



**National Institute of
Standards and Technology**
U.S. Department of Commerce

Special Publication 500-293
(Draft)

US Government Cloud Computing Technology Roadmap Volume II Release 1.0 (Draft)

Useful Information for Cloud Adopters

Lee Badger, Robert Bohn, Shilong Chu, Mike Hogan, Fang Liu, Viktor Kaufmann, Jian Mao, John Messina, Kevin Mills, Annie Sokol, Jin Tong, Fred Whiteside and Dawn Leaf

NIST Cloud Computing Program
Information Technology Laboratory



National Institute of Standards and Technology • U.S. Department of Commerce

This page left intentionally blank

NIST Special Publication 500-293
(Draft)

US Government Cloud Computing Technology
Roadmap Volume II Release 1.0 (Draft)

Useful Information for Cloud Adopters

Lee Badger, Robert Bohn, Shilong Chu, Mike Hogan,
Fang Liu, Viktor Kaufmann, Jian Mao, John Messina,
Kevin Mills, Annie Sokol, Jin Tong, Fred Whiteside and
Dawn Leaf

Information Technology Laboratory

Cloud Computing Program
Information Technology Laboratory
National Institute of Standards and
Technology
Gaithersburg, MD 20899

November 2011



U.S. Department of Commerce
John E. Bryson, Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Under Secretary for Standards and Technology and Director

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analysis to advance the development and productive use of information technology. This Special Publication 500-series reports on ITL's research, guidance, and outreach efforts in Information Technology and its collaborative activities with industry, government, and academic organizations.

National Institute of Standards and Technology Special Publication 500-293 (Draft)

Natl. Inst. Stand. Technol. Spec. Publ. 500-293, 85 pages (November 2011)

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

Acknowledgments

The authors, Shilong Chu, Fang Liu, Viktor Kaufmann, Jian Mao, and Jin Tong, of Knowcean Consulting Incorporated (under contract through the Federal Integrated Product Team, Space & Naval Warfare (SPAWAR) Systems Center Atlantic, Lee Badger, Robert Bohn, Mike Hogan, John Messina, Kevin Mills, Annie Sokol, Fred Whiteside and Dawn Leaf of the National Institute of Standards and Technology (NIST), gratefully acknowledge and appreciate the broad contributions from members of the NIST Cloud Computing US Government Target Business Use Case, Reference Architecture and Taxonomy, Standards Acceleration to Jumpstart the Adoption of Cloud Computing (SAJACC), Security, and Standards Roadmap Working Groups.

We especially acknowledge Carolyn French, Romaine Hines, Michaela Iorga, Elizabeth Lennon, and Peter Mell of NIST for providing technical input and detailed editorial support.

Appendix A presents an extended list of contributors that participated and contributed to NIST public working groups.

Appendix B lists references that were consulted in the development of this document.

Additional acknowledgments will be added upon the final publication of this guideline.

Table of Contents

Executive Summary	11
1 Introduction	13
1.1 NIST Cloud Computing Program Background.....	13
1.2 NIST Cloud Computing Program Vision.....	13
1.3 Intended Audience and Use	14
1.4 Document Organization	14
2 NIST Cloud Computing Definition and Reference Architecture	16
2.1 Revisiting the Definition.....	16
2.2 NIST Cloud Computing Reference Architecture.....	18
2.2.1 Conceptual Model	19
2.2.2 Cloud Computing Actors	19
2.2.3 Architecture Components.....	23
2.3 NIST Cloud Computing Taxonomy.....	24
3 Cloud Computing Use Cases and Requirements.....	26
3.1 Target Business Use Case and High-Level Requirements.....	26
3.1.1 Business Use Case Template.....	28
3.1.2 Business Use Case Summaries.....	28
3.1.3 Business Use Case Analysis.....	31
3.2 SAJACC Use Cases and Technical Requirements.....	44
4 Cloud Computing Standards and Gap Analysis	46
4.1 Cloud Computing Standards	46
4.2 Cloud Computing Standards Gaps and USG Priorities.....	47
4.3 Accelerating the Development and the Use of Cloud Computing Standards	48
5 High-Priority Security Requirements.....	49
5.1 Understanding Security in the Cloud Context	50
5.1.1 Cloud Service Model Perspectives.....	50
5.1.2 Implications of Cloud Deployment Models	50
5.1.3 Shared Security Responsibilities.....	50
5.1.4 Developing Security Architecture for Cloud Systems	51

5.2	Challenging Security Requirements and Risk Mitigations	51
5.3	Process-Oriented Requirements.....	52
5.3.1	Application of NIST SP 800-53-style Controls and Compliance	52
5.3.2	Cloud Audit Assurance and Log Sensitivity Management	53
5.3.3	Cloud Certification and Accreditation Guidelines	56
5.3.4	Clear eDiscovery Guidelines.....	56
5.3.5	Cloud Privacy Guidelines.....	57
5.3.6	Clarified Security Control Roles and Responsibilities.....	59
5.3.7	Trustworthiness of Cloud Operators	59
5.3.8	Business Continuity and Disaster Recovery	60
5.4	Focused Technical Requirements	61
5.4.1	Technical Continuous Monitoring Capabilities	62
5.4.2	Visibility for Customers.....	63
5.4.3	Control for Customers.....	63
5.4.4	Data Protection Concerns.....	65
5.4.5	Risk of Account Compromise	66
5.4.6	Identity and Access Management (IAM) and Authorization	67
5.4.7	Multi-tenancy Risks and Concerns	69
5.4.8	Cloud-Based Denial of Service	70
5.4.9	Incident Response	71
6	Other Related Work	72
6.1	Cloud Data Issues	72
6.1.1	Operational Data Functions.....	72
6.1.2	Informational Data and Data Services	74
6.2	Service-Level Agreement Taxonomy	75
6.3	Reliability Research in Cloud-based Complex Systems.....	77
7	Summary and Next Steps	78
Appendix A. USG Federal Cloud Computing Standards Working Group Members and Interagency, Academic, Standards Organizations, and Industry Contributors		79
Appendix B. Useful References.....		83

Executive Summary

The first release of the Special Publication 500-293, United States Government *USG Cloud Computing Technology Roadmap* document consists of two volumes. Consistent with the NIST Cloud Computing program strategy, the roadmap focuses on both strategic and tactical objectives related to cloud computing.

Volume I, *High-Priority Requirements to Further USG Cloud Computing Adoption*, frames the discussion and introduces the roadmap in terms of summarized strategic requirements that must be met for USG agencies to further cloud adoption. The roadmap strategic elements can be characterized as “high-priority technical areas” which are enablers for cloud computing in both the short and long term.

Volume II, *Useful Information for Cloud Adopters*, provides information for those actively working on strategic and tactical cloud computing initiatives, including but not limited to, government cloud adopters.

This volume presents a summary of the work completed from November 2010 through September 2011 through the NIST Cloud Computing program and collaborative effort to develop a USG Cloud Computing Technology Roadmap.

This document presents a representative sample of the work that was completed and documented through this effort. Additional working documents, special publications, meeting and other collaboration artifacts can be found on the NIST Cloud Computing Web site <http://www.nist.gov/itl/cloud/index.cfm>.

Volume II:

- Introduces a conceptual model, the NIST Cloud Computing Reference Architecture and Taxonomy;
- Presents USG target business use cases and technical use cases in the cloud;
- Identifies existing interoperability, portability, and security standards that are applicable to the cloud computing model and specifies high-priority gaps for which new or revised standards, guidance, and technology need to be developed;
- Discusses security challenges in the context of cloud computing adoption, high-priority security requirements, and current and future risk mitigation measures requirements; and
- Provides insight into the rationale for the list of candidate Priority Action Plans (PAPs) recommended for voluntary self-tasking by government and private sector organizations, listed in Volume I.

The document presents a subset of the analysis that drove the rationale for the requirements introduced in Volume I of this NIST Special Publication, titled *High-Priority Requirements to Further USG Agency Cloud Computing Adoption*.

The following Figure 1 shows the relationship between the high-priority requirements in Volume I and the key NIST-led activities and contributing sources that are summarized here in Volume II.

	Cloud Computing Standards Roadmap Working Group	Cloud Computing Reference Architecture and Taxonomy Working Group	Cloud Computing Security Working Group	Cloud Computing Target USG Business Use Cases Working Group	Adoption of Cloud Computing Acceleration to Jumpstart the	NIST Special Publications: 800-125 Security for Virtualization, 800-144 Guidelines for Security and Privacy in Cloud Computing, 800-146 Cloud Computing Synopsis and Recommendations	NIST Complex Information System Measurement Project: Koala, IaaS Computing Simulation
Requirement 1: International voluntary consensus based interoperability, portability and security standards	X				X	X	
Requirement 2: Solutions for high priority security requirements	X		X	X	X	X	
Requirement 3: Technical specifications for high quality service level agreements		X	X			X	
Requirement 4: Clear & consistently categorized cloud services		X					
Requirement 5: Frameworks to support federated community clouds		X		X		X	
Requirement 6: Technical security solutions de-coupled from organizational policy			X				
Requirement 7: Defined unique government requirements and solutions			X	X		X	
Requirement 8: Collaborative parallel “future cloud” development initiatives				X			
Requirement 9: Defined & implemented reliability design goals			X	X		X	X
Requirement 10: Defined & implemented cloud service metrics		X	X	X			X

Figure 1: Relationship between Volume I Requirements and Work Presented in Volume II

1 Introduction

1.1 NIST Cloud Computing Program Background

The National Institute of Standards and Technology plays a technology leadership role in accelerating the federal government's secure adoption of cloud computing. In this role, NIST, in close consultation and collaboration with standards bodies, the private sector, and other stakeholders, is leading the efforts to develop the necessary standards and guidelines that will facilitate the secure, rapid adoption of cloud computing.

The NIST Cloud Computing Program was formally launched in November 2010, and supports the US federal government effort to incorporate cloud computing, where appropriate, as a replacement for, or enhancement of, the traditional information systems and application models. The NIST Cloud Computing Program operates in coordination with other federal cloud computing efforts and is integrated within the Federal Cloud Computing Strategy.¹

For more information regarding the program's scope and objectives, the reader is referred to Volume I of this NIST Special Publication 500-293: *High-Priority Requirements to Further USG Agency Cloud Computing Adoption*.

In order to leverage the expertise of the broad cloud computing stakeholder community, NIST has established the following Public Working Groups:

- Cloud Computing Reference Architecture and Taxonomy Working Group
- Cloud Computing Target Business Use Cases Working Group
- Cloud Computing SAJACC Technical Use Cases Working Group
- Cloud Computing Standards Roadmap Working Group
- Cloud Computing Security Working Group

The groups are listed in the same sequence that their respective subject matter is presented in this document. The order does not imply priority or chronological sequencing.

1.2 NIST Cloud Computing Program Vision

NIST seeks to provide thought leadership and guidance around the cloud computing model to catalyze its use within industry and government, and to shorten the adoption cycle, which will enable near-term cost savings and increased ability to quickly create and deploy safe and secure enterprise solutions. Additionally, NIST is committed to fostering cloud computing practices that support interoperability, portability, and security requirements that are appropriate and achievable for various usage scenarios, by focusing on the necessary standards, specifications, and guidance that must be in place for these requirements to be met.

The first release of the USG Cloud Computing Technology Roadmap is presented as a two-volume NIST Special Publication 500-293 document. The process and document together are the mechanism used to define and communicate the high-priority USG interoperability, portability, and security requirements for cloud computing, and to identify the necessary associated standards, guidance, and technology.

¹ Office of Management and Budget, U.S. Chief Information Officer, Federal Cloud Computing Strategy, Feb. 8, 2011. Online: www.cio.gov/documents/Federal-Cloud-Computing-Strategy.pdf.

This document, Volume II of the Special Publication, focuses on work that helped to identify the USG high-priority interoperability, portability, and security requirements which are introduced in Volume I and summarizes work in the following areas:

- Introduction of an overall cloud computing conceptual model in the form of the NIST Cloud Computing Reference Architecture and Taxonomy. This technical reference can be used to understand, discuss, categorize, and compare different cloud service offerings, and to facilitate the communication and analysis of the security, interoperability, and portability candidate standards and reference implementations.
- Presentation of a template and an initial set of USG target business and technical use cases that describe how government agencies seek to use cloud computing, and presentation of key, specific technical requirements that surfaced through these use cases.
- Identification of existing interoperability, portability, and security standards and guidance that are applicable to the cloud computing model, and identification of high-priority gaps for which new or revised standards, guidance, and technology need to be developed.
- Identification of the high-priority security requirements that challenge the adoption of cloud computing and presentation of proposed mitigation strategies.
- Discussion of considerations and activities related to cloud Service-Level Agreements (SLAs).

1.3 Intended Audience and Use

This publication is intended for a diverse audience:

- **US Policy Makers, US Federal CIO Council, and those with identified key roles identified in the Federal Cloud Computing Strategy** – as a technology-oriented reference to inform policy and planning.
- **USG Agencies** – as a useful tool in the context of the USG Federal Cloud Computing Strategy risk-based management decision framework.
- **Cloud Computing Stakeholders (Academia, Government, Industry, Standards Developing Organizations)** – as a consolidated presentation of USG cloud computing technology perspectives and work, including a unifying cloud computing reference model, a set of documented technical requirements, and a list of identified gaps in standards, guidance, and technology.

1.4 Document Organization

Consistent with the NIST Cloud Computing program strategy, the roadmap focuses on both strategic and tactical objectives related to cloud computing. The strategic roadmap elements can be characterized as “high-priority technical areas” which are enablers for cloud computing in both the short and long term. The tactical work not only supports strategic goals, but is intended to support cloud adopters in the interim deployment period as the cloud computing model is maturing.

This initial release of the roadmap special publication consists of two volumes.

Volume I is aimed at interested parties who wish to gain a general understanding and overview of the background, purpose, context, work, results, and next steps of the USG Cloud Computing Technology Roadmap initiative. Volume I reflects the collective inputs of USG agencies through the Federal CIO Council-sponsored *Cloud Computing Standards and Technology Working Group*.

Volume I, *High-Priority Requirements to Further USG Cloud Computing Adoption*, frames the discussion and introduces the roadmap in terms of:

- Prioritized strategic and tactical requirements that must be met for USG agencies to further cloud adoption;
- Interoperability, portability, and security standards, guidelines, and technology that must be in place to satisfy these requirements; and
- Recommended list of Priority Action Plans (PAPs) as candidates for development and implementation, through voluntary self-tasking by the cloud computing stakeholder community, to support standards, guidelines, and technology development.

This volume, Volume II, *Useful Information for Cloud Adopters*, is designed to be useful at the tactical level to those actively working on cloud computing initiatives, including but not limited to, US government cloud adopters. Volume II summarizes the work completed to date, explains the assessment findings based on this work, and highlights how these findings support the key requirements in the roadmap introduced in Volume I.

The Executive Summary of this volume includes a chart that shows the correlation between the set of high-priority USG requirements presented in Volume I, and the NIST projects and public working group efforts and findings summarized in Volume II.

The remainder of Volume II is organized into the following sections: Section 2 presents the NIST cloud computing definition and reference architecture. Section 3 presents USG cloud computing requirements through business use cases and technical use cases. Section 4 summarizes cloud computing technology standards and gap analysis. Section 5 discusses cloud computing security and presents a list of security impediments and corresponding mitigations.

