billing month for a given Microsoft Azure subscription.

"**Downtime**" is the total accumulated minutes that are part of Maximum Available Minutes that have no SAP HANA on Azure Connectivity. Downtime excludes Announced Single Instance Maintenance.

Monthly Uptime Percentage: The Monthly Uptime Percentage for SAP HANA on Azure Single Instance is calculated using the following formula

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

The following Service Levels and Service Credits are applicable to Customer's use of SAP HANA on Azure Single Instances:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%
<95%	100%

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$100\% - Average\ Error\ Rate$$

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 / Definitions

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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.

Table of Contents / Definitions

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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.

Table of Contents / Definitions

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Service-Bus Service – Relays

Additional Definitions:
"**Message**" refers to any user-defined content sent or received through Service Bus Relays, Queues, or Topics, using any protocol supported by Service Bus.
"**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 Customer in a given Microsoft Azure subscription during a billing month.

Downtime: Is the total accumulated Deployment Minutes, across all Relays deployed by Customer 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 for Relays is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

SQL Data Warehouse Database

Additional Definitions:
"**Database**" means any SQL Data Warehouse Database.
"**Maximum Available Minutes**" is the total number of minutes that a given Database has been deployed in Microsoft Azure during a billing month in a given Microsoft Azure subscription.
"**Client Operations**" is the set of all documented operations supported by SQL Data Warehouse.

Downtime: is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Database is unavailable. A minute is considered unavailable for a given Database if more than 1% of all Client Operations completed during the minute return an Error Code.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%

Monthly Uptime Percentage	Service Credit
< 99%	25%

Table of Contents / Definitions

SQL Database Service (Basic, Standard and Premium Tiers)

Additional Definitions:
"**Database**" means any single or elastic Basic, Standard, or Premium Microsoft Azure SQL Database.
"**Maximum Available Minutes**" is the total number of minutes that a given Database has been deployed in in Microsoft Azure during a billing month in a given Microsoft Azure subscription.

Downtime: is the total accumulated minutes during a billing month for a given Microsoft Azure subscription during which a given Database is unavailable. A minute is considered unavailable for a given Database if all continuous attempts to establish a connection to the Database within the minute fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

SQL Server Stretch Database

Additional Definitions:
"**Database**" means one instance of SQL Server Stretch Database.
"**Maximum Available Minutes**" is the total number of minutes that a given Database has been deployed in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated minutes across all Databases deployed by Customer 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 Customer to establish a connection to the Database within the minute fail.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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.
"**Blob Storage Account**" is a storage account specialized for storing data as blobs and provides the ability to specify an access tier indicating how frequently the data in that account is accessed.
"**Cool Access Tier**" is an attribute of a Blob Storage Account indicating that the data in the account is infrequently accessed and has a lower availability service level than data in other access tiers.
"**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, file shares, tables, or queues; clearing of queues; and copying blobs or files 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
PutFile and GetFile	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)
CopyFile	Ninety (90) seconds (where the source and destination files are within the same storage account)
PutBlockList GetBlockList	Sixty (60) seconds
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, File, 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:

$$100\% - Average\ Error\ Rate$$

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%

Service Credit – LRS, GRS and RA-GRS (write requests) Blob Storage Accounts (Cool Access Tier):

Monthly Uptime Percentage	Service Credit
< 99%	10%
< 98%	25%

Service Credit – RA-GRS (read requests) Blob Storage Accounts (Cool Access Tier):

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 98%	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:

$$Monthly\ Uptime\ \% = \frac{Total\ Transaction\ Attempts - Failed\ Transactions}{Total\ Transaction\ Attempts}$$

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.

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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.
"**Availability Zone**" is a fault-isolated area within an Azure region, providing redundant power, cooling, and networking.
"**Data Disk**" is a persistent virtual hard disk, attached to a Virtual Machine, used to store application data.
"**Fault Domain**" is a collection of servers that share common resources such as power and network connectivity.
"**Operating System Disk**" is a persistent virtual hard disk, attached to a Virtual Machine, used to store the Virtual Machine's operating system.
"**Single Instance**" is defined as any single Microsoft Azure Virtual Machine that either is not deployed in an Availability Set or has only one instance deployed in an Availability Set.
"**Virtual Machine**" refers to persistent instance types that can be deployed individually or as part of an Availability Set.
"**Virtual Machine Connectivity**" is bi-directional network traffic between the Virtual Machine and other IP addresses using TCP or UDP network protocols in which the Virtual Machine is configured for allowed traffic. The IP addresses can be IP addresses in the same Cloud Service as the Virtual Machine, IP addresses within the same virtual network as the Virtual Machine or public, routable IP addresses.