A third volume, *Technical Considerations for USG Cloud Computing Deployment Decisions*, is under development, and in keeping with the NIST transparent and collaborative process, is currently available as a working document. Volume III is being developed as an interagency project through the Federal Cloud Computing Standards and Technology Working Group, and will leverage the NIST-led cloud computing program public working group process. Volume III is intended to serve as a guide for decision makers who are planning and implementing cloud computing solutions by explaining how the technical work and resources in Volume II can be applied, consistent with the Federal Cloud Computing Strategy “Decision Framework for Cloud Migration.” The current version of the working document defines and proposes a methodology for defining a representative sample of common cloud computing planning and deployment scenarios, presents the initial candidate set of 12, presents a process for applying the technical work, and proof-of-concept examples of how this can be accomplished. Volume III was initiated in parallel, but is logically dependent on the technical work contained in Volume II, and will necessarily be completed and presented as part of the roadmap special publication in a subsequent release.

The Volume I and Volume II draft special publications, as well as the working document under development as Volume III, are publically available through the NIST ITL Cloud Computing Web site, as are all of the NIST Cloud Computing special publications and work-in-progress documents, <http://www.nist.gov/itl/cloud/index.cfm>.

2 NIST Cloud Computing Definition and Reference Architecture

Cloud computing is an emerging computing model which has evolved as a result of the maturity of underlying prerequisite technologies. There are differences in perspective as to when a set of underlying technologies becomes a “cloud” model. In order to categorize cloud computing services, and to expect some level of consistent characteristics to be associated with the services, cloud adopters need a consistent frame of reference. The NIST Cloud Computing Reference Architecture and Taxonomy document defines a standard reference architecture and taxonomy that provide the USG agencies with a common and consistent frame of reference for comparing cloud services from different service providers when selecting and deploying cloud services to support their mission requirements. At a certain level of abstraction, a cloud adopter does not need to repeatedly interpret the technical representation of cloud services available from different vendors. Rather the use of a common reference architecture by the cloud service providers can be an efficient tool that ensures consistent categorization of the services offered.

Highlights: The NIST cloud computing definition identifies three distinct service models, i.e., Software as a Service, Platform as a Service, and Infrastructure as a Service.

In late 2010, the NIST Cloud Computing Reference Architecture project team surveyed and completed an analysis of existing cloud computing reference models, and developed a vendor-neutral reference architecture which extends the NIST cloud computing definition.

This effort leveraged a collaborative process through the NIST Cloud Computing Reference Architecture and Taxonomy working group. Through a discussion and validation process, the NIST cloud computing reference architecture project team and working group analyzed the intricacies of different types of cloud services and confirmed the need for “Clear and Consistently Categorized Cloud Services” - NIST USG Cloud Computing Technology Roadmap Volume I, Requirement 4.

The NIST cloud computing definition and reference architecture provide a technical basis for discussing “Frameworks to support federated community clouds” - Volume I, Requirement 5.

The companion NIST cloud computing taxonomy effort has also identified the need for: “Technical specification for high-quality service-level agreements” – Volume I, Requirement 3, and “defined and implemented cloud service metrics” – Volume I, Requirement 10.

See NIST Special Publication 800-145, *The NIST Definition of Cloud Computing*, and NIST Special Publication 500-292, *NIST Cloud Computing Reference Architecture*.

2.1 Revisiting the Definition

This document uses the NIST SP 800-145, *The NIST Cloud Computing Definition*, to explain characteristics of cloud computing. For the convenience of the reader, the following is excerpted from NIST SP 800-145:

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

This definition lists **five essential characteristics** that are common among all cloud computing services:

On-demand self-service: A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service’s provider.

Broad network access: Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and personal digital assistants [PDAs]).

Resource pooling: The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the subscriber generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Rapid elasticity: Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured Service: Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Service Models

Cloud Software as a Service (SaaS): The capability provided to the consumer to use the provider's applications running on a 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). 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.

Cloud Platform as a Service (PaaS): The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. 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.

Cloud Infrastructure as a Service (IaaS): The capability provided to the consumer is 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 the operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models

Based on how exclusive the cloud infrastructure is operated and made available to a consumer, cloud services can also be categorized by a series of deployment models:

Private cloud: The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premise or off-premise.

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

considerations). It may be managed by the organizations or a third party and may exist on-premise or off-premise.

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

Hybrid cloud: The cloud infrastructure is a composition of two or more clouds (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).

2.2 NIST Cloud Computing Reference Architecture

The NIST cloud computing reference architecture is a logical extension to the NIST cloud computing definition. This extension provides a common frame of reference to help USG and other cloud computing stakeholders to:

- Gain a further understanding of the technical and operational intricacies of cloud computing;
- Communicate cloud consumers' requirements precisely;
- Categorize and compare cloud services objectively; and
- Analyze security, interoperability, and portability requirements systematically in order to better inform solution implementations.

The reference architecture describes a conceptual model comprising abstract architectural elements and their relations or interactions, such as

- Cloud computing actors and how they interact with each other in their activities;
- System components and how these components are orchestrated to deliver the computing services;
- Management functionalities that are required to support the life cycle of operations; and
- Other cross-cutting aspects such as security and privacy associated with these elements.

The reference architecture is a high-level, abstract model not tied to any specific cloud technology or vendor product, that focuses on the requirements of “what” cloud services provide and not on “how to” design and implement these services.

The reference architecture also provides a companion cloud computing taxonomy detailing the definitions and relationships of a control vocabulary.

A cloud solution provider may use this reference architecture to guide the development of real architectures from different viewpoints (such as application architecture, middleware architecture, data architecture, and network architecture), given constraints imposed by the organization's operational and technical environments. The reference architecture has a direct benefit for the cloud consumer as well. By mapping the various cloud solution products to the architectural components defined in the reference architecture, a cloud consumer can understand and compare cloud service offerings and make informed decisions. For other stakeholders, such as academia and Standards Developing Organizations (SDOs), the reference architecture can help frame issues and provide a common baseline for research.

As described above, the NIST Cloud Computing Reference Architecture Project Team surveyed and completed an initial analysis of existing cloud computing reference architectures and reference models. On this basis, the project team developed a straw man model of architectural concepts. This effort leveraged a collaborative process from the NIST Cloud Computing Reference Architecture and Taxonomy Working Group, active between November 2010 and April 2011. This process involved broad

participation from the industry, academic, SDOs, and private and public sector cloud adopters. The project team iteratively revised the reference model by incorporating comments and feedback received from the working group. This section summarizes version 1.0 of the reference architecture and taxonomy.

2.2.1 Conceptual Model

Figure 2 presents the NIST cloud computing reference architecture, which identifies the major actors, their activities, and their functions in cloud computing. The diagram depicts a generic high-level architecture and is intended to facilitate the understanding of the requirements, uses, characteristics, and standards of cloud computing.




Figure 2: The Conceptual Reference Model

2.2.2 Cloud Computing Actors

As shown in Figure 2, the NIST cloud computing reference architecture defines five major actors: cloud consumer, cloud provider, cloud carrier, cloud auditor, and cloud broker. Each actor is an entity (a person or an organization) that participates in a transaction or process or performs tasks in cloud computing.

2.2.2.1 Cloud Consumer

The cloud consumer is the principal stakeholder that uses the cloud computing services. A cloud consumer represents a person or organization that maintains a business relationship with, and uses the service from, a cloud provider. A cloud consumer browses the service catalog from a cloud provider, requests the appropriate service, sets up service contracts with the cloud provider, and uses the service. The cloud consumer may be billed for the service provisioned, and needs to arrange payments accordingly.

Cloud consumers use Service-Level Agreements (SLAs) for specifying the technical performance requirements to be fulfilled by a cloud provider. SLAs can cover terms regarding the quality of service, security, and remedies for performance failures. A cloud provider may also list in the SLAs a set of restrictions or limitations, and obligations that cloud consumers must accept. In a mature market environment, a cloud consumer can freely choose a cloud provider with better pricing and more favorable

terms. Typically a cloud provider's public pricing policy and SLAs are nonnegotiable, although a cloud consumer who expects to have heavy usage might be able to negotiate for better contracts.

SaaS applications are made accessible via a network to the SaaS consumers. The consumers of SaaS can be organizations that provide their members with access to software applications, end users who directly use software applications, or software application administrators who configure applications for end users. SaaS consumers can be billed based on the number of end users, the time of use, the network bandwidth consumed, the amount of data stored, or the duration of stored data.

PaaS consumers employ the tools and execution resources provided by cloud providers to develop, test, deploy, and manage the operation of PaaS applications hosted in a cloud environment. PaaS consumers can be application developers who design and implement application software, application testers who run and test applications in a cloud-based environment, application deployers who publish applications into the cloud, and application administrators who configure, monitor, and manage applications deployed in a cloud. PaaS consumers can be billed according to the number of PaaS users, the processing, storage, and network resources consumed by the PaaS application, and the duration of the platform usage.

IaaS clouds provide cloud consumers with virtual computers, network-accessible storage, network infrastructure components, and other fundamental computing resources, on which IaaS consumers can deploy and run arbitrary software. IaaS can be used by system developers, system administrators, and IT managers who are interested in creating, installing, monitoring, and managing services and applications deployed in an IaaS cloud. IaaS consumers can be billed according to the amount or duration of the resources consumed, such as CPU hours used by virtual computers, volume and duration of data stored, network bandwidth consumed, or the number of IP addresses used for certain intervals.

2.2.2.2 Cloud Provider

A cloud provider is the entity (a person or an organization) responsible for making a service available to interested parties. A cloud provider acquires and manages the computing infrastructure required for providing the services, runs the cloud software that provides the services, and makes the arrangements to deliver the cloud services to cloud consumers through network access.

For SaaS, the cloud provider deploys, configures, maintains, and updates the operation of the software applications on a cloud infrastructure. The SaaS cloud provider is mostly responsible for managing the applications, security, and the cloud infrastructure, while the SaaS cloud consumer has limited administrative control of the applications.

For PaaS, the cloud provider manages the computing infrastructure for the platform and runs the cloud software that provides the components of the platform, such as runtime software execution stack, databases, and other middleware components. The PaaS cloud provider typically also supports the development, deployment, and management process of the PaaS cloud consumer by providing tools such as integrated development environments (IDEs), development versions of cloud software, software development kits (SDKs), and deployment and management tools. The PaaS cloud consumer has control over the applications and possibly over some of the hosting environment settings, but has no or limited access to the infrastructure underlying the platform such as network, servers, operating systems (OSs), or storage.

For IaaS, the cloud provider acquires the physical computing resources underlying the service, including the servers, networks, storage, and hosting infrastructure. The cloud provider runs the cloud software necessary to render the necessary computing resources to the IaaS cloud consumer through a set of service interfaces and computing resource abstractions, such as virtual machines and virtual network interfaces. In return, the IaaS cloud consumer uses these computing resources, such as a virtual computer,

for fundamental computing needs. Compared to SaaS and PaaS consumers, an IaaS consumer has access to more fundamental forms of computing resources and thus has control over more software components in an application stack, including the OS. The IaaS cloud provider, on the other hand, has control over the physical hardware and cloud software that make the provisioning of these infrastructure services possible, for example, the physical servers, network equipment, storage devices, host OS, and hypervisor software for virtualization.

A cloud provider's activities span five major areas including service deployment, service orchestration, cloud service management, security, and privacy.

2.2.2.3 Cloud Auditor

A cloud auditor is a party that can perform an independent examination of cloud service controls with the intent to express an opinion thereon. Audits are performed to verify conformance to standards through a review of objective evidence. A cloud auditor can evaluate the services provided by a cloud provider such as security controls, privacy impact, and performance.

Auditing is especially important for federal agencies. The Federal Cloud Computing Strategy document published in February 2011 pointed out that "agencies should include a contractual clause enabling third parties to assess security controls of cloud providers." Security controls are the management, operational, and technical safeguards or countermeasures employed within an organizational information system to protect the confidentiality, integrity, and availability of the system and its information. For security auditing, a cloud auditor can make an assessment of the security controls in the information system to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to the security requirements for the system. The security auditing should also assess the compliance with the specified regulation and with the security policy. For example, an auditor can be tasked with ensuring that the correct policies are applied to data retention according to relevant rules for the jurisdiction. The auditor may ensure that fixed content has not been modified and that the legal and business data archival requirements have been satisfied.

A privacy impact audit can help federal agencies comply with applicable privacy laws and regulations governing an individual's privacy, and to ensure confidentiality, integrity, and availability of an individual's personal information at every stage of development and operation.

2.2.2.4 Cloud Broker

As cloud computing evolves, the integration of cloud services can be too complex for cloud consumers to manage. A cloud consumer may request cloud services from a cloud broker, instead of contacting a cloud provider directly. A cloud broker is an entity that manages the use, performance, and delivery of cloud services and negotiates relationships between cloud providers and cloud consumers.

In general, a cloud broker can provide services in three categories:

- *Service Intermediation*: A cloud broker enhances a given service by improving some specific capability and providing value-added services to cloud consumers. The improvement can be managing access to cloud services, identity management, performance reporting, enhanced security, etc.
- *Service Aggregation*: A cloud broker combines and integrates multiple services into one or more new services. The broker provides data integration and ensures the secure data movement between the cloud consumer and multiple cloud providers.

- *Service Arbitrage*: Service arbitrage is similar to service aggregation except that the services being aggregated are not fixed. Service arbitrage means a broker has the flexibility to choose services from multiple agencies. The cloud broker, for example, can use a credit-scoring service to measure and select an agency with the best score.

2.2.2.5 Cloud Carrier

A cloud carrier acts as an intermediary that provides connectivity and transport of cloud services between cloud consumers and cloud providers. Cloud carriers provide access to consumers through network, telecommunication, and other access devices. For example, cloud consumers can obtain cloud services through network access devices, such as desktop computers, laptops, mobile phones, and other mobile Internet devices (MIDs). The distribution of cloud services is normally provided by network and telecommunication carriers or a transport agent, where a transport agent refers to a business organization that provides physical transport of storage media, such as high-capacity hard drives.

2.2.2.6 Scope of Control between Provider and Consumer

The cloud provider and cloud consumer share the control of resources in a cloud system. As illustrated in Figure 3, different service models affect an organization's control over the computational resources and thus what can be done in a cloud system. Figure 3 shows these differences using a classic software stack notation comprised of the application, middleware, and OS layers. This analysis of delineation of controls over the application stack increases understanding of the responsibilities of parties involved in managing the cloud application.




Figure 3: Scope of Controls between Provider and Consumer

- The application layer includes software applications targeted at end users or programs. The applications are used by SaaS consumers, or installed/managed/maintained by PaaS consumers, IaaS consumers, and SaaS providers.
- The middleware layer provides software building blocks (e.g., libraries, database, and Java virtual machine) for developing application software in the cloud. The middleware is used by PaaS consumers, installed/managed/maintained by IaaS consumers or PaaS providers, and hidden from SaaS consumers.

- The OS layer includes operating system and drivers, and is hidden from SaaS and PaaS consumers. An IaaS cloud allows one or multiple guest OSs to run virtualized on a single physical host. Generally, consumers have broad freedom to choose which OS is to be hosted among all the OSs that could be supported by the cloud provider. The IaaS consumers should assume full responsibility for the guest OS(s), while the IaaS provider controls the host OS.

2.2.3 Architecture Components

This section describes the architectural elements with which cloud actors interact, including an abstraction of the system components that orchestrate together to deliver the service capabilities, the different deployment models of these infrastructure components, and the management activities cloud providers engage in with cloud consumers.

2.2.3.1 Service Orchestration

Service Orchestration refers to the composition of system components to support the cloud provider activities in arrangement, coordination, and management of computing resources in order to provide cloud services to cloud consumers. Figure 4 shows a generic stack diagram of this composition that underlies the provisioning of cloud services.




Figure 4: Cloud Provider - Service Orchestration

A three-layered model is used in this representation to depict the grouping of the three types of system components that cloud providers need to compose to deliver their services.

In the model shown in Figure 4, the top is the service layer, where cloud providers define interfaces for cloud consumers to access the computing services. Access interfaces of each of the three service models are provided in this layer. It is possible, though not necessary, that SaaS applications can be built on top of PaaS components, and PaaS components can be built on top of IaaS components. The optional dependency relationships among SaaS, PaaS, and IaaS components are represented graphically as components stacking on each other; while the angling of the components represents that each of the service component can stand by itself. For example, a SaaS application can be implemented and hosted on

virtual machines from an IaaS cloud, or it can be implemented directly on top of cloud resources without using IaaS virtual machines.

The middle layer in the model is the resource abstraction and control layer. This layer contains the system components that cloud providers use to provide and manage access to the physical computing resources through software abstraction. Examples of resource abstraction components include software elements such as hypervisors, virtual machines, virtual data storage, and other computing resource abstractions. The resource abstraction needs to ensure efficient, secure, and reliable usage of the underlying physical resources. While virtual machine technology is commonly used at this layer, other means of providing the necessary software abstractions are also possible. The control aspect of this layer refers to the software components that are responsible for resource allocation, access control, and usage monitoring. This is the software framework that ties together the numerous underlying physical resources and their software abstractions to enable resource pooling, dynamic allocation, and measured service. Various open source and proprietary cloud software are examples of this type of middleware.

The lowest layer in the stack is the physical resource layer, which includes all the physical computing resources. This layer includes hardware resources, such as computers (CPU and memory), networks (routers, firewalls, switches, network links, and interfaces), storage components (hard disks), and other physical computing infrastructure elements. It also includes facility resources, such as heating, ventilation and air conditioning (HVAC), power, communications, and other aspects of the physical plant.

Following system architecture conventions, the horizontal positioning, i.e., the layering, in a model represents dependency relationships – the upper layer components are dependent on adjacent lower layer to function. The resource abstraction and control layer exposes virtual cloud resources on top of the physical resource layer and supports the service layer where cloud services interfaces are exposed to cloud consumers. Cloud consumers do not have direct access to the physical resources.

2.2.3.2 Cloud Service Management

Cloud Service Management includes all of the service-related functions that are necessary for the management and operation of those services required by or proposed to cloud consumers. Cloud service management can be described from the perspective of business support, provisioning and configuration, and from the perspective of portability and interoperability requirements.

2.3 NIST Cloud Computing Taxonomy

The NIST Cloud Computing taxonomy was developed in conjunction with the reference architecture and describes key cloud computing concepts, the relationships between these concepts, and their given context. The taxonomy organizes the key concepts into four levels:

- *Level 1: Role*, which indicates a set of obligations and behaviors as conceptualized by the associated actors in the context of cloud computing;
- *Level 2: Activity*, which entails the general behaviors or tasks associated to a specific role;
- *Level 3: Component*, which refers to the specific processes, actions, or tasks that must be performed to meet the objective of a specific activity; and
- *Level 4: Sub-component*, which presents a modular part of a component.

The taxonomy can be used as a source for developing a controlled vocabulary of cloud computing terms that will provide an increased clarification and standardization of the cloud computing terminology. Details about this taxonomy and the related vocabulary can be found on the NIST cloud computing

reference architecture and taxonomy collaboration site: <http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/ReferenceArchitectureTaxonomy>.

3 Cloud Computing Use Cases and Requirements

Although use cases have been traditionally employed as a system analysis tool that links the actors to the system functions, the same methodology has also been widely used within business architectures for such purposes as describing business processes of an enterprise, actors corresponding to these processes, and organizational participants. Using well-defined elements such as actors, conditions, and activity flows, a use case can systematically reveal the requirements and constraints which can subsequently direct system architecture and design.

- The NIST projects and working groups apply use case methodology to define business and technical operational scenarios and requirements.
-
- Business use cases document scenarios at the functional mission level. The use case describes the business goal with no assumptions as to how cloud computing technology (deployment model constraints) will be deployed to achieve that goal. These business use cases can then be explored by walking through the considerations of planning and deploying candidate cloud computing service and deployment model options, issues, and constraints. While this process has documented business use cases that are in pilot or operational deployment stage, the objective of the Target focuses on those business use cases that agencies have identified as an opportunity, but consider to be difficult to implement, or have a perceived impediment to implementation.
-
- The second case where use case methodology is applied is definition of technical use cases in the Standards Acceleration to Jumpstart the Adoption of Cloud Computing (SAJACC) effort. These use cases are designed to facilitate the qualitative testing of standards through the use of third-party APIs implemented in adherence to candidate specifications and emerging standards. Of necessity, each SAJACC use case represents a single activity, such as the deletion of data, and the actions needed to successfully execute that activity (receive the request, respond to the request, execute the request, etc.).

A business use case is decomposed into a list of high-level requirements, then into successively more detailed requirements, until it can ultimately be mapped to technical requirements that are required to identify and execute the appropriate SAJACC use cases.

3.1 Target Business Use Case and High-Level Requirements

The main objective of the NIST Cloud Computing Target Business Use Case (TBUC) Project is to work with federal CIOs to identify and document application and service use cases for potential migration to a cloud environment. As described in the Federal Cloud Computing Strategy, NIST is working with agencies to define target business use cases that are complex, or have technical hurdles or standards gaps that need to be overcome. The high-level requirements that are extracted from the target business use cases are the primary deliverables for that project area.

At the time of this writing, the business use cases from agencies and departments summarized here have perceived impediments or obstacles that prevent their immediate implementation or require workarounds. These business use cases focus attention on the areas where technical and procedural gaps are assessed and prioritized to propose recommendations for mitigation. After target business use cases are developed,

Highlights: Target business use cases of federal agencies were captured to understand security, interoperability, and portability requirements. These business use cases and the cross-cutting requirements extracted were developed as part of the iterative and complementary process used to identify the strategic requirements in Volume I of the Technical Roadmap.

Specifically, this section summarizes the use cases which were used as the basis for defining the following high-priority requirements listed in the **NIST USG Cloud Computing Technology Roadmap Volume I high-priority requirements**:

- “Solutions for High-priority Security Requirements” - Requirement 2;
- “Frameworks to support federated community clouds” Requirement 5;
- “Defined unique government regulatory requirements, technology gaps, and solutions”- Requirement 7;
- “Collaborative parallel ‘future cloud’ development initiatives”-Requirement 8;
- “Defined and implemented reliability design goals” – Requirement 9; and
- “Defined and implemented cloud service metrics” – Requirement 10.

they are analyzed to determine which business requirements are pertinent to the cloud. These business requirements are examined to determine their relevance to security, portability, and interoperability needs, and whether they are mission-specific requirements or cross-cutting requirements. The final step is to determine the relationship of the business requirements to the SAJACC technical use cases.

A template to capture target business use cases was created and is described in the next section. An initial portfolio of target business use cases using this template was developed using two methods. The most common approach is documentation via interviews with agency and department CIOs identified through the Federal CIO Council Cloud Computing Executive Steering Committee and Cloud First Task Force. Information is gathered about the business use case through information provided by agencies, after which NIST-led interviews of key members of the cloud effort are conducted to flesh out the business use case and identify areas of concern. Alternatively, participants in the NIST-chaired public Cloud Computing Business Use Case Working Group (CCBUCWG) volunteer to document and obtain agency sponsorship of business use cases that might be of interest. Sponsoring federal agencies develop the business use cases and submit them to the project team as contributions. As business use cases are drafted, they are presented to the Cloud Computing Business Use Case Working Group for review and comment.

As requirements are identified and areas of research are prioritized, NIST works with federal agencies, industry, SDOs, and academia to identify options for addressing challenges, using the vendor-neutral reference architecture and taxonomy as a frame of reference. This research results in the definition of new or augmented standards, guidance, and technology requirements where appropriate. The portfolio of target business use cases can also be used by Federal CIOs to aid them in considering their projects. As federal CIOs identify new business use cases, it is helpful to the broader community to add them to this portfolio. As the portfolio of business use cases is expanded, trends and commonalities become more apparent, permitting prioritization of research areas.

3.1.1 Business Use Case Template

In order to identify common themes across business use cases, a template for documenting business use cases was created with input from the CCBUCWG. The template was designed to organize how the business use case was documented, ensure that the documenter articulated how the project met the NIST definition for cloud computing, and to encourage consideration of the various elements of the NIST Cloud Computing Reference Architecture.

The template consists of five major sections: description, background, cloud computing concept of operations, analysis, and concerns and challenges. The description is a brief, one-paragraph summary of the purpose and goals of the business use case. The background provides an explanation of how the business use case is currently solved, along with any definitions and descriptions needed to understand the business use case generally. The cloud computing concept of operations examines how a cloud implementation would work and identifies the key requirements that a cloud implementation would need to meet.

The analysis section incorporates the NIST definition of cloud computing and the reference architecture, leading the documenter to consider the service model, delivery model, the five essential characteristics, and the NIST focus areas of security, portability, and interoperability. Finally, any concerns or challenges expressed by the sponsor are captured.

3.1.2 Business Use Case Summaries

3.1.2.1 NIST IT Service Management

Delivery Model: Private Cloud

Service Model: SaaS

Agency: National Institute of Standards and Technology

FISMA² Impact Level: Moderate

NIST is interested in moving its service ticketing system to the cloud as part of a larger move to an IT Service Management model for providing services to end users. One of the main drivers for moving the trouble ticket system to the cloud is to allow IT to focus its resources on applications that directly implement functional aspects of the NIST mission. Moving non-core applications to the cloud eliminates the need to patch and update software and servers.

A longer-term goal of this implementation is to enable other service groups (such as telecommunications, security, and building maintenance) within NIST to use this tool as well. In this way, a single service request can be routed to appropriate service providers within NIST in a seamless way. The use of a cloud application would provide flexibility in the timing of deployments and the availability of system resources for testing and training.

3.1.2.2 US Census Virtual Desktop Infrastructure

Delivery Model: Private Cloud

Service Model: SaaS

Agency: United States Bureau of the Census

FISMA Impact Level: Moderate

² Federal Information Security Management Act (FISMA) of 2002.

The United States Bureau of the Census proposes to use cloud technology to comply with the Telework Enhancement Act of 2010 and to improve productivity by eliminating the need to use the SafeBoot device encryption tool. The benefits of this approach are in realizing a decreased cost of delivering computing and support services, creating a mobile workforce capable of using a variety of devices, and improving security by limiting the loss of sensitive data through the loss or theft of a mobile device or by malicious software. Specifically, the use of a Virtual Desktop Infrastructure (VDI) will reduce the high cost associated with providing and maintaining desktop service. The US Census expects to use a private cloud environment for its cloud effort.

Securing sensitive data is critical to enabling telework. By running virtual machines on a server and ensuring that all data resides on network storage, data can be properly secured. Finally, end-user compliance with security policies can be improved through managed personalization of the desktop environment.

The security infrastructure that enables single-sign-on and two-factor authentication is also an essential part of the solution and will be deployed in the same private cloud.

3.1.2.3 USAID Virtual Desktop Infrastructure (VDI)

Delivery Model: Community Cloud

Service Model: SaaS

Agency: US Agency for International Development

FISMA Impact Level: Moderate

USAID is interested in migrating to the cloud to provide IT services for its users distributed across the globe. The plan (in-progress) is to move email, office productivity, and some business applications into a cloud-based infrastructure and implement a cloud-based VDI to enable secure access to the services. This migration will decrease the cost of delivering computing and support services, create a mobile workforce that will use a variety of devices, and improve security by limiting the loss of sensitive data through the loss or theft of a mobile device or by malicious software. Specifically, the VDI will reduce the high cost associated with providing and maintaining desktop service, and by moving IT services into the cloud, help to reduce the need and the cost associated with developing and maintaining data centers. USAID expects a hybrid cloud environment that uses both private cloud and community cloud for its cloud effort. The security infrastructure that enables single-sign-on and two-factor authentication is also an essential part of the solution and will be deployed in the same private cloud.

3.1.2.4 USAID Office Productivity

Delivery Model: Community Cloud

Service Model: SaaS

Agency: US Agency for International Development

FISMA Impact Level: Moderate Internal, Low External

USAID OCIO plans to use Google Apps service for government to provide cloud-based email and document management service for USAID users. This service is expected to be deployed in an outsourced community cloud and accessed through the VDI or directly through the Internet. The other business applications are expected to be deployed in an on-site private cloud at the beginning and will later be migrated into an outsourced private cloud. These cloud-based applications will be accessed through the cloud-based VDI. The security infrastructure that enables single-sign-on, two-factor authentication, and identity management is an essential part of the solution and will be deployed in the same on-site private cloud.

3.1.2.5 FGDC Geospatial Cloud

Delivery Model: Community Cloud, Public Cloud

Service Model: PaaS

Agency: Federal Geospatial Data Committee

FISMA Impact Level: Moderate and Low, depending on need.

The Federal Geographic Data Committee and the General Services Administration (GSA) Cloud Computing Program Management Office operate the GeoCloud project on behalf of a wide range of federal agencies to explore the impact and possibilities of a geospatial computing-oriented cloud. The initiative seeks to define and investigate cloud savings, best practices, and lessons learned by migrating, benchmarking, and operating a set of ten existing public-access geospatial projects from six currently participating agencies –US Geologic Survey (USGS), National Oceanic and Atmospheric Administration (NOAA), Bureau of the Census, Environmental Protection Agency (EPA), Department of Agriculture (USDA), and Department of the Interior (DOI) with interest from the Department of Homeland Security (DHS).

The overall plan is to define, construct, and maintain a set of common geospatial platforms to support the project, using a joint agency platform model. Once platforms are in place and under maintenance, each project team will evaluate their application on its matching platform, document the steps needed to ensure security and performance, and track lessons learned along the way. To date, two platforms have been defined; one has been hardened and constructed and operates on Amazon's AWS public cloud. The project teams are beginning their exploration and sandbox phase to discover and document the processes needed to maintain these existing applications in the cloud.

Some agency geospatial applications, targeted for the public cloud, have data storage or processing needs that appear to make them more cost-effective in a community cloud setting. These applications will be piloted on similar shared platforms in a community facility housed in the US Geologic Survey.

3.1.2.6 NOAA Email

Delivery Model: Community Cloud

Service Model: SaaS

Agency: National Oceanic and Atmospheric Administration

FISMA Impact Level: Moderate

NOAA envisions using a cloud-based Unified Messaging Service (UMS) to replace NOAA's existing in-house-hosted email and calendaring systems and its installation of Blackberry Enterprise Server. The UMS would decrease system maintenance responsibilities for NOAA, and provide users with new features as they become available in the cloud-based solution. Additionally, NOAA expects to expand collaboration capabilities through increased use of integrated messaging and collaboration tools, and, optionally, to obtain archival and eDiscovery capabilities.

3.1.2.7 FAA eDiscovery

Delivery Model: Community Cloud

Service Model: SaaS

Agency: Federal Aviation Administration, Federal Cloud Computing Working Group

FISMA Impact Level: Moderate

The FAA, in conjunction with the Federal Cloud First Task Force and other federal agencies, is seeking a cloud-based eDiscovery solution, motivated by the agency's moving email to a cloud-based solution. This solution would be composed of an archive, identification and collection capability, data preservation

capability, and the processing and export of content. The objective is to implement a cloud-based eDiscovery solution that can analyze both in-house and cloud-based email systems because of the time that the project will take to migrate the FAA's email from in-house systems to the cloud. During the migration of email, the ability to respond to eDiscovery and FOIA requests is necessary.