Monthly Uptime Calculation and Service Levels for Virtual Machines in Availability Zones

"**Maximum Available Minutes**" is the total accumulated minutes during a billing month that have two or more instances deployed across two or more Availability Zones in the same region. Maximum Available Minutes is measured from when at least two Virtual Machines across two Availability Zones in the same region have both been started resultant from action initiated by Customer to the time Customer has initiated an action that would result in stopping or deleting the Virtual Machines.
"**Downtime**" is the total accumulated minutes that are part of Maximum Available Minutes that have no Virtual Machine Connectivity in the region.
"**Monthly Uptime Percentage**" for Virtual Machines in Availability Zones is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$Monthly\ Uptime\ \% = \frac{(Maximum\ Available\ Minutes\text{-}Downtime)}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:
The following Service Levels and Service Credits are applicable to Customer's use of Virtual Machines deployed across two or more Availability Zones in the same region:

Monthly Uptime Percentage	Service Credit
< 99.99%	10%
< 99%	25%
< 95%	100%

Monthly Uptime Calculation and Service Levels for Virtual Machines in an Availability Set

Maximum Available Minutes: 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.

Downtime: The total accumulated minutes that are part of Maximum Available Minutes that have no Virtual Machine Connectivity.

Monthly Uptime Percentage: for Virtual Machines is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes in a billing month for a given Microsoft Azure subscription. Monthly Uptime Percentage is represented by the following formula:

$$Monthly\ Uptime\ \% = \frac{(Maximum\ Available\ Minutes\text{-}Downtime)}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:
The following Service Levels and Service Credits are applicable to Customer's use of Virtual Machines in an Availability Set:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%

Monthly Uptime Percentage	Service Credit
< 99%	25%
< 95%	100%

Monthly Uptime Calculation and Service Levels for Single-Instance Virtual Machines

"**Minutes in the Month**" is the total number of minutes in a given month.

Downtime: is the total accumulated minutes that are part of Minutes in the Month that have no Virtual Machine Connectivity.

Monthly Uptime Percentage: is calculated by subtracting from 100% the percentage of Minutes in the Month in which any Single Instance Virtual Machine using premium storage for all Operating System Disks and Data Disks had Downtime.

$$Monthly\ Uptime\ \% = \frac{(Minutes\ in\ the\ Month\ -\ Downtime)}{Minutes\ in\ the\ Month}\ x\ 100$$

Service Credit:

The following Service Levels and Service Credits are applicable to Customer's use of Single-Instance Virtual Machines:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%
< 95%	100%

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 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 for a given VPN Gateway is calculated as Maximum Available Minutes less Downtime dived by the Maximum Available Minutes in a billing month for the VPN Gateway. The Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

The following Service Levels and Service Credits are applicable to Customer's use of each VPN Gateway:

Basic Gateway for VPN or ExpressRoute Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Standard, High Performance, VpnGw1, VpnGw2, Gateway for VPN / Standard, High Performance, Ultra Performance Gateway for ExpressRoute Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.95%	10%
< 99%	25%

Table of Contents / Definitions

Visual Studio App Center Build Service

Additional Definitions:
"**Build Service**" is a feature that allows customers to build their mobile applications in Visual Studio App Center.
"**Maximum Available Minutes**" is the total number of minutes for which Build Service has been deployed by Customer for a given Microsoft Azure subscription during a billing month.
"**Downtime**" is the total number of minutes within Maximum Available Minutes 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 Customer throughout the minute either result in an Error Code or do not return a response within one minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Visual Studio App Center Build Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

The following Service Levels and Service Credits are applicable to Customer's use of the Visual Studio App Center Build Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Visual Studio App Center Test Service

Additional Definitions:
"**Test Service**" is a feature that allows customers to upload and run tests for their mobile applications on physical devices running in Visual Studio App Center.
"**Maximum Available Minutes**" is the total number of minutes for which Test Service has been deployed by Customer for a given Microsoft Azure subscription during a billing month.

Downtime: The total number of minutes within Maximum Available Minutes during which the Test Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to the Test Service to perform operations initiated by Customer throughout the minute either result in an Error Code or do not return a response within one minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Visual Studio App Center Test Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

The following Service Levels and Service Credits are applicable to Customer's use of the Visual Studio App Center Test Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Visual Studio App Center Push Notification Service

Additional Definitions:
"**Push Notification Service**" is a feature that enables customers to push messages to specific devices configured to receive such messages using Visual Studio App Center.
"**Maximum Available Minutes**" is the total number of minutes for which Push Notification Service has been deployed by Customer for a given Microsoft Azure subscription during a billing month.

Downtime: The total number of minutes within Maximum Available Minutes during which Push Notification Service is unavailable. A minute is considered unavailable if all continuous HTTP requests to Push Notification Service to perform operations initiated by Customer throughout the minute either result in an Error Code or do not return a response within one minute.

Monthly Uptime Percentage: The Monthly Uptime Percentage for the Visual Studio App Center Push Notification Service is calculated as Maximum Available Minutes less Downtime divided by Maximum Available Minutes multiplied by 100. Monthly Uptime Percentage is represented by the following formula:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

The following Service Levels and Service Credits are applicable to Customer's use of the Visual Studio App Center Push Notification Service. Free tier service is not covered by this SLA.

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Visual Studio Team Services – Build Service

Additional Definitions:
"**Build Service**" is a feature that allows customers to build their applications in Visual Studio Team Services.
"**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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Visual Studio Team Services – 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Visual Studio Team Services – User Plans Service

Additional Definitions:
"**Build Service**" is a feature that allows customers to build their applications in Visual Studio Team Services.
"**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 Team Services 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

Microsoft Azure Plans

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:

$$\frac{User\ Minutes\ -\ Downtime}{User\ Minutes}\ x\ 100$$

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 B2C

Additional Definitions:
"**Deployment Minutes**" is the total number of minutes for which an Azure AD B2C directory has been deployed during a billing month.
"**Maximum Available Minutes**" is the sum of all Deployment Minutes across all Azure AD B2C directories in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated minutes across all Azure AD B2C directories deployed by Customer in a given Microsoft Azure subscription during which the Azure AD B2C service is unavailable. A minute is considered unavailable if either all attempts to process user sign-up, sign-in, profile editing, password reset and multi-factor authentication requests, or all attempts by developers to create, read, write and delete entries in a directory, fails to return tokens or valid Error Codes, or do not return responses within two minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\ -Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Service Level Exceptions: No SLA is provided for the Free tier of Azure Active Directory B2C.

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 Information Protection Premium

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:

$$\frac{User\ Minutes\ -\ Downtime}{User\ Minutes}\ x\ 100$$

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 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 Azure 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.
"**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**": For a specific Protected Instance configured for On-Premises-to-Azure replication in a given billing month is two hours.

Service Credit:

Monthly Recovery Time Objective	Service Credit
> 2 hours	100%

Additional Terms: Monthly Recovery Time Objective and Service Credits are calculated for each Protected Instance used by you.

Table of Contents / Definitions

Azure 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 Azure 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 Azure 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: Is the total accumulated Failover Minutes in which the Failover of a Protected Instance is unsuccessful due to unavailability of the Azure 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes} \ x\ 100$$

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

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.

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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes} \ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	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.
"**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 by Customer 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 number of minutes within Maximum Available Minutes 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:

$$\frac{Maximum\ Available\ Minutes\text{-}Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

StorSimple Data Manager

Additional Definitions:

"**Total Requests**" is the set of all requests, other than Excluded Requests, to perform operations against StorSimple Data Manager service during a billing month for a given Microsoft Azure subscription.

"**Excluded Requests**" is the set of requests that result in an HTTP 4xx status code.

"**Failed Requests**" is the set of all requests within Total Requests that either return an Error Code or fail to return a Success Code within 60 seconds.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ Requests\text{-}Failed\ Requests}{Total\ Requests}\ x\ 100$$

Service Credit:

Monthly Uptime Percentage	Service Credit
< 99.9%	10%
< 99%	25%

Table of Contents / Definitions

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:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -\ Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -\ Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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

Microsoft Cloud App Security

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:

$$\frac{User\ Minutes\ -\ Downtime}{User\ Minutes}\ x\ 100$$

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%	10%
< 99%	25%

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 Cloud App Security) that provide updates via API (application programming interface) to any services licensed as part of the Service subscription.

Table of Contents / Definitions

Microsoft Flow

Downtime: Any period of time when users' flows have no connectivity to Microsoft's Internet gateway.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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: No SLA is provided for any free of charge tier of Microsoft Flow.

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:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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 PowerApps

Downtime: Any period of time when users are unable to read or write any portion of data in Microsoft PowerApps to which they have appropriate permissions.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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: No SLA is provided for any free of charge tier of Microsoft PowerApps.

Table of Contents / Definitions

Microsoft Stream

Downtime: Any period of time when users are unable to upload, playback, delete video or edit video metadata when they have appropriate permissions and content is valid excluding unsupported scenarios[1].