3.1.2.8 In-depth Email

Delivery Model: All

Service Model: All

Agency: N/A

FISMA Impact Level: Moderate, Low

Currently available collaboration solutions tend to fall into one of two categories. The first category is a single client-based solution (e.g., Outlook/Exchange, Zimbra, Mobile.me) and provides a number of integrated functions within the client interface (e.g., email, calendar, address book, etc.). The second category is an amalgamation of a variety of separate tools, sometimes integrated within the mail client framework using plugins (e.g., Thunderbird supports a variety of calendar plug-ins).

In the majority of cases, Email/User Collaboration tools are services hosted within the organization and are designed to connect to user client systems. Web-based email, while a frequent functional offering is typically a casual use offering (leveraged when users travelling or it is inconvenient to access a work system). Despite its current low utilization, Web-based systems offer enhanced security and administrative controls. These solutions are pertinent to a Secure/Classified environment.

As laptops and 'Smart' mobile devices become more common, there is more pressure to make the user collaboration tools work within this extended usage paradigm. Ensuring that data and security models are adhered to in the mobile environment is critical.

3.1.3 Business Use Case Analysis

Mission requirements are extracted from the business use cases. Mission requirements are high-level requirements that must be met to successfully support the primary goals of the business use case. Cross-cutting system requirements which relate to security, portability, and interoperability are also identified. These system requirements are used to inform high-level strategic USG requirements in cloud adoption. Complementary tactical efforts, such as technical requirement analysis from the SAJACC effort and cloud security impediments analysis, benefit from these source requirements.

3.1.3.1 Mission Requirements

The portfolios of target business use cases help to identify the following mission requirements in USG agency migration to cloud computing:

Requirement		Description
1	eDiscovery	Meet eDiscovery requirements, identify electronic records meeting search criteria, and retrieve both the records and their metadata. Archives of responsive Electronically Stored Information (ESI) such as documents and spreadsheets should be portable among eDiscovery solutions. These ESI must retain metadata during migration between ESI-producing platforms.

2	FOIA	Meet the requirements of the Freedom of Information Act (FOIA) for identifying and responding to records requests. As with eDiscovery, archives of responsive ESI must be portable between eDiscovery solutions, and metadata should be retained when migrating from one ESI-producing platform to another.
3	Email	Move agency email services to the cloud to provide improved operating efficiency, in some cases consolidating several different email installations into a single cloud-based solution.
4	Workforce Mobility	Provide mobile access to all IT services, enabling secure access from any device and any place where there is sufficient network bandwidth.
5	Collaboration	Enable secure sharing and authoring of documents with partners, including nongovernmental organizations (NGOs) and foreign governments. The purpose is to allow the creation of common workspaces either within the agency, across agencies, or with partners of agencies on a project-by-project basis.
6	Common Geospatial Platform	Provide agencies with the ability to create and deploy geospatial applications rapidly and efficiently.
7	Security Audit Information Collection	Enable the capture, identification, and mitigation of security events. Security audit information needs to be captured at both a high level for monitoring purposes and at a level of detail sufficient to allow forensic analysis of any security incidents that occur. Furthermore, it is necessary to retain the information for a time sufficient to meet the forensic analysis needs of the cloud service procured.
8	Telework Enhancement Act Compliance	Provide secure telework options to employees. While the Workforce Mobility mission requirement is concerned with enabling appropriate IT services to be accessed from anywhere on any device, this mission requirement applies to allowing employees to work from home, providing agencies with greater control over data and security.
9	Provisioning, Monitoring, Trouble Ticketing Integration	Enable integration of IT support and monitoring tools for both traditional systems and cloud-based systems. Provisioning users through a common interface is necessary to avoid increased maintenance burdens as the number of cloud systems an agency has subscribed to increases. Trouble ticket management and visibility would encounter similar problems as the number of systems increases.

Table 1: Mission Requirements from Target Business Use Cases

3.1.3.2 Mapping Mission Requirements to Business Use Cases

The analysis of the business use cases begins with the identification of mission requirements that are distilled from a closer look at the primary goals of each business use case. They address not only what the business use case is trying to achieve, but also those elements deemed particularly important. The table below shows how different mission requirements can be traced to specific targeted business use cases.

		Business Use Cases							
		NIST ITSM	Census VDI	USAID VDI	USAID Office Productivity	FGDC Geospatial	NOAA Email	FAA eDiscovery	In-depth Email
Mission Requirements	eDiscovery							X	
	FOIA							X	
	Email						X		X
	Workforce Mobility			X	X		X	X	
	Collaboration				X				
	Common Geospatial Platform					X			
	Security Audit Information Collection	X							
	Telework Enhancement Act Compliance		X						
	Provisioning, Monitoring, Trouble Ticketing Integration	X			X		X	X	X

Table 2: Mapping Mission Requirements to Business Requirements

The next step of this analysis is construction of a matrix to correlate mission requirements to system requirements. System requirements are composed of requirements classified as cross-cutting elements, necessary in different cloud adoption scenarios and consequently considered an evolving product of business use case analysis. System requirements are critical in order for the mission requirements to be fully realized within the framework of the USG Cloud Computing Technology Roadmap. Cross-cutting system requirements can be broken down further into the generalized categories of security, interoperability, and portability.

Throughout the process of capturing mission requirements in each use case and decomposing them into system requirements, the roadmap priorities for USG cloud computing adoption are reassessed. BUC

mission and cross-cutting system requirements are instrumental in determining the highest priorities to further USG Cloud Computing Technology Adoption. Preliminary analysis has prompted and paved the way for further work to:

- Identify and provide solutions for high-priority security requirements (Requirement 1, Volume I);
- Develop frameworks to support federated community cloud, (Requirement 5, Volume I);
- Define unique government regulatory requirements, technology gaps, solutions (Requirement 7, Volume I);
- Identify the collaborative parallel strategic “future cloud” development initiatives (Requirement 8, Volume I);
- Define and implement reliability design goals (Requirement 9, Volume I); and
- Define and implement cloud service metrics (Requirement 10, Volume I).

The following sections provide illustrative examples that originate in the targeted business use cases for each category: security, interoperability, and portability.

3.1.3.3 Cross-cutting Security System Requirements

Security system requirements include those that pertain to information security. These include protection of information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction to help ensure integrity, confidentiality, and availability.

Requirement/Details		Description
1	<u>Identity Management</u>	A means of integrating Identity Management in the cloud with the cloud consumer’s Identity Management solution is necessary. Agencies that participated in the collection of business use cases typically require that a user be authenticated by the agency network, at which time access to cloud applications is provided. Cloud-based applications should be integrated into an identity management framework to avoid separate management of user identities in the cloud.
	<i>Single Sign-On (SSO)</i>	Upon authentication through the cloud consumer’s identity management solution, users should be able to access all cloud services without further authentication. Analysis of the use cases shows that systems with needs to migrate to the cloud tend to be integrated with a single sign-on (SSO) infrastructure. To prevent the loss of current functionality, the ability to integrate with an agency’s SSO solution is necessary.
	<i>Strong Authentication</i>	Most of the analyzed business use cases were for applications that were considered to be of a FISMA impact level of moderate, necessitating the use of strong authentication. Cloud providers will need to provide strong authentication to support systems with a FISMA impact level of moderate, such as two-factor authentication techniques using disconnected tokens or Homeland Security

		Presidential Directive (HSPD)12-compliant Common Access Cards,.
	<i>User Provisioning</i>	Cloud providers need to deliver standards-based APIs to allow the provisioning of users, either individually or in bulk. As the number of cloud services to which a cloud consumer is subscribed increases, the time spent on user maintenance will rapidly increase without the availability of interfaces that allow user management to be automated.
	<i>Access Policy Management</i>	A standard policy management interface is needed to permit creation, deletion, and maintenance of access policies from a standardized management tool. Well-maintained policies are a necessity for maintaining secure systems. As the number of cloud services to which a cloud consumer has subscribed increases, maintenance of access policies across cloud services becomes difficult without a standard interface to permit the use of a standard management tool.
2	<u>Security Audit Information</u>	Security audit data must be maintained for every aspect of the cloud service, as defined by contract and dependent on the impact level of the service, for use in the analysis of security incidents when they are discovered. High-level summaries of security audit information provide enough information to determine when an event took place, and detailed logs provide the information needed to perform a forensic analysis of the incident.
	<i>Availability of High-level Security Audit Data</i>	High-level security data must be captured and transferred to the cloud consumer on a regular basis, as defined within the contract. Capturing security audit information is required by Federal Information Processing Standards (FIPS) 200. These data are used to analyze security events of interest. While this information is readily available in traditional environments, the multi-tenant nature of cloud services requires additional cooperation between the cloud provider and the cloud consumer.
	<i>Availability of Detailed Security Audit Data</i>	Detailed security data must be captured and stored by the cloud provider so that forensic analysis of security breaches can be undertaken in cooperation with the cloud consumer. The need to capture detailed security audit information at the level required to carry out a forensic analysis is required by FIPS 200. The multi-tenant nature of cloud services requires cooperation between the cloud provider and the cloud consumer(s) affected by a security incident.
	<i>Security Audit Data</i>	Both high-level and detailed security audit data must be provided in a standards-based format so that cloud consumers could analyze the

	<i>Format and Exchange</i>	data. These data would be transferred at intervals defined in the contract. FIPS 200 requires that security audit information be captured and used for analysis of security incidents.
	<i>Security Audit Data Retention</i>	The cloud provider shall retain security audit data per cloud consumer requirements. FIPS 200 requires that security audit information is to be retained for a period of time sufficient to perform incident investigation in the event of a security breach. A cloud consumer needs to be notified of all security breaches that occur within the systems providing the cloud service.
	<i>Security Audit Data Monitoring</i>	The cloud provider needs to monitor security audit data with the frequency needed to rapidly identify and respond to security incidents, and notify the customer promptly in the event of a security breach. In addition to security breaches arising in contracted cloud services or in traditional systems operated by the cloud consumer, the multi-tenant nature of cloud services means that security incidents may originate with another customer at that cloud provider.
3	<u>Encryption</u>	Encryption is required for systems that are assigned a FISMA impact level of moderate or above. Most of the business use cases have been identified as systems that have a FISMA impact level of moderate. FISMA requires encryption of data, both at rest and in transit, to meet security requirements of moderate and above systems. In this way, even if devices are lost or stolen or transmissions intercepted, data remains protected.
	<i>Encryption of Data at Rest</i>	Systems at FISMA moderate or higher shall encrypt data using FIPS 140-2-validated encryption modules. Keys must be managed separately from data and require higher privileges. Encryption keys shall be changed on a regular basis, decrypting data and re-encrypting with the new key. To protect mobile devices from loss or theft, FISMA requires that data be encrypted if any of the systems on the mobile device have an impact rating of moderate.
	<i>Encryption of Data in Transit</i>	Encryption of data in transit is a FISMA requirement for moderate impact systems. This encryption protects data, including usernames and passwords, from interception. This is especially important when using untrusted network environments, such as open wireless access points at coffee shops, or public computer terminals in a library.
	<i>Multi-tenant Encryption</i>	Where encryption keys are required, the cloud provider must provide a FIPS 140-2-validated encryption algorithm for cloud consumers to establish their own encryption keys rather than the encryption keys. The cloud consumer remains responsible for

		establishing the encryption key whether or not the cloud provider is acting as a cloud broker. In the multi-tenant environment of cloud systems, not only does data need to be protected from other cloud consumers but from the cloud provider as well.
4	<u>Physical Security</u>	FISMA security standards not only apply to security protocols implementable using hardware or software, but also to the physical security of the facilities used to house the equipment and services. Physical security includes all measures whose purpose is to prevent physical access to a building, resource, or stored information. These physical security requirements apply to third parties engaged by cloud brokers.
	<i>Inspection of Premises</i>	The cloud provider shall make all facilities involved in providing the cloud service available for inspection by the cloud consumer or the cloud auditor, as required by FISMA. Cloud service implementations using third parties to provide some aspect of a service must allow inspection of the third party premises. This permits the evaluation of the physical security to meet FISMA moderate impact security requirements.
	<i>Physical Data Center Location</i>	The cloud provider shall limit the facilities in which the cloud consumer's data reside to the continental United States when requested. Limiting the physical data center location simplifies meeting FISMA moderate requirements as international travel by inspectors is not required, nor is understanding local laws regarding data ownership, privacy, and security necessary. The decreased visibility into data center locations with cloud implementations is a concern to US agencies.
5	<u>Assessment and Authorization</u>	The cloud provider shall work with the cloud consumer to obtain certification that the service being provided meets the requirements for the stated FISMA data classification. The Assessment and Authorization (formerly known as the Certification and Authorization) process is a security review and an approval to operate a system or to interact with other government systems. Cloud-based systems introduce additional complexity because the responsibility for documenting security controls is shared between the cloud provider and the cloud consumer.

Table 3: Cross-cutting Security Requirements

3.1.3.4 Cross-cutting Interoperability System Requirements

Interoperability relates to communication and data transfer between different systems. System requirements related to interoperability reflect the desire of federal agencies to automate processes between systems to the greatest degree possible. Interoperability decreases the need for manual intervention or providing the same information to multiple systems.

Requirement/Details		Description
1	<u>eDiscovery and FOIA</u>	eDiscovery interfaces shall be standard for both cloud and non-cloud systems. eDiscovery requests do not differentiate between cloud-based and traditional systems; both sources must be searched for responsive ESI. In order to avoid multiple interfaces, depending on which application or cloud service was obtained, standards are necessary to enable a single interface. The capability of capturing this information is more complex in cloud-based systems.
	<i>eDiscovery Search</i>	The ability to search various messaging, document repositories, and application databases for eDiscovery and FOIA purposes must be provided, including the search of metadata. The ability to search all sources needs to be independent of whether the solution being searched is in the cloud or directly managed. Due to the multi-tenant nature of cloud services, this capability is currently immature.
2	<u>Integrated Mobile Device Support</u>	The cloud provider shall provide support for heterogeneous clients, including mobile devices, thin and zero clients, Web clients, and thick clients. The option to allow the use of the different devices shall be configurable through a standard policy management interface. A single interface used to configure all devices eliminates the need to swap between programs when configuring different devices.
3	<u>Email Integration in Cloud Services</u>	The cloud provider shall provide a means of integrating application email capabilities with the email systems of the cloud consumer. There should be no need to separately define users within the cloud application; the appropriate information should be received through the bulk provisioning interface. Ensuring that email is appropriately configured and relayed provides the cloud consumer with the traceability required for complying with eDiscovery laws and regulations.
4	<u>Help Desk and Trouble Ticketing Management</u>	The cloud provider shall provide a means of integrating application email capabilities with the email systems of the cloud consumer. There should be no need to separately define users within the cloud application; the appropriate information should be received through the bulk provisioning interface. Ensuring that email is appropriately configured and relayed provides the cloud consumer with the traceability required for complying with eDiscovery laws and regulations.
	<i>Interface for Opening and Routing Trouble</i>	The cloud provider shall provide a standard interface for opening trouble tickets, enabling cloud consumers to open trouble tickets using automated tools or to route trouble tickets from any general

	<i>Tickets</i>	ticketing solutions that the cloud consumer may be using. Complexity is decreased for a cloud consumer using multiple cloud services if there is a single point for the creation, update, and monitoring of trouble tickets.
	<i>Interface for Notification of Ticket Updates and Status Changes</i>	The cloud provider shall provide a standard interface for receiving updates on tickets that are not closed so that automated tools or general ticketing solutions could be updated. Cloud consumers that have automated reporting of problems and outages through their ticketing systems need to integrate cloud provider ticketing with their systems.
	<i>Ticket Interface to Email</i>	The cloud provider shall allow the cloud consumer to update trouble tickets using email for those individuals without access to a primary interface. Agencies that provide the ability to email problem reports that automatically open tickets have been identified.
	<i>Interface for Event Management System Opening and Update of Tickets</i>	The cloud provider shall notify the cloud consumer's event management system when appropriate through a standard interface, updating status as appropriate. Monitoring of all system event information through a single interface is necessary for a unified view of important events throughout all applications that are used by the cloud consumer. Moving a particular system to the cloud does not remove the responsibility of the cloud consumer to monitor and understand events in their systems.
5	<u>Collaboration Standards</u>	Standard document formats are needed for portability and interoperability. Metadata such as privileges, creation and modification dates, etc., are needed to ensure that privileges, traceability, and information needed to meet eDiscovery requirements are retained. Many agencies have documents that are stored in old or obsolete formats. The ability to convert these documents to more recent formats while retaining all metadata is critical to allow these documents to be ported to the cloud.
	<i>Document Migration Path</i>	The cloud provider shall provide the ability to bulk convert files, including metadata, from old or obsolete formats to current formats. When implementing a collaboration solution in the cloud, agencies must be able to migrate from old or obsolete file formats to current file formats. Metadata need to be retained for eDiscovery and security purposes. The use of cloud services for office productivity solutions increases the frequency and complexity of changing providers.
	<i>External Collaboration</i>	The cloud provider shall provide a means for cloud subscriber users to not only collaborate internally, but also to collaborate with

		external partners. The sharing of documents in a secure and compliant way with external organizations is frequently cited as a requirement for a collaboration solution.
6	<u>Billing and Reporting Interoperability</u>	Billing and usage reporting should be standardized across systems to enable cloud consumers to make meaningful comparisons of costs and benefits across multiple cloud implementations.
7	<u>VM Management Interoperability</u>	Virtual machine management interoperability is required so that platforms running in services provided by multiple cloud providers can be stopped, started, terminated, and maintained using a single interface.

Table 4: Cross-cutting Interoperability Requirements

3.1.3.5 Cross-cutting Portability System Requirements

Portability system requirements identify needs for moving data between systems. Portability needs arise when upgrading software or when migrating between two competing systems. Ending a contract for a cloud service, whether by the cloud consumer or the cloud provider, results in additional considerations, such as what must occur with data held by the cloud provider.

Requirement/Details		Description
1	<u>Email Data Portability</u>	Standards for moving email data must include metadata for purposes of eDiscovery. Existing consensus-based standards for email, calendaring, contacts, tasks, and notes should be fully supported to ensure portability between different vendors. Retention of metadata when moving email between different implementations or providers needs to be supported. As not all standards for email are fully supported by all vendors, the complexity of migrations is increased.
	<i>Data Export</i>	The cloud provider shall provide a method for exporting email, calendar entries, tasks, notes, contacts, and saved instant messages to a standard format, retaining initial and current metadata. Export to fully supported standard formats simplifies migrations and enhances data portability. Retention of initial and current metadata allows agencies to more easily meet eDiscovery regulations.
	<i>Data Import</i>	The cloud provider shall provide a method for importing email, calendar entries, tasks, notes, contacts, and saved instant messages from a standard format, retaining initial and current metadata. Support for standard formats increases the portability of standard email capabilities across vendors. Retaining metadata during the import process enables compliance with federal eDiscovery requirements.

2	<u>Data Deletion</u>	Ensuring that data are completely deleted decreases the likelihood of security breaches in the future, and ensures that federal agencies are meeting security and privacy statutes. Traditionally, the owner of the data is responsible for the hardware on which data were stored and backups made, and ensured that data were destroyed prior to disposal of hardware. In the cloud, the cloud consumer must rely on the cloud provider to ensure deletion of data from all appropriate components (such as hard disks and tapes).
	<i>Deletion of Business Data</i>	At the termination of a contract, the cloud provider must return all business data to the cloud consumer, and ensure that the data are irrevocably deleted from all of their systems. Ensuring deletion of all data at the termination of a contract ensures that the cloud provider does not have any future obligation to the cloud consumer. The cloud consumer does not need to worry about potential security or privacy breaches at their former cloud provider.
	<i>Deletion of Logs, Usage Data, and Audit Data</i>	At the termination of a contract, the cloud provider must delete all usage data from all services that could be traced back to an agency or user. This information could provide useful information to third parties about usage patterns and implementation that the cloud consumer may not want released. In a traditional implementation, the agency was able to directly control data and its use; in a cloud implementation, the accountability remains but the direct control is lost.
	<i>Code Escrow</i>	In the event of a cloud provider exiting or de-supporting a cloud solution, to support the ability to set up this solution to another cloud so that the solution can be used or migrated in the future, the cloud provider shall put a copy of all of the source code required to re-create the system in escrow. Federal cloud consumers must meet statutory data retention requirements. Additionally, it is incumbent upon federal cloud consumers to ensure continuity of operations. The ability to rapidly re-create the environment if a cloud provider is no longer able to provide access to an appropriate environment and version of the system is needed.
3	<u>Portability for eDiscovery and FOIA Purposes</u>	Federal agencies must meet various statutes regarding eDiscovery and FOIA that are in place today. In order to meet eDiscovery obligations, metadata need to be retained even as the underlying ESI are migrated from one vendor to another. It is easier to retain metadata in a traditional environment as more operations retain the information than when switching cloud vendors.
	<i>Portability of Responsive Electronically Stored</i>	For ESI deemed responsive to be portable, it is necessary to ensure that information regarding implemented litigation holds and whether a specific record was deemed responsive to one or more

	<i>Information</i>	searches is retained. The ESI themselves must be exportable in formats defined in discovery or FOIA case law. The cloud environment differs in that retention of historic information is likely to require migration.
	<i>Portability of Metadata Required for eDiscovery and FOIA</i>	The migration of ESI shall retain metadata as per consensus-based standards. Standards ensure that discovery tools provide agencies with the ability to extract metadata from ESI in a manner consistent with eDiscovery and FOIA requirements across applications or systems. The need to rely on cloud providers having appropriate metadata necessitates the use of standards.
	<i>Export of Electronically Stored Information for eDiscovery and FOIA</i>	The cloud provider shall provide the ability for eDiscovery tools to produce ESI deemed to be responsive in standard formats, such as native, tiff, jpg, and pdf. The format in which responsive ESI is provided to requesting parties is determined through negotiation. Supporting multiple formats for export of ESI is necessary to produce what is expected to the requesting parties.
4	<u>Portability of Virtual Desktops</u>	The ability to move virtual desktops between vendors and cloud providers must be provided. Virtual desktops are not currently portable between vendors. Once a cloud consumer makes a decision to virtualize the desktop environment, the virtualization stack is very difficult to migrate to a different implementation.
	<i>Moving Virtual Desktops Between Vendors</i>	The cloud provider shall implement a standard format for virtual desktops. A standard format based on consensus-based standards allows virtual desktops to be moved seamlessly from one implementation to another.
	<i>Migration of Virtual Desktops</i>	The cloud provider shall use standard interfaces that assign, start, and stop virtual desktops. Migration of a virtual desktop should not require additional configuration on the part of the cloud consumer's administrators to allow the user of the desktop to use the desktop in the new environment. Agencies have thousands of users, and configuration changes would make migrations very difficult and time-consuming.
	<i>Accessibility of Virtual Desktops from Heterogeneous Devices</i>	The cloud provider shall make virtual desktops accessible via any device, including mobile devices, pads, thin and zero clients, and standard fat clients. Enabling access of virtual desktops from any device would significantly increase the mobility of the cloud consumer's workforce. Cloud consumers use virtual desktops not only for increased control over the desktop, but also to provide their users with the increased accessibility through mobile computing.
	<i>Virtualization of</i>	The cloud provider shall provide a means for virtualizing legacy

	<i>Legacy Software</i>	software packages. Legacy software is a significant problem for many cloud consumers. In many agencies, there may be a lot of legacy applications used by only a few people each that, if virtualized, would allow better support and monitoring. Virtualizing these legacy applications removes the dependency on aging hardware platforms and enables organizations to continue to offer the utility of this software on modernized computing infrastructure.
5	<u>Portability of Virtual Machines</u>	Static virtual machine portability is required so that the maintained platform images can be freely migrated between cloud implementations without the need for parallel development or maintenance.

Table 5: Cross Cutting Portability Requirements

3.1.3.6 Mapping System Requirements to Mission Requirements

The table below shows the system requirements and which mission requirements provided the genesis for each. The same system requirement could arise from one or more mission requirements.

		Mission Requirements								
		eDiscovery	FOIA	Email	Workforce Mobility	Collaboration	Common Geospatial Platform	Security Data Collection	Telework Compliance	Monitoring and Ticketing
Security Requirements	Identity Management	X	X	X	X	X	X		X	X
	Security Audit Information							X		
	Encryption	X	X	X	X	X			X	X
	Physical Security									X
	Assessment and Authorization			X	X	X	X		X	X
Interoperability Requirements	eDiscovery and FOIA	X	X							
	Integrated Mobile Device Support									X
	Email Integration in Cloud Services				X					X
	Help Desk and Trouble Ticketing Management									X
	Collaboration Standards					X				
	Billing and Reporting Interoperability						X			
	VM Management Interoperability						X			

Portability Requirements	Email Data Portability	X	X	X						
	Data Deletion									X
	Portability for eDiscovery and FOIA Purposes	X	X			X				
	Portability of Virtual Desktops				X				X	
	Portability of Virtual Machines						X			

Table 6: Mapping System Requirements to Mission Requirements

3.2 SAJACC Use Cases and Technical Requirements

The Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC) project focuses on cloud consumers' technical requirements to generate concrete data about how different kinds of cloud

Highlights: SAJACC refers to a tactical project, process, and portal.

The SAJACC project develops technical requirements, and identifies and defines and supports a process and portal that can be used to test system interfaces that meet or partially meet these requirements.

The results of the tests are analyzed to capture portability, interoperability, or security implications.

This section presents rationale and support for NIST USG Cloud Computing Technology Roadmap Volume I, High-Priority Requirements:

- “International voluntary consensus-based interoperability, portability, and security standards” - Requirement 1, and
- “Solutions for High-priority Security Requirements” - Requirement 2.

system interfaces can support portability, interoperability, and security. By showing worked examples, the SAJACC project seeks to facilitate SDOs in their efforts to develop high-quality standards that address these important needs.

Since its inception in May 2010, SAJACC has evolved to be an operational process and portal which includes iteratively:

- developing a set of cloud system use cases that express selected portability, interoperability, and security concerns that cloud users may have;
- selecting a small set of existing cloud system interfaces that can be used for testing purposes;
- developing a test driver, for each use case and selected system interface, that represents (to the extent possible) the operation of the use case on the selected system interface;
- running the test drivers and documenting the extent each test driver can run on each selected system interface; and documenting any portability, interoperability, or security implications of the test run; and
- publishing all use cases, test codes, and test

results on the openly accessible NIST Cloud Computing Collaboration Portal, for use by SDOs and other interested parties.

The set of technical use cases developed by the SAJACC project describes how groups of users and their resources may interact with one or more cloud computing systems to achieve specific goals. Each of the goals expressed in the use cases are usually a small atomic unit of work. This use case methodology has been widely used in software and system engineering as a tool to express technical requirements. It

describes actors (who are involved) and goals (what to achieve), success scenarios (how to achieve the goals), failure conditions, and failure handling.

The process of documenting cloud computing technical requirements using SAJACC use cases is ongoing; however, the first set of published SAJACC use cases includes three categories: management, interoperability, and security, as shown in Table 7 below.

Management	Interoperability	Security
• Open An Account • Close An Account • Terminate An Account • Copy Data Objects Into a Cloud • Copy Data Objects Out of a Cloud • Erase Data Objects In a Cloud • Allocate VM Instance • Manage Virtual Machine Instance State • Query Cloud-Provider Capabilities and Capacities	• Copy Data Objects between Cloud-Providers • Dynamic Operation Dispatch to IaaS Clouds • Cloud Burst From Data Center to Cloud • Migrate a Queuing-Based Application • Migrate (fully-stopped) VMs from one cloud-provider to another	• User Account Provisioning • User Authentication in the Cloud • Data Access Authorization Policy Management in the Cloud • User Credential Synchronization Between Enterprises and the Cloud • eDiscovery • Security Monitoring • Sharing of Access to Data in a Cloud

Table 7: SAJACC Use Cases

Through an open process, the SAJACC project has also collected and generated a catalog of system interfaces that can be used to address the technical requirements expressed in these use cases. The SAJACC project has developed a generic testing framework and implemented test drivers for an initial set of use cases using identified system interfaces. This testing mechanism has demonstrated how cloud consumers' technical requirements can be implemented using existing public interfaces and also helped to surface issues and gaps in existing system interfaces. The set of use cases, test drivers, and testing results provide concrete data to support the development of high-quality standards to address portability, interoperability, and security concerns expressed by the consumers. The SAJACC project plans to continue to develop technical use cases and update existing ones with community input. The project will also continue to develop demonstrable test drivers to show how existing system interfaces can be used to implement requirements, and identify issues and gaps to feed ongoing standardization efforts.

4 Cloud Computing Standards and Gap Analysis

Cloud Computing owes its existence to a sizable collection of standards that have been developed to facilitate communication, data exchange, and security. As Cloud Computing gains momentum, many other standards are emerging to focus on technologies that support cloud computing, such as virtualization. SDOs and others are developing cloud computing conceptual models, standards roadmaps, use cases, etc. The NIST Cloud Computing Standards Roadmap Working Group is leveraging this existing, publicly available work, plus the work of the other NIST working groups, to identify standards, standards gaps, and standardization priorities.

Highlights: To support US government requirements for interoperability, portability, and security in cloud computing, the NIST public Cloud Computing Standards Roadmap Working Group has surveyed the existing standards landscape for security, portability, and interoperability standards/models /studies/use cases, etc., relevant to cloud computing.

An inventory of Cloud Computing Relevant Standards has been compiled, and only three emerging cloud standards have been identified to date.

The findings confirm the need for these: USG Cloud Computing Technology Roadmap Volume I high-priority requirements:

- “International Voluntary Consensus based Interoperability, Portability & Security Standards” – Requirement 1, and
- “Solutions for high priority security requirements” – Requirement 2.

See NIST Special Publication 500-291, *NIST Cloud Computing Standards Roadmap*.

As identified in Volume I of the Technology Roadmap, standards will play an important role in cloud computing, particularly in interoperability, portability and security. The analysis of cloud computing standards, and resulting gaps, is closely correlated to the entire cloud strategy:

- The standards, as listed in Section 4.1, are aligned to and categorized by the NIST conceptual model and reference architecture as referenced in Section 2;
- The use cases in Section 3 and the revealed USG cloud computing requirements provided references in prioritization on the standards gaps are listed in Section 4.2.
- Recommendations for accelerating the development and use of cloud computing standards, presented in Section 4.3, are in accordance with the Priority Action Plans presented in Volume I of the Technology Roadmap.

4.1 Cloud Computing Standards

Standards are already available in support of many of the functions and requirements for cloud computing portability, interoperability, and security. While many of these standards were developed for pre-cloud computing technologies, such as those designed for Web services and the Internet, they can

also support the functions and requirements of cloud computing. Other standards are now being developed in specific support of cloud computing functions and requirements, such as virtualization.

To assess the state of standardization in support of cloud computing, the NIST Cloud Computing Standards Roadmap Working Group has compiled an Inventory of Standards Relevant to Cloud Computing (URL: <http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/StandardsInventory>). This inventory is being maintained and the information could be used to update the *NIST Cloud Computing Standards Roadmap* accordingly.