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{User\ Minutes\ -Downtime}{User\ Minutes}\ x\ 100$$

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%

Service Level Exceptions: No SLA is provided for any free of charge tier of Microsoft Stream.

[1]Unsupported Scenarios could include playback on unsupported devices / OS, client side network issues, and user errors.

Table of Contents / Definitions

Minecraft: Education Edition

Downtime: Any period of time when users are unable to access Minecraft: Education Edition.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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

Power BI Embedded

Deployment Minutes: is the total number of minutes for which a given workspace collection has been provisioned during a billing month.

Maximum Available Minutes: is the sum of all Deployment Minutes across all workspace collections provisioned by a customer in a given Microsoft Azure subscription during a billing month.

Downtime: is the total accumulated Deployment Minutes, during which the workspace collection is unavailable. A minute is considered unavailable for a given workspace collection if all continuous attempts within the minute to read or write any portion of Power BI Embedded data result in an Error Code or do not return a response within five minutes.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\ -Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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%	10%
< 99%	25%

Table of Contents / Definitions

Power BI Premium

"**Capacity**" means a named capacity provisioned by an admin through the Power BI Premium capacity admin portal. A Capacity is a grouping of one or more nodes.
"**Maximum Available Minutes**" is the total number of minutes that a given Capacity has been instantiated during a billing month in a given tenant.

Downtime: The total accumulated minutes during a billing month for a given Capacity during which a given Capacity is unavailable. A minute is considered unavailable for a given Capacity if all attempts to view Power BI reports or dashboards within the minute fail due to system errors.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\ -Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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%	10%
< 99%	25%

Table of Contents / Definitions

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:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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:

$$\frac{Total\ number\ of\ minutes\ in\ a\ month\ -Downtime}{Total\ number\ of\ minutes\ in\ a\ month}\ x\ 100$$

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

Windows Desktop Operating System

Additional Definitions:
"**Maximum Available Minutes**" is the total accumulated minutes during a billing month for Windows Defender Advanced Threat Protection portal. Maximum Available Minutes is measured from when the Tenant has been created resultant from successful completion of the on-boarding process.
"**Tenant**" represents Windows Defender Advanced Threat Protection customer specific cloud environment.

Downtime: The total accumulated minutes that are part of Maximum Available Minutes in which the Customer unable to access any portion of a Windows Defender Advanced Threat Protection portal site collections for which they have appropriate permissions and customer has a valid, active, license.

Monthly Uptime Percentage: The Monthly Uptime Percentage is calculated using the following formula:

$$\frac{Maximum\ Available\ Minutes\ -Downtime}{Maximum\ Available\ Minutes}\ x\ 100$$

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%	10%
< 99%	25%

Service Level Exceptions: This SLA does not apply to any trial/preview version Tenants.

Table of Contents / Definitions

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%

BMC RoD

BMC response targets are established as follows:

Severity	Schedule	Target Response Times
S1	24 hours a day, 7 days a week (including published holidays)	15 Minutes
S2	**Local business hours: 7:00 AM to 7:00 PM, Monday to Friday** (excluding published holidays)	30 minutes
S3	**Local business hours: 7:00 AM to 7:00 PM, Monday to Friday** (excluding published holidays)	4 business hours
S4	**Local business hours: 7:00 AM to 7:00 PM, Monday to Friday** (excluding published holidays)	16 business hours

BMC does not offer service credits for responses that miss the above targets.

Enterprise iPaaS - SLA

State of Utah/NASPO

July 2018

Table of Contents

Submitting Support Requests

The Purchasing Entity may request support by way of a Support Request containing the information described under this section.

Each Support Request shall include a description of the problem, the Purchasing Entity's initial view of the Severity level, and the start time of the Incident.

The Purchasing Entity shall provide the Offerer with:

- prompt notice of any Incidents; and
- such output and other data, documents, information, assistance and (subject to compliance with all Purchasing Entity's security and encryption requirements notified to the Offeror in writing) remote access to the Purchasing Entity's system, as are reasonably necessary to assist the Offeror to reproduce operating conditions similar to those present when the Purchasing Entity detected the relevant fault and to respond to the relevant support request.

All support shall be provided from the Offeror's office.

The Purchasing Entity acknowledges that, to properly assess and resolve Support Requests, it may be necessary to permit the Offeror direct access at the Purchasing Entity's premises and to its system, files, equipment and personnel. The Purchasing Entity shall provide such access promptly.

Service Levels

The Offeror shall:

- prioritise all Support Requests based on its reasonable assessment of the severity level of the Incident reported; and
- respond to all Support Requests in accordance with the Response Times specified in the table set out below:

Severity level of Incident	Definition	Service Level Response
1 Critical Extensive	**Business Critical:** Outage of a Small, Medium or Large Platform Instance leading to total loss of service of all APIs and/or integrations deployed on the Enterprise iPaaS Platform.	2 hours
2 High Significant	**Business Impact:** Impact on key functionality and/or performance degradation of the Enterprise iPaaS Platform but Purchasing Entity critical business functions are still operational. No acceptable workaround is available.	6 hours
3 Medium Moderate	**Operational Impact:** Moderate impact on usage of Enterprise iPaaS Platform but remains operational. A workaround is available to improve the situation until the issue is fully resolved.	12 hours

4 Low Minor	**Minor Impact:** Minor impact on usage of Enterprise iPaaS Platform. Includes minor, cosmetic, or documentation related issues. No impact on Enterprise iPaaS Platform features.	24 hours

On receipt of a Support Request by the Offeror iPaaS Support team, the validity and severity of the Incident is assessed, with a response provided within the above timescales to either accept or return the Incident, and amend the Incident severity in line with the definitions if required.

The maximum number of Support Requests which the Purchasing Entity can raise in each month is as follows:

Number of Platform Instances	Purchasing Entity Incident Allowance Per Month
1-4	60
5-10	120
11-15	180

Actual resolution time will depend on the nature of the fault underlying the Incident.

Subscription and Support Services Exclusions

The following matters are excluded from, and do not form part of, the Subscription and Support Services:

- enhancements to the Enterprise iPaaS Platform that are not part of a planned Release;
- the provision of Enterprise iPaaS Software other than as specified in Section Service Solution;
- the resolution of Incidents due to improper Use or a failure by the Purchasing Entity to comply with a Purchasing Entity Responsibility;
- support of any version of the Enterprise iPaaS Platform which has been discontinued by the Offeror;
- support of any version of the Enterprise iPaaS Platform if the Purchasing Entity refuses to accept an Upgrade Release in respect of a Platform Instance within the forty (40) Business Days' window;
- any Incident caused by the Cloud Hosting Provider services on which the relevant Platform Instance is hosted (and for these purposes, the Purchasing Entity acknowledges that where any Incident relates to any services supplied by the Cloud Hosting Provider, the Offeror's ability to respond and resolve any Incidents will be limited by the service response and support offered by the Cloud Hosting Provider and the service levels as set forth in the Cloud Hosting Provider's terms of service will apply),
- each, an "**Exclusion**".

Where the root cause of an Incident, in the reasonable opinion of the Offeror, is attributable to any one or more of the exclusions set out in Section Service Solution (above), the Service Levels shall not apply to the Incident.

The Offeror may reasonably determine that any services are Out-of-scope Services. If the Offeror makes any such determination, it shall promptly notify the Purchasing Entity of that determination. For the purposes of this Document, "**Out-of scope Services**" shall mean any services provided by

the Offeror in connection with any apparent problem regarding the Enterprise iPaaS Platform reasonably determined by the Offeror not to have been caused by a fault, but rather by a Purchasing Entity cause (including, without limitation, any improper use, misuse or unauthorised alteration of the Enterprise iPaaS Platform by the Purchasing Entity) or a cause outside the Offeror's control (including any investigational work resulting in such a determination).

The Purchasing Entity acknowledges that the Offeror is not obliged to provide Out-of-scope Services.

Service Availability

Subject to the exclusions set out in Section Subscription and Support Services Exclusions, the Offeror shall use reasonable commercial endeavours (subject to planned outages detailed below) to ensure that the Enterprise iPaaS Platform is Available 99.95% of the time during the Support Hours.

The actual Availability shall be stated as a percentage and calculated as follows:

$$\left(\frac{(\mathrm{x}-\mathrm{y})}{\mathrm{x}}\right) \times 100$$

where:

x = the total number of minutes during the Measurement Period less the Scheduled Downtime.

y = the total number of minutes during the Measurement Period where the Enterprise iPaaS Platform is Unavailable (other than unavailability due to Scheduled Downtime or an Exclusion).

Any planned outages as described in section Patching, Maintenance and Upgrades. will be excluded from the above instances and Capgemini will be relieved of its service level obligations during any instance of Force Majeure.

Service Credits

The Purchasing Entity must inform the Offeror of Unavailability of the Enterprise iPaaS Platform within 10 days of the end of the month in which the Purchasing Entity determines the Enterprise iPaaS Platform was Unavailable.