Using the taxonomy developed by the NIST Cloud Computing Reference Architecture and Taxonomy Working Group, cloud computing relevant standards have been mapped to the requirements of portability,

interoperability, and security. The *NIST Cloud Computing Standards Roadmap*, First Edition, NIST SP 500-291, includes a mapping of cloud computing standards. The status of the standard maturity level will be closely monitored and updated.

4.2 Cloud Computing Standards Gaps and USG Priorities

There are emerging challenges in some areas of cloud computing that have been addressed by technology vendors and service providers' unique innovations. New service model interactions and the distributed nature in resource control and ownership in cloud computing have resulted in new standards gaps. Additionally, standardization gaps from some pre-cloud computing era gaps are being brought to the forefront by cloud computing. Areas of standardization gaps are identified from examining cloud computing standards.

As described in the Federal Cloud Computing Strategy, cloud computing business use cases have various priorities. The requirements expressed in these high-priority target business use cases can be used to prioritize the standardization gaps. For example, various USG groups have identified data center consolidation using virtualization technologies as one of the primary goals in the next few years. Migrating collaboration applications, including email messaging (email, contact and calendars) and online office productivity applications, to the cloud is also an early target of government cloud operation.

Table 8 summarizes the areas of standardization gaps and standardization priorities based on USG cloud computing adoption requirements. The NIST cloud computing reference architecture is used as the framework of reference to identify these gaps in need of standardization, and then a broad set of USG business use cases are used to identify the priorities of standardization that will maximize the benefits and meet the more urgent needs of government consumers.

Area of Standardization Gaps	Priorities for Standardization Based On USG Requirements
SaaS Functional Interfaces, e.g., • Data format and interface standards for email and office productivity • Metadata format and interface standards for eDiscovery	High priorities on: • SaaS application specific data and metadata format standards to support interoperability and portability requirement when migrating high-value, low-risk applications to SaaS.
SaaS Self-Service Management Interfaces, e.g., • Interface standards related to user account and credential management	<i>(n.b. this requirement is not yet identified as a priority by a specific USG target business use case)</i>
PaaS Functional Interfaces, e.g., • Standards of data format to support database serialization and de-serialization	<i>(n.b. this requirement is not yet identified as a priority by a specific USG target business use case)</i>
Business Support, Provisioning and Configuration, e.g., • Standards for describing cloud Service-Level agreement and quality of services • Standards for describing and discovering cloud service resources • Standards for metering and billing of service consumptions and usage.	High priorities on: • Resource description and discovery standards to support data center consolidation using private and community IaaS clouds

Security and Privacy, e.g., • Standards for identity provisioning and management across different network and administration domains • Standards for secure and efficient replication of identity and access policy information across systems • Single sign-on interface and protocol standards that support strong authentication • Standards in policies, processes, and technical controls in supporting the security auditing, regulation, and law compliance needs	High priorities on: • Security auditing and compliance standards to support secure deployment, assess, and accreditation process for cloud specific deployment • Identity and access management standards to support secure integration of cloud systems into existing enterprise security infrastructure
---	--

Table 8: Cloud Computing Standards Gaps and USG Priorities

4.3 Accelerating the Development and the Use of Cloud Computing Standards

There is a fast-changing landscape of cloud computing relevant standardization under way in a number of SDOs. While there are only a few approved cloud computing specific standards at present, federal agencies are encouraged to participate in specific cloud computing standards development projects that support their service priorities. Specific recommendations for government agencies are:

Recommendation 1 – Contribute Agency Requirements

Agencies should contribute clear and comprehensive user requirements for cloud computing standards projects.

Recommendation 2 – Participate in Standards Development

Agencies should actively participate in cloud computing standards development projects that are of high priority to their agency missions.

Recommendation 3 – Encourage Testing to Accelerate Technically Sound Standards-Based Deployments

Agencies should support the concurrent development of conformity and interoperability assessment schemes to accelerate the development and use of technically sound cloud computing standards and standards-based products, processes, and services.

Recommendation 4 – Specify Cloud Computing Standards

Agencies should specify cloud computing standards as a factor in procuring cloud services and assess cases when multiple vendors offer standards-based implementations and there is evidence of successful interoperability testing. In such cases, agencies should ask vendors to show compliance to the specified standards.

Recommendation 5 – USG-Wide Use of Cloud Computing Standards

To support USG government requirements for interoperability, portability, and security in cloud computing, in coordination with and under the cognizance of the federal Enterprise Architecture program, the Federal Standards and Technology Working Group should recommend specific cloud computing standards for USG-wide use.

Recommendation 6 – Dissemination of Information on Cloud Computing Standards

A listing of standards relevant to cloud computing should be posted and maintained.

5 High-Priority Security Requirements

Industry surveys and polls consistently show that security, privacy, and compliance are among the greatest concerns of organizations considering adopting cloud solutions. For USG agencies, such concerns are often heightened due to the sensitivity of information being handled and the gravity of the consequences of failing to protect such information. Indeed, cloud computing characteristics do bring unique security challenges such as:

Federal managers are sensitive to challenging security requirements that may become obstacles to the adoption of cloud computing, and the need to understand and consider possible mitigations.

The Security Requirements list reported in this section was produced by the NIST-led public Cloud Computing Security Working Group, and reviewed with the Federal Cloud Computing Standards and Technology Working Group, and other interagency stakeholders.

This section presents rationale that supports the NIST USG Cloud Computing Technology Roadmap Volume I high-priority requirements:

- “Solutions for High-priority Security Requirements” - Requirement 2;
- “Technical specifications to enable development of consistent, high-quality SLAs” - Requirement 3;
- “Technical security solutions which are decoupled from organizational policy decisions”- Requirement 6;
- “Defined unique government regulatory requirements, technology gaps, and solutions”- Requirement 7;
- Defined and implemented reliability design goals” – Requirement 9; and
- Defined and implemented cloud service metrics” – Requirement 10.

See also NIST Special Publication 800-144: *Guidelines on Security and Privacy in Public Cloud Computing*, and NIST Special Publication 800-146: *Cloud Computing Synopsis and Recommendations*.

- Broad network access, a prerequisite for moving IT assets into the cloud, has the potential to introduce new cyber threats;
- The (perceived) lack of visibility and control over the IT assets often runs counter to the existing security policies and practices that assume complete organizational ownership and physical security boundaries;
- Multi-tenancy is prevalent in real-world cloud solutions and a source of concern related to segmentation, isolation, and incident response.

Such challenges, however, are not insurmountable. The key to secure cloud computing lies in understanding the security requirements in the particular cloud architectural contexts and mapping them to proper security controls and practices in technical, operational, and management dimensions. In addition, cloud computing brings new benefits to security architectures and solutions, resulting in services that could be made more robust and resilient. For example:

- Well-defined resource abstraction layers (infrastructure, platform, and software apps) bring more architectural flexibility, allowing for application of more effective security countermeasures at each layer, resulting in better “defense in depth” compared with traditional, rigid security controls relying on physical attributes (such as specific devices, MAC addresses, etc.).
- A cloud provider may achieve better “economies of scale” in applying security improvements to many consumers. For example, a new control designed to remedy one consumer’s vulnerability may be more quickly applied for all consumers.

5.1 Understanding Security in the Cloud Context

Though constantly facing new threats and incorporating new technological advances, network and information security is generally a well-understood and well-researched domain with a rich body of knowledge both in theory and in practice. Cloud-based services can leverage existing analyses of security architectures to address security requirements such as authentication, authorization, availability, confidentiality, identity management, integrity, audit, continuous monitoring, incident response, and security policy management.

However, while these security requirements are not new, they need to be analyzed using cloud-specific perspectives and characteristics. One approach is to leverage the Cloud Computing Reference Architecture to better understand how and why security needs to be looked at differently in the cloud, using the cloud model definition and perspectives.

5.1.1 Cloud Service Model Perspectives

The three service models identified by the NIST cloud computing definition, i.e., SaaS, PaaS, and IaaS, present consumers with different types of service management operations and expose different entry points into cloud systems, which in turn also create different attack surfaces for adversaries. Hence, it is important to consider the impact of cloud service models and their different issues in security design and implementation. For example, SaaS provides users with accessibility of cloud offerings using a network connection, normally over the Internet and through a Web browser. There has been an emphasis on Web browser security in SaaS cloud system security considerations. Cloud consumers of IaaS are provided with virtual machines (VMs) that are executed on hypervisors on the hosts; therefore, hypervisor security for achieving VM isolation has been studied extensively for IaaS cloud providers that use virtualization technologies.

5.1.2 Implications of Cloud Deployment Models

One way to look at the security implications from the deployment model perspective is the differing level of exclusivity of tenants in a deployment model. A private cloud is dedicated to one consumer organization, whereas a public cloud could have unpredictable tenants coexisting with each other; therefore, workload isolation is less of a security concern in a private cloud than in a public cloud. Another way to analyze the security impact of cloud deployment models is to use the concept of access boundaries. For example, an on-site private cloud may or may not need additional boundary controllers at the cloud boundary when the private cloud is hosted on-site within the cloud consumer organization's network boundary, whereas an out-sourced private cloud tends to require the establishment of such perimeter protection at the boundary of the cloud.

5.1.3 Shared Security Responsibilities

The cloud provider and the cloud consumer have differing degrees of control over the computing resources in a cloud system. Compared to traditional IT systems, where one organization has control over the whole stack of computing resources and the entire life cycle of the systems, cloud providers and cloud consumers collaboratively design, build, deploy, and operate cloud-based systems. The split of control means both parties now share the responsibilities in providing adequate protections to the cloud-based systems. Security is a shared responsibility. Security controls, i.e., measures used to provide protections, need to be analyzed to determine which party is in a better position to implement these controls. This analysis needs to include considerations from a service model perspective, where different service models imply different degrees of control between cloud providers and cloud consumers. For example, account management controls for initial system privileged users in IaaS scenarios are typically performed by the

IaaS Provider whereas application user account management for the application deployed in an IaaS environment is typically not the provider’s responsibility.

5.1.4 Developing Security Architecture for Cloud Systems

As shown in previous sections, many other factors will affect the security in the cloud. The NIST Cloud Computing Security Working Group will continue to work on guidelines to support a framework for developing cloud security architecture for cloud systems.

5.2 Challenging Security Requirements and Risk Mitigations

Given the landscape of rapidly changing cloud industry solutions and emerging cloud security standards, it is premature to provide a definitive, overarching architecture framework, and implementation guidance for cloud security. As part of the roadmap initiative, the NIST Cloud Computing Security Working Group has taken the first step in identifying a list of likely security impediments to cloud adoption, and the available strategies for mitigating the risks inherent to the selected security requirements.

The NIST security requirements and risk mitigations list describes the security issues that the NIST Cloud Computing Security Working Group has identified as challenging for the cloud computing adopters, and provides, when available, strategies for mitigating their effects.




Figure 5: Challenging Security Requirements to Mitigation Mapping

Figure 5 illustrates the approach. For each identified requirement, there is a brief textual description of the nature of the challenge created by the unsatisfied requirement and, when available, a set of mitigations that can address or facilitate solutions for this challenge. Each mitigation may briefly describe a strategy for mitigating the security requirement, it may point to other existing work where the security requirement is addressed, or both.

This document, Volume 2 of the *USG Cloud Computing Technology Roadmap*, provides a high-level summary of requirement challenges and mitigations. It is not intended to serve the purpose of detailed security guidance. More detailed security guidance exists in the form of special publications which are referenced in this document and the NIST *Challenging Security Requirements for USG Cloud Computing Adoption* which is being developed in an open collaborative process through the working group. The working document is available through the working group Web site: <http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/CloudSecurity>.

The following list of security requirements and mitigations is grouped in two categories: Process-Oriented Requirements and Focused Technical Requirements. The following two sections summarize the contents of the requirements and mitigations.

5.3 Process-Oriented Requirements

The following requirements rely primarily upon human-centered processes, procedures, and guidance for risk mitigation.

5.3.1 Application of NIST SP 800-53-style Controls and Compliance

Description: The requirement addresses the need for clarity in how NIST SP 800-53 security and privacy controls can be applied in cloud-based information systems.

Importance: Federal system owners must ensure that systems processing federal data are assessed and authorized to operate. Migration of systems or services to the cloud environment does not affect the authorizing official's responsibility and authority.

Mitigation: NIST Risk Management Framework

FISMA and Office of Management and Budget (OMB) policy require cloud service providers handling federal information or operating information systems on behalf of the federal government to meet the same security and privacy requirements as federal agencies. Security and privacy requirements for cloud service providers including the security and privacy controls for information systems processing, storing, or transmitting federal information are expressed in appropriate contracts or other formal agreements using the Risk Management Framework and associated NIST security standards and guidelines.

Organizations can require cloud service providers to implement all steps in the Risk Management Framework described in NIST SP 800-37 with the exception of the security authorization (to operate) step, which remains an inherent federal responsibility that is directly linked to the management of risk related to the use of all IT services, including cloud services.

Organizations determine the security category of the information that will be processed, stored, or transmitted within the cloud-based information system in accordance with FIPS Publication 199. This security categorization drives the selection of appropriate security and privacy controls that will be required to be implemented by cloud service providers. Since many security and privacy controls have shared responsibility for implementation depending on the cloud service model chosen (e.g., IaaS, PaaS, SaaS), organizations should provide in their contracts and Service-Level Agreements with cloud service providers, the specific allocation of those responsibilities.

Organizations should also ensure that the assessment of required security and privacy controls is carried out by qualified independent, third-party assessment organizations that are able to assert if the cloud service providers deliver appropriate evidence of control effectiveness. This evidence is used by organizations to make initial authorization decisions. Organizations should also develop a continuous monitoring strategy and ensure that the strategy is implemented by the cloud service provider including defining how the security and privacy controls will be monitored over time (e.g., frequency of monitoring activities, rigor and extent of monitoring activities, and the data feeds provided to the organization from the cloud service provider). The continuous monitoring data feeds will be used by the organization for ongoing authorization decisions as part of its enterprise-wide risk management program.

The assurance or confidence that the risk from using cloud services is at an acceptable level depends on the trust that the organization places in the external service provider. In some cases, the level of trust is based on the amount of direct control the organization is able to exert on the cloud service provider with regard to employment of security and privacy controls necessary for the protection of federal information and the cloud service as well as the evidence brought forth as to the effectiveness of those controls. The level of control is usually established by the terms and conditions of the contract or Service-Level Agreement with the cloud service provider (e.g., negotiating a contract or agreement that specifies detailed security and privacy controls for the provider).

The Federal Risk and Authorization Management Program (FedRAMP) is being implemented by the Federal CIO Council and GSA in order to reduce the compliance burden for agencies and suppliers in terms of time and cost, while still satisfying the requirements described above. This includes defining minimum security and privacy requirements for cloud-based information systems. FedRAMP has identified a set of requirements that must be in place to satisfy security and privacy controls from NIST SP 800-53 as defined for low- and moderate-impact information processed, stored, and transmitted within cloud-based information systems delivering cloud services. Continuous monitoring controls are also defined. A conformity assessment program will provide opportunities to obtain independent, third-party assessment services to determine security and privacy control effectiveness. FedRAMP also follows the NIST Risk Management Framework as described in NIST SP 800-37.

References: NIST SP 800-53 (as amended), NIST SP 800-37 (as amended), FedRAMP URLs.