Subject to the Service Credit Cap (in excess of which, no Service Credits shall be payable) and the notice period as defined in the above paragraph, if during any Measurement Period the actual Availability is lower than 99.95% then the Offeror shall:

- credit the Purchasing Entity’s account; or
- if the Subscription Charges have been paid by the Purchasing Entity in advance, reimburse the Purchasing Entity,
- in each case, by an amount calculated as follows:

$$\frac{(99.95-a)}{100} \times b$$

where:

a = the actual Availability achieved (expressed as a percentage);

b = means: (i) the total Subscription Charges payable during the relevant Measurement Period; or (ii) where the Subscription Charges have been paid by the Purchasing Entity in advance of the services, the proportion of those Subscription Charges paid in advance which would have

been payable (on a pro-rata basis) had the Subscription Charges been payable monthly and arrears.

Platform Requests

Subject to paragraph below, during the Subscription Term, the Purchasing Entity iPaaS Account Owner may raise a Platform Request via the Enterprise iPaaS Platform support portal to request the addition or removal of a Platform Instance, or adjust the size of an existing Platform Instance between the following pre-defined sizes ("Micro", "Small", "Medium", and "Large").

Notwithstanding paragraph above:

- the Purchasing Entity shall not be permitted to reduce the number of Platform Instances below one (1) Micro Platform Instance during the Subscription Term;
- the Purchasing Entity shall always be charged for the Shared Foundation Service platform component defined in Section Service Solution Table A;
- the Offeror shall use reasonable commercial endeavours to fulfill the Purchasing Entity's request under the paragraph paragraph above within one (1) Business Day of receipt of the Platform Request;
- the Purchasing Entity may make no more than five (5) Platform Requests in any one (1) calendar month.

Patching, Maintenance and Upgrades

In respect of all Releases:

- the Purchasing Entity must nominate an initial Platform Instance that the Release will first be deployed into prior to deployment of the Release to all Platform Instances;
- the Offeror will notify the Purchasing Entity of the implementation timetable for a Release and any associated Maintenance Window required;
- the Offeror can provide the following minimum notice periods to the Purchasing Entity:
  - Critical Patch – one (1) Business Day;
  - Maintenance Release – five (5) Business Days;
  - Upgrade Release – forty (40) Business Days;
- for a Critical Patch or Maintenance Release, no Purchasing Entity User Testing is required. The Offeror will deploy the Release to each Platform Instance as per the communicated implementation timetable;
- for an Upgrade Release, the Purchasing Entity must undertake User Testing of the Release within the nominated initial Platform Instance within forty (40) Business Days;
- the acceptance criteria for an Upgrade Release is that the iPaaS Platform shall continue to comply with the capabilities described in Section Service Solution Table A;
- once the User Testing of an Upgrade Release is successfully completed, the Offeror will deploy the Release to each Platform Instance as per the communicated implementation timetable;
- if the Purchasing Entity refuses to accept an Upgrade Release in respect of a Platform Instance within the forty (40) Business Days' window, then the Offeror shall be relieved from performing the Support Services in respect of such Platform Instance and the Service Levels (and provision of Service Credits) will no longer apply to that Platform Instance.

Backup and Restore

The Offeror can store backup copies of Purchasing Entity data for each Enterprise iPaaS Platform Instance on an independent storage volume within the same Cloud Hosting Provider hosting location;

Back-ups will be undertaken of the incremental data changes every 4 hours, with a rolling retention period of 7 days, with an auto-overwrite of data from "oldest to newest", commencing on the 8th day to optimize storage volumes; and

Upon a system failure of an Enterprise iPaaS Platform Instance, the Offeror can restore from the latest backup copy to replace the existing Platform Instance.

About Capgemini

A global leader in consulting, technology services and digital transformation, Capgemini is at the forefront of innovation to address the entire breadth of clients' opportunities in the evolving world of cloud, digital and platforms. Building on its strong 50-year heritage and deep industry-specific expertise, Capgemini enables organizations to realize their business ambitions through an array of services from strategy to operations. Capgemini is driven by the conviction that the business value of technology comes from and through people. It is a multicultural company of 200,000 team members in over 40 countries. The Group reported 2017 global revenues of EUR 12.8 billion.

Learn more about us at www.capgemini.com

People matter, results count.

ServiceNow IT Service Management

ServiceNow response targets by priority follows.

Priority	Target Level of Effort	Target Response Times
P1	Continuously, 24 hours per day, 7 days per week	30 minutes
P2	Continuously, but not necessarily 24 hours per day, 7 days per week	2 hours
P3	As appropriate during normal business hours	1 business day
P4	Varies	N/A

virtustream

SERVICE LEVEL AGREEMENT

1. DEFINITIONS

For the purpose of this Service Level Agreement, the following terms shall have the corresponding definitions:

"Availability" means the period in a month that the Services are available, excluding Scheduled Downtime and shall be calculated as follows:

Availability = Maximum Availability – Unscheduled Downtime / Maximum Availability x 100

"Incident" means a report issued to Virtustream by Customer informing Virtustream that the Services are experiencing a Service Failure.

"Maximum Availability" means the total number of minutes in a calendar month, less the Scheduled Downtime.

"Scheduled Downtime" means, the number of hours during a month for which the Services are scheduled to be unavailable in order for Virtustream to perform maintenance or other scheduled services.

"Service Level" means the Availability of the Service in a calendar month.

"Unscheduled Downtime" means the inability of Customer to access the Services as a result of an Incident due to a cause within the control of Virtustream; provided, however, an application server instance being unavailable shall not be considered a Service Failure if the application environment is still available.

2. INCIDENT PRIORITIZATION

All Incidents that are reported to Virtustream or that Virtustream otherwise becomes aware of will initially be assigned a priority by Virtustream as set forth below. Internal escalation for Incidents resources shall be determined by Virtustream based on the priority level assigned to the Incident by Virtustream. The priority/severity level may be adjusted as agreed to by the parties.

Priority/ Severity	Definition	Time to Respond [Note 1]	Customer Communication Interval	Level of Effort
1	Major part of the Services is unavailable/not operating correctly, affecting multiple users. No workarounds are in place, and business operations are not possible. OR Incident has a critical impact on the business (e.g., loss of the Exchange production server impacting all users).	30 minutes	Every 30 minutes	Immediate and continuous effort until the issue is resolved or a workaround is developed

Priority/ Severity	Definition	Time to Respond [Note 1]	Customer Communication Interval	Level of Effort
2	Part of the Services is unavailable/not operating correctly, affecting users in a single function. No workarounds are in place, and business operations in this function are not possible/severely impacted. OR Incident has a serious impact on part of the business (e.g., a configuration change is impacting a small subset of users).	60 minutes	Every 60 minutes	Continuous effort until the issue is resolved or a workaround is developed
3	Part of the Services is unavailable/not operating correctly, affecting users in a single function. Workarounds are in place, but business operations are impacted, although not severely. OR Incident has a temporary impact on users and is non-critical or is a development issue (e.g., email is slow to deliver).	4 hours	Updates provided as available	Work until issue is resolved or a workaround is developed during business hours
4	Incident that is causing inconvenience to the business, but not impacting operations. OR Incident has a minor impact on users or business, or issue is a request for further information.	1 US business day	Not applicable	Will be addressed during the next general update to the services

Note 1: Time to Respond is measured as the time between the proper notification of an Incident, and the Incident being acknowledged within the Service Management System.

3. SERVICE LEVEL CREDITS

Virtustream shall provide Customer with the Service Level Credit if the Services fail to satisfy any of the Service Levels set out herein. The Service Level shall commence thirty (30) days following commencement of the Steady State phase of the SOW. Each of the Service Level Credits shall be based on the fees paid for the applicable service, as set out in this SOW. The aggregate Service Level Credits for all Service Levels in any month shall not exceed 15% of the total monthly recurring charges ("MRCs") for such month set out in Exhibit 1, above. The Service Level Credit shall be Customer's sole and exclusive remedy and Virtustream's sole and exclusive liability for Unscheduled Downtime.

3.1 *Cloud Platform Services - Core and Basic µVM Service Level Credits (including vHANA)*

µVM Availability on Enterprise Core µVM	µVM Availability on Enterprise Basic µVM	µVM Availability on DMZ Core VM	µVM Availability on DMZ Basic VM	Service Level Credit [Note 2]
99.95% – 99.999%	99% - 99.5%	99.5% - <99.9%	99% - <99.4%	1%
99.5% - 99.94%	98% - 98.99%	98% - <99.5%	98% - <99%	3%
95% - 99.4%	95% - 97.99%	95% - <98%	95% - <98%	5%
90% - 94.99%	90% - 94.99%	90% - <95%	90% - <95%	10%
Below 90%	Below 90%	Below 90%	Below 90%	15%

Note 2: The Service Level Credit shall be calculated based on the applicable MRCs paid for the Core and Basic µVM and related products set out in Exhibit 1 to the SOW

3.2 *Application Managed Services*

3.2.1 SAP Support Service Credits

Provided that the Customer is subscribed to Virtustream's AMS SAP Support service, and subject to any Service Level Exception, Virtustream will provide to Customer the service level credits indicated below for the availability of SAP applications.