5.3.2 Cloud Audit Assurance and Log Sensitivity Management

Description: Mechanisms to gain assurance that:

- Important events are monitored;
- Sensitive/private audit logs are appropriately protected;
- Integrity of audit data used for initial or continuous auditing purposes, e.g., audit logs, data collected by Security Content Automation Protocol (SCAP), etc., is protected; and
- Audit data interchange incompatibility is addressed.

Technical Considerations: The cloud model introduces another party, the Cloud Service Provider's auditor, into an organization's computing model. This fact introduces important questions about monitoring and auditing requirements:

- Who is doing any particular monitoring or auditing task?
- Who is informed of the results of a particular monitoring or auditing task, and when?
- What is an appropriate level of abstraction and summarization in the aforementioned results?

It is important to note the distinction between monitoring and reporting. This requirement addresses the monitoring task and how the results from that activity such as raw log data or aggregated reports are handled. Section 5.9 of this document discusses the reporting requirements and guidance aimed at standardizing the reporting function. Monitoring a system for anomalies is in the purview of the system operator. Monitoring will produce results that can be compiled in a report and delivered to other stakeholders of the system.

Cloud providers may be required to store and/or forward log data to designated collection points or aggregation storage media. Whichever option for the handling of system log data is chosen, in order to assure the data is secure, steps must be taken to protect the data in transit and at rest. There is any number of methods for deployment of encryption to protect the data while ensuring it can be accessed when requested. Data may be forwarded to external entities for automated inspection. An IPSec-like encryption method may provide the best performance but may not be suitable for highly mobile data scenarios.

Practical Example: Operational requirements for the monitoring or auditing of cloud environments can vary significantly depending on many factors. For example, a private cloud restricted to limited physical locations may not be as inherently mobile as a public cloud where data may be relocated more dynamically. In such a private cloud scenario, monitoring sensors could be deployed without the concerns of iterative relocating or modifying of sensors. In a public cloud, multi-tenancy concerns could emerge depending on the characteristics of the data monitored and/or captured. If those data are moved

dynamically, providers and subscribers may face challenges in ensuring that subscribers are able to monitor and receive reports specific to their data.

In a public cloud scenario, the provider has operational control of the environment and may offer a baseline of monitoring services. SLAs or contracts should be used to ensure that specific requirements for monitoring and metrics are satisfied. In any SLA or contract with the cloud service provider, the customer should specify measurable monitoring and reporting standards. The contract should specify the measures to be taken if any SLA requirements are not met. The requirement for a periodic review of the SLAs and their parameters should be defined in the contract. Monitoring tasks also do not absolve the customer of responsibility to monitor and audit aspects of the information system that the customer operates or manage.

Importance: Cloud Auditing and Continuous monitoring is identified as a requirement for all federal systems.

Solution Maturity: Immature. While effective monitoring solutions have been in use for some time, the high mobility inherent to the cloud computing environment and multi-tenancy provide unique challenges in the implementation of mechanisms to monitor specific data.

Mitigation 1: Risk management framework

The NIST Risk Management Framework (RMF) (SP 800-37) provides guidance to federal system owners to take a risk-based approach to securing systems. This approach is operationally focused and is intended to facilitate the monitoring, documenting, and mitigation of threats on a regular if not near real-time basis. Continuous monitoring is step 6 of SP 800-37's 6-step risk management framework. While many vendors are seeking to offer automated vulnerability monitoring tools, it is important to realize that there is more to an effective continuous monitoring program than automated tools.

The FedRAMP program's Proposed Security Assessment and Authorization document (<https://info.apps.gov/sites/default/files/Proposed-Security-Assessment-and-Authorization-for-Cloud-Computing.pdf>) describes an effective continuous monitoring program as one that includes:

- Configuration management and control processes for information systems;
- Security impact analyses on proposed or actual changes to information systems and environments of operation;
- Assessment of selected security controls (including system-specific, hybrid, and common controls) based on the defined continuous monitoring strategy;
- Security status reporting to appropriate officials; and
- Active involvement by authorizing officials in the ongoing management of information system-related security risks.

It is important to note that there is a distinction between the continuous monitoring controls requirements identified in FedRAMP controls set, currently implemented mechanisms to perform continuous monitoring functions, and target or future continuous monitoring solutions and standards which are being defined and developed. They are not one and the same, although the current continuous monitoring mechanisms and future continuous monitoring solutions may be applied to satisfy the FedRAMP controls requirements.

Sufficiency Comment: The RMF and 800-53 provide adequate guidance and controls related to the securing of audit data.

Mitigation 2: Audit Data Interchange

The Cybersecurity Information Exchange Techniques (CYBEX) project was launched by the International Telecommunication Union - Telecommunication Standardization Sector (ITU-T). Cybex provides for the structured exchange at known assurance levels of information about the measurable "security state" of systems and devices, about vulnerabilities, about incidents such as cyber attacks, and about related knowledge "heuristics." The CYBEX initiative imports more than 20 "best of breed" standards for platforms developed over the past several years by government agencies and industry to enhance cyber security and infrastructure protection. Pulling these platforms together in a coherent way provides for:

- "Locking down" on-line systems to minimize vulnerabilities;
- Capturing incident information for subsequent analysis when harmful incidents occur; and
- Discovering and exchanging related information with some degree of assurance.

The CYBEX Model includes:

- Architecting cyber security information to support exchange;
- Identifying and discovering cyber security information and entities;
- Establishing trust and policy agreement between exchanging entities;
- Requesting and responding with cyber security; and
- Assuring the integrity of the cyber security information exchange.

Real-time Inter-network Defense (RID) [RFC6045, RFC6046] provides a proactive inter-network communication method to facilitate sharing incident handling data while integrating existing detection, tracing, source identification, and mitigation mechanisms for a complete incident handling solution. Organizations have a need for RID and related standards in cloud computing to communicate quickly and efficiently with their providers on incident information. The escalation points from detection to investigation and mitigation may vary based on SLAs, but the transfer of the information must be standardized (globally) to enable the use of various vendor platforms for the secure and standardized exchange of incident information. The incident information may be exchanged for the purpose of situational awareness or be for an investigation that is associated with a request to mitigate or stop the incident. Incidents may also be benign and require quick reporting and mitigation methods. Examples include configuration issues or availability issues caused by operations problems. These incidents may also be communicated via the described protocols.

References:

- CSA Cloud Audit - <http://cloudataudit.org/page5/page5.html>
- CSA/ CSC - Cloud Trust Protocol - http://assets1.csc.com/lef/downloads/Digital_Trust_in_the_Cloud.pdf
- The FedRAMP document: <https://info.apps.gov/sites/default/files/Proposed-Security-Assessment-and-Authorization-for-Cloud-Computing.pdf>
- NIST 800-53 AU9 – Protection of audit Information
- PCI DSS 10.5.5 – File Integrity Monitoring
- ISO27001 10.10.3 – Protection of Log Information
- NIST SP 800-92 - *Guide to Computer Security Log Management*
- CSA CCM SA-14 – Audit Logging / Intrusion Detection
- CYBEX Overview - http://www.itu.int/dms_pub/itu-t/oth/0A/0D/T0A0D00001D0004PDFE.pdf

5.3.3 Cloud Certification and Accreditation Guidelines

Description: How to certify and accredit cloud solutions with confidence.

Importance: USG departments and agencies, to effectively manage information security risks inherent in all modern computing technologies, must have a high degree of trust and confidence in the entities providing new and innovative technologies, including cloud technologies and services.

Mitigation: FedRAMP was initiated to provide a cost-effective, risk-based approach for the assessment and authorization of federal cloud services. Establishing clear and concise expectations for security and privacy based on current threats, taking advantage of innovative, open, and state-of-the-practice solutions for the protection of federal information in cloud-based information systems, and ensuring a high degree of transparency in security and privacy solutions, will promote a climate of trust between consumers and providers of cloud services.

References:

- <http://www.fedramp.gov>

5.3.4 Clear eDiscovery Guidelines

Description: Mechanism to provide access to data in response to lawful authority while protecting customer privacy. Mechanism to ensure service providers are preserving electronic records with sufficient evidential weight and chain of custody controls.

Importance: Meeting electronic discovery requests can pose a challenge when electronically stored information (ESI) is in the cloud.

Mitigation 1: When procuring a cloud service, customers must gain an understanding of how the cloud provider processes electronic discovery and litigation holds. The customer should acquire knowledge of key issues – such as the length of time the provider takes to enforce a litigation hold (i.e., prevent the modification and/or destruction of pertinent evidence) or respond to an electronic discovery request and what steps are required to invoke these processes, types of logs and metadata retained including life cycles of same, dependencies on other providers, evidentiary chain of custody and storage, and additional processing fees that may be incurred. Having a subject-matter expert discuss these processes with the cloud provider is preferable to a checklist, due to the variances of cloud environments and the specialized knowledge requirements around electronic discovery and preservation of evidence. Specific wording or clauses may need to be inserted into the cloud contract to ensure that cloud providers share the burden for failure to properly secure and maintain evidence once a hold or request has been properly initiated.

Mitigation 2: Customers should undertake the effort to map significant business processes and ESI created, processed, and/or stored as a result that would have a high likelihood of being the target of an electronic discovery request. Where possible, the proactive collection, indexing, and storage of ESI that has a reasonable expectancy of falling within the scope of future litigation or discovery requests (such as email) may lessen the dependency on cloud providers – particularly if the ESI can be stored on systems under the direct control of the customer. A records retention policy defining the forms of ESI routinely collected and archived, as well as ESI formats not retained, can assist in refining the scope of this effort.

Mitigation 3: Providers should undertake the effort to understand the requirements for lawful intercept, National Security Letters, Subpoena, and eDiscovery. Providers must make a timely response and provide

information for a specific tenant without collateral information from other tenants. Providers must be able to locate and provide access to data or communication channels that are specific to a single tenant.

References:

- Federal Rules of Civil Procedure (2010).

5.3.5 Cloud Privacy Guidelines

Description: This requirement addresses the need to build predictability and confidence in the degree to which cloud solutions provide privacy data and Personally Identifiable Information (PII) protection.

Importance: The Privacy Act of 1974, 5 U.S.C. § 552a As Amended (<http://www.justice.gov/opcl/privstat.htm>) and The Computer Matching and Privacy Protection Act of 1988 (http://www.irs.gov/irm/part11/irm_11-003-039.html) require the protection of personal information held by agencies. Additionally, in the commercial arena, the FTC's Fair Information Practices have established a framework under which individuals can depend upon certain privacy-related rights and expectations when engaging in business transactions with both online and brick-and-mortar merchant entities. (<http://www.ftc.gov/reports/privacy3/fairinfo.shtm>) The OMB Memorandum M03-22 established the guidance for federal agencies to implement the E-Government Act of 2002. This guidance provided for individual agencies to develop Privacy Impact Assessments (PIAs) to enable them to understand the privacy implications of the data that they were managing within their systems and to ensure that the proper controls were in place to protect the data according to established law.

Mitigation 1: Establish and maintain the confidence of those for and about whom federal agencies manage personal data. Cloud Customers (federal agencies) should, in the case of cloud services as in the case of other computing models, consistently assess the scope of the Personally Identifiable Information that they manage within services for which they are responsible. This requires the application of PIA processes in order to determine the degree of risk associated with the type of data that is being maintained. For instance, health information (under the Health Insurance Portability Accountability Act and Health Information Technology for Economic and Clinical Health [HITECH requirements]) needs to be assessed in the context of the public, hybrid public/private, community and private cloud models at all service levels.

Mitigation 2: Ensure that Cloud Providers protect the personal information to the requisite levels of protection a) that have been established for all of the federal agencies' systems, and b) are finalized to the degree necessary to define cloud-specific controls. Service-Level Agreements and other legal instruments need to be established between the Cloud Customer and the Cloud Provider, given that the Cloud Customer is still responsible for the protection of the data.

Mitigation 3: Ensure that cross-jurisdictional Privacy issues are addressed and incorporated in agencies' cloud deployments if the data that will be collected, managed, retained, or otherwise processed falls under the scope of global Data Protection regulations.

References:

General Privacy Laws Governing Federal Agencies

- Privacy Act of 1974 <http://www.justice.gov/opcl/privstat.htm>
- E-Government Act of 2002 http://frWebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=107_cong_public_laws&docid=f:publ347.107.pdf

OMB Privacy Guidance and Policies

- Privacy Act Implementation, Guidelines and Responsibilities
http://www.whitehouse.gov/sites/default/files/omb/assets/omb/inforeg/implementation_guidelines.pdf
- OMB Circular No. A-130, Management of Federal Information Resources
http://www.whitehouse.gov/omb/circulars_a130_a130trans4
- OMB Memorandum M-99-18, Privacy Policies on Federal Web Sites
http://www.whitehouse.gov/omb/memoranda_m99-18
- OMB Memorandum M-03-22, OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002 http://www.whitehouse.gov/omb/memoranda_m03-22
- OMB Memorandum M-06-15, Safeguarding Personally Identifiable Information
<http://www.whitehouse.gov/sites/default/files/omb/assets/omb/memoranda/fy2006/m-06-15.pdf>
- OMB Memorandum M-07-16, Safeguarding Against and Responding to the Breach of Personally Identifiable Information
<http://www.whitehouse.gov/sites/default/files/omb/assets/omb/memoranda/fy2007/m07-16.pdf>
- OMB Memorandum M-10-22, Guidance for Online Use of Web Measurement and Customization Technologies http://www.whitehouse.gov/sites/default/files/omb/assets/memoranda_2010/m10-22.pdf
- OMB Memorandum M-10-23, Guidance for Agency Use of Third-Party Websites and Applications
http://www.whitehouse.gov/sites/default/files/omb/assets/memoranda_2010/m10-23.pdf
- Other OMB Guidelines Additional Guidance from OMB regarding Privacy Regulations
http://www.whitehouse.gov/omb/inforeg_infopoltech#prm

Department of Justice

- DOJ Privacy Act Regulations, “Protection of Privacy and Access to Individual Records Under the Privacy Act of 1974,” 28 C.F.R. pt. 16 subpart D. <http://www.justice.gov/opcl/regulations.htm>
- DOJ Privacy Act Regulations, “Exemption of Records Systems Under the Privacy Act,” 28 C.F.R. pt. 16 subpart E. http://www.access.gpo.gov/nara/cfr/waisidx_10/28cfr16_10.html
- Incident Response Procedures for Data Breaches Involving Personally Identifiable Information
<http://www.justice.gov/opcl/breach-procedures.pdf>
- DOJ Overview of Privacy Act <http://www.justice.gov/opcl/1974privacyact-overview.htm>

Department of Homeland Security

- http://www.dhs.gov/xlibrary/assets/privacy/privacy_pia_guidance_may2007.pdf

U.S. Security and Exchange Commission

- <http://www.sec.gov/about/privacy/piaguide.pdf>

FDIC

- <http://fcx.fdic.gov/about/privacy/assessments.html>

Department of Education

- <http://www2.ed.gov/notices/pia/index.html>

Department of Defense

- <http://www.dtic.mil/whs/directives/corres/pdf/540016p.pdf>

5.3.6 Clarified Security Control Roles and Responsibilities

Description: Mechanism to define who (among cloud actors such as customer and provider) is responsible for the implementation of required security controls. Intuitively, it seems that the actor most able to observe and configure a specific portion of a cloud implementation would be the best positioned to implement a relevant control.

Importance: The data owner (cloud customer) is responsible for compliance with laws and regulations including the proper security controls for their data, regardless of its location or the involvement of other parties. The data owner's ability to implement security controls is often limited when customer data is off-premise and under the control of a third party. Cloud providers/brokers/carriers have increasing responsibilities for implementing and maintaining security depending on the cloud deployment and service models.