Service Level (SAP Application Production)	Service Level Credit [Note 3]
99.85% – 99.90%	1%
99.75% - 99.84%	3%
99.50% - 99.74%	5%
99.30% - 99.49%	7%
99.00% - 99.29%	10%
<99.00%	15%

Service Level (SAP Application Non-Production)	Service Level Credit [Note 3]
99.10% - 99.40%	1%
98.70% - 99.09%	3%
98.30% - 98.69%	5%
97.90% - 98.29%	7%
97.50% - 97.89%	10%
< 97.50%	15%

Note 3: The Service Level Credit shall be calculated based on the applicable MRCs paid for the Application Managed Services and related products set out in Exhibit 1, above.

3.2.2 SAP System Application Response Times

Monthly **Production** Average SAP Dialog Response time, as measured by transaction ST03N, will be less than 1,000 milliseconds, excluding custom programs and transactions usually referred to as Z-programs. Customer must be using Tier 1 storage and Enterprise Core uVMs for the Production environment.

4 SERVICE LEVEL CREDIT POLICIES

In the event of Unscheduled Downtime, Virtustream shall promptly address such failure as provided herein:

- If a single Incident results in Virtustream's failure to meet more than one Service Level, Customer shall receive only the highest one of the multiple Service Credits applicable to such Service Level defaults. This shall not affect the Customer's entitlement to Service Credits, as applicable, for any other Service Level defaults that have a root cause other than the specific Incident referred to above.
- Virtustream will credit any Service Level Credit against the charges otherwise payable by Customer to Virtustream for the applicable Services on the next invoice. If no further charges are due and owing to Virtustream, Virtustream shall pay to the Customer the applicable Service Level Credit within forty-five (45) calendar days of the date such credit was incurred.
- If Customer's accounts receivable balance for the Services is not current in the month in which the Service Failure occurred, Customer shall not be entitled to a Service Level Credit and Virtustream will be excused for its failure to meet the Service Level.

5 SERVICE LEVEL EXCEPTIONS

Virtustream shall not be liable for any failure to meet the Service Levels to the extent that one or more of the following caused such failure:

- Failure of the Customer (including any of the Customer's third party service providers) to perform any of its responsibilities under the Agreement or SOW;
- Any act or omission of the Customer (including the Customer's third party service providers), including the Customer's lack of email or telephone availability or delays due to lack of the Customer's response;
- Failure of the Customer's hardware, software, product or equipment;
- Failure of the Customer to secure the proper access rights or maintenance and support services with respect to any component of the Services (e.g., hardware, software, network, maintenance) that does not fall under Virtustream's scope of services that are contracted with the Customer;
- Scheduled Downtime, emergency maintenance or a Force Majeure event;
- Customer's reprioritization of the tasks to be performed by Virtustream where such reprioritization causes Virtustream to miss a Service Level;
- Viruses, provided that the infected Virtustream-provided Services had virus protection for which the virus protection software updates were current;
- Against the advice of Virtustream, the Customer elected to purchase a base commitment to the Services that is not sufficient to run the Customer's system;
- Claims of performance degradation not substantiated through Customer provided diagnostic testing results;
- Failure to meet Service Levels while operating under a business continuity or disaster recovery plan dependent upon customer contract;
- Failures outside the In-Scope environment;
- Infringements of third-party proprietary rights; and
- Resolution delays due to lack of Customer response, including delays resulting from Customer not providing access or login credentials to access non-Virtustream systems relevant to providing a resolution,, in which case the measurement time shall be suspended for the period of the delayed response.

6. RESPONSE TO UNSCHEDULED DOWNTIME

In the event of Unscheduled Downtime, Virtustream shall promptly address such failure as provided herein:

- Promptly investigate and report on the causes of such problem based on the assigned severity level;
- Upon Virtustream's determination of the cause of such failure, it will provide to Customer a preliminary report citing the cause of such failure.
- If Virtustream determines that the failure was due to Virtustream and it is a 'Priority 1' (P1) Incident, then Virtustream will provide a root cause analysis (RCA) as soon as practical after such failure
- Workarounds or fixes are provided for Incidents categorized as P2, P3, or P4, but no RCA will be provided
- Correct or undertake remedial efforts such Service Failure that is Virtustream's fault or responsibility as provided herein;
- Advise Customer of the status of remedial efforts being undertaken with respect to such problem;
- Demonstrate that the causes of such problem (if due to Virtustream's fault or responsibility) has been, or shall be, corrected.
- If applicable, Virtustream will take long-term corrective action using reasonable commercial efforts to minimize the re-occurrence of such failure to prevent any recurrence of such problem (that is Virtustream's fault or responsibility).