Mitigation 1: Provider-subscriber guidelines

Guidance that documents roles and responsibilities definitions for cloud provider and consumer helps provide the required clarity. Such guidance can be used in specifying the responsibilities for protection in contract terms between a system owner and a cloud provider.

Mitigation 2: Cloud type/service selection

In cases where a larger degree of direct control over security roles/responsibilities and the ability to implement security controls is needed, cloud customers may consider utilization of a service type and/or a deployment type which will allow that requirement to be fulfilled.

References:

- CSA Cloud Controls Matrix, which included controls from frameworks such as: ISO 27001/27002, ISACA COBIT, PCI, NIST 800-143, Jericho Forum and NERC CIP
- NIST Special Publication 800-146, Cloud Computing Synopsis and Recommendations

5.3.7 Trustworthiness of Cloud Operators

Description: This requirement addresses the need to ensure that individuals with physical and logical access to subscriber data are properly vetted and screened periodically to ensure trustworthiness.

Importance: For cloud service consumers, it is critical to be able to confirm the security practices of their service providers' operations. This is necessary to maintain and improve the security posture of their data and operational services. Cloud consumers need to know and understand what cloud providers are doing and if they are effectively performing those functions. In addition, cloud consumers must be able to randomly and independently verify their cloud service providers' practices.

Mitigation 1: Operator human resources practices

Through standardized, consistent SLAs of high quality and completeness, consumers can specify requirements such as background screening requirements for operator staff, require regular training to ensure that operator employees (including contractors and third-party users) understand responsibilities related to specific consumer requirements, and apply best practices. It is also reasonable for consumers

and operators to define and confirm application of separation of duties and processes to monitor unauthorized activities by malicious insiders.

Mitigation 2: Operator self-certification and third party verification

To gain consumers' trust, cloud operators may pursue self-certification of compliance with legal and regulatory requirements (consistent with SAS 70 or ISO 27002 compliant certification systems). Third-party independent audit of operators' information security management can be applied to policies and specific management and technical controls.

Mitigation 3: Operator transparency

Consumers need to trust and verify that cloud operators offer the appropriate level of security and governance for their data and applications. Operator transparency implies a commitment to communicate security information (policies, practices and incident responses) to consumers and to advise them as to risks and risk mitigations.

Mitigation 4: Improved knowledge base through reviews of services provided by government, consumer, and industry groups

References:

- FedRAMP repository of authorized cloud providers (<http://www.fedramp.gov>).
- Reviews and insights into the cloud hosting companies (<http://www.cloud-hosting-providers.com/>).
- List of cloud servers (<http://www.bestcloudserver.com/>).
- List of cloud hosting providers (<http://www.cloudhostingreviewer.com/>).

5.3.8 Business Continuity and Disaster Recovery

Description: In traditional IT operations, business continuity planning (more specifically, contingency planning) is complex, and the effectiveness of its implementation is difficult to test and verify. More often than not, when disasters occur, unexpected disruptions create confusion and result in less efficient recovery practices. Cloud computing increases complexity to the IT infrastructure and obfuscates responsibility between cloud provider and customer. This elevates the level of concern related to business continuity and disaster recovery in a new paradigm such as cloud computing.

Importance: Identifying an effective Contingency and Disaster Recovery Plan is imperative to securing information systems and is a required deliverable of the Risk Management Framework and Certification and Accreditation Process.

Mitigation 1: Consistent policies and procedures, as in the case of all IT services. This includes taking action to:

- Develop a contingency plan for a cloud-based application or system using guidelines in NIST SP 800-34 Rev 1 and in Domain 9: Contingency Planning, Federal Cloud Security Guidelines (if published);
- Determine ownership, data sensitivity, cloud service and deployment models, roles and responsibilities;
- Specify Recovery Point Objective (RPO) and Recovery Time Objective (RTO);
- Set recovery priorities and map resource requirements accordingly;
- Provide a road map of actions for activation, notification, recovery procedures, and reconstitution;
- Enforce policies and procedures through SLAs;

- Incorporate the customer contingency plan into cloud provider's overall contingency plan;
- Establish management succession and escalation procedures between cloud provider and customer; and
- Reduce the complexity of the recovery effort.

Mitigation 2: Ensure that requirements traditionally met through the following clustering and redundancy mechanisms are addressed:

- Shared storage clusters;
- Hardware-level clustering;
- VM clusters; and
- Software clustering (application servers and database management systems).

Mitigation 3: Ensure requirements met traditionally through alternate sites and backup are addressed. NIST SP 800-53 Rev3 recommends:

- Alternate storage and processing sites;
- Alternate telecommunication services;
- Information system backup;
- Provide cold, warm and hot backup sites (economies of scale);
- Outsource information system backup to a cloud backup service;
- Use multiple cloud providers; and
- Supplement cloud provider's backup schemes with customer's non-cloud sites.

Mitigation 4: Ensure effective testing and exercises are conducted. This includes exercising the contingency plan periodically to verify its effectiveness (including personnel training) and confirming that it is updated to reflect changes in any of the dependent factors.

The service provider and consumer should plan to perform joint contingency plan testing and exercises against high-level disruptions to discover deep-rooted risks.

The service provider and consumer should plan to perform joint testing in business and service provider production-like environments to exercise contingency plans.

References:

- NIST SP 800-34 Rev 1: Contingency Planning Guide for Federal Information Systems
- NIST SP 800-53 Rev 3: Recommended Security Controls for Federal Information Systems and Organizations
- NIST SP 800-144: DRAFT Guidelines on Security and Privacy in Public Cloud Computing
- Federal Cloud Security Guidelines (2011)

5.4 Focused Technical Requirements

This section describes potential security impediments and risk mitigations, where the focus is on technical mechanisms rather than human processes.

5.4.1 Technical Continuous Monitoring Capabilities

Description: The assessment is that there are insufficient technical continuous monitoring capabilities to the extent necessary to support monitoring of cloud environments. This requirement is especially challenging in the case of multi-data center clouds which use many different security tools. The audit data from diverse security tools must be normalized and aggregated to provide situational awareness to support low-level security operations. This data then needs to be further aggregated to provide the perspective needed to support high-level operational mission assessments and management decisions. The data needs to reflect the security posture of the cloud as well as the security posture of customer's mission supported by the cloud services.

Practical Example: Questions exist regarding how specific information can be obtained and obsessed related to the security posture of an environment in which a subscribers' data may reside. Existing monitoring solutions were not designed for highly mobile environments or multi-tenant environments with potentially largely disparate monitoring and reporting requirements.

Importance: Cloud providers must be able to gain situational awareness of their cloud environment and to provide evidence to their customers that the cloud infrastructure is secure. Also important is the ability to provide customer feedback on the security posture related to their use of cloud services.

Solution Maturity: Much of the foundation for addressing this requirement exists in the subject area of security automation standards. This is especially true for asset, configuration, and vulnerability management. However, the higher-level model needed to provide situational awareness is still immature.

Mitigation 1: The CAESARS Framework Extension effort (under development). This joint NIST, NSA, and DHS effort is planned to provide a reference model for data normalization, aggregation, and situational awareness. In the short term, the effort is focused on binding to the Security Content Automation Protocol in order to provide continuous monitoring capabilities for asset, configuration, and vulnerability management.

CyberScope is designed to be a secure Web-based application that collects automated and manual data from federal agencies, used to assess and report the agencies' IT security posture. CyberScope receives live data feeds and that provided through data entry by agency staff. CyberScope is designed as a central repository, accessible by agencies through a standard interface and format. Through this interface, agencies provide data to the OMB, which then compiles and generates reports to other agencies, as required by the FISMA.³

The information that OMB requires to be reported through CyberScope is broader in scope than the status of individual assets. The latter is the focus of the CAESARS reference architecture. Nevertheless, the CAESARS reference architecture can directly support the achievement of some continuous monitoring objectives by ensuring that the inventory, configuration, and vulnerabilities of systems, services, hardware, and software are consistent, accurate, and complete. A fundamental underpinning of both the CAESARS reference architecture and the CyberScope reporting objectives is full situational awareness of all agency IT assets.⁴

³ <https://www.cippguide.org/2010/11/02/cyberscope/>

⁴ <http://www.dhs.gov/xlibrary/assets/fns-caesars.pdf>

Sufficiency Comment: When adopted and implemented, the CEASARS framework will allow agencies to implement CM more rapidly by leveraging CM-compliant tools, eliminating the need for custom integration efforts. This is envisioned to more effectively support the Cloud Computing paradigm.

References:

- CAESARS Framework Extension: A Continuous Monitoring Technical Reference Architecture, Draft NIST IR 7756, http://esrc.nist.gov/publications/drafts/nistir-7756/Draft-nistir-7756_feb2011.pdf
- NIST SP 800-137, Information Security Continuous Monitoring for Federal Information Systems and Organizations

5.4.2 Visibility for Customers

Description: Mechanism to define how cloud subscribers (customers) can observe their workloads to become aware of their security, compliance, privacy, health, and general status. Mechanism to determine how subscribers can instruct the cloud service provider regarding the information in which they are interested.

Importance: Cloud subscribers are ultimately liable for security, compliance, and privacy. Security/compliance/privacy regulations specify that that ultimate liability cannot be outsourced. Providers do not currently attempt to accept full responsibility through their SLAs.

Providers may compensate for the subscription cost of an outage, but not the actual damage or resulting loss of business.

Mitigation 1: Implementing Audit mechanisms

For example, CloudAudit.org is a Cloud Security Alliance standardization initiative that allows subscribers to tell providers what information they require and in what format. The maturity is unclear in terms of implementation.

Mitigation 2: Monitoring

Security Content Automation Protocol (SCAP) is an alert format standard specified to enable providers to provide alerts to subscribers in a standard format.

References:

- www.cloudaudit.org
- <http://scap.nist.gov>

5.4.3 Control for Customers

Description: The assessment is that customers have limited control over security policies enforced by cloud providers on their behalf. There is little automation available to help customers to implement technical controls (policies) in their applications which are deployed in cloud models. A mechanism is needed to allow cloud customers to maintain effective control over their workload, given that the protection mechanisms and the location of the workloads may not be known to them. The requirement is a mechanism that allows customers to communicate to the cloud provider regarding the security policies that are to be enforced at various control layers such as data object, VMs/Applications, virtual network, and geographic location.

Importance: Moving IT services to the cloud model necessitates some degree of ceding control over how information is protected and where it resides. It is important to identify information assets and control needs and to adopt cloud models accordingly. Customers and providers need to be able to define and enforce security policies at various control layers.

Mitigation 1: Selection and Use of Appropriate Cloud Models

Different service models and different deployment models affect the degree of customer control.

Mitigation 2: Control Data Objects

Access control over data objects is a widely used and mature function. Customers need to verify that providers protect data at rest, in transit, and when it is processed. Protection measures include:

- Establishing and maintaining data ownership;
- Using of authorization management standards/systems to specify and enforce access controls based on the attributes of the user and the data object, and the context of the access request;
- Maintaining change history records; and
- Managing the data life cycle.

Mitigation 3: Control of VMs and Applications

Consumers need to take steps to verify and cloud service providers need to:

- Perform and verify that VM hardening is implemented based on federal and generally accepted standards;
- Use automated tools to assess and report VM baseline security configurations and patch updates (including dormant and rolled back);
- Sanitize and protect virtual machine images; and
- Secure APIs (based on externalized, unified and fine-grained authorization management, for example) to allocate, start, stop and de-allocate VMs/applications.

Mitigation 4: Control Virtual Network

Consumers need to take steps to verify and cloud service providers need to:

- Apply protection mechanisms to intra-host virtual network (vSwitches/vLANs) that are similar to those applied to physical networks (for example, firewall, IDS, and antivirus); and
- Make virtual network traffic visible to physical network security and monitoring devices;

Mitigation 5: Control of Geographic Location

Consumers need to take steps to verify and cloud service providers need to:

- Consider contingency and availability criteria when identifying and select data center locations; and
- Enforce and verify security and compliance constraints for trans-border data flow in self-service, data replication, workload management, and cloud bursting.

References:

- www.modeldrivensecurity.org

- www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml

5.4.4 Data Protection Concerns

Description: The loss of confidentiality, integrity, or availability of customer data results in a variety of impacts. Cloud customers need to understand the extent of the data protection that a cloud offers (even if limited) in order to make rational risk-based decisions regarding cloud data storage and processing services.

FIPS 199 provides a categorization scheme (low-impact, moderate-impact, high-impact) for data and systems and describes the impacts in terms of confidentiality, integrity, and availability. The suitability of a cloud to store or process customer data varies depending on the data security impact level and on the extent that the cloud service provider can offer assurance that the data is protected. The technical ability to protect data varies depending on how the data is accessed. A number of access scenarios are possible, including:

- In transit to or from a provider: Data that a customer wishes to upload into a cloud must be protected in transit; similarly, data that a customer wishes to download from a cloud must be protected in transit;
- Passively stored with no shared access: Data should be accessed only by the originating customer and needs to be protected against access attempts by all other entities, while preserving the availability for the originating customer;
- Passive stored with selective shared access: Data should be accessed only by entities that have been authorized by the originating customer for specific access modes (e.g., read, write, delete) and needs to be protected against access attempts by unauthorized entities or accesses in unauthorized modes, while preserving availability for authorized customers;
- Passively stored public access: Data should be accessible anonymously in some authorized modes (e.g., read) but not accessed in other modes except by authorized customers;
- Actively processed: Data is accessed by a computation running in a cloud (e.g., a VM, PaaS, or SaaS application) but otherwise may not be shared or may be shared selectively;
- Account termination: Data should be maintained for a fixed period of time; and
- Deletion: There is authorized erasure of customer data.

Importance: High. If cloud services do not offer robust protection of customer data, migration to cloud computing will be limited to low-impact data and applications.

Mitigation 1: Consumers need to take steps to verify and cloud service providers need to implement data management measures to ensure the integrity and availability of information which is in transit, being processed, and in storage. Another consideration of cloud usage is data segregation and isolation, to address the risk that data may be comingled between organizations. Data encryption can be used to address the requirement of data confidentiality in various states. Data management measures include:

Data at rest:

- Those to prevent data tampering, copying, alteration, and deletion;
- Applying hashes or certificates to ensure authenticity; and
- Implementing method(s) to support search and to update encryption algorithms.

Data processing:

- Define the requirements for treatment of information which is processed within the cloud; and
- Implement processes to prevent data leakage.

Data in transit:

- Deploy remote VPN connections/Public ISP access;
- Assess properties of mobile wireless devices;
- Assess intranet, cross-agency or cross-department data transfer; and
- Directly encrypt data, using hashing/signatures.⁵

Mitigation 2: Consumers need to take steps to verify and cloud service providers need to employ a comprehensive Information Life Cycle Management Program to help assure the protection and proper handling of data throughout the various phases of data management. Cloud providers are responsible for managing some phases of the SDLC program but federal officials are ultimately responsible for ensuring that mechanisms for enforcement and oversight are in place and adhered to.

The Cloud Security Alliance has developed a useful model of information life cycle management, which defines the phases of Create, Store, Use, Share, Archive, and Destroy⁶, as shown in Figure 6. The security requirements in this life cycle are defined based on the types of data.



Figure 6: Information Life Cycle Management Phases

This simple model of Create, Store, Use, Share, Archive, and Destroy can use adapted security controls from NIST SP 800-64 and NIST SP 800-53Rev3. This is one example of a private sector model, which is useful for formulating additional pertinent controls.

References:

- <http://www.cloudsecurityalliance.org/csaguide.pdf>.
- Guidelines for Secure Use of Cloud Computing by Federal Departments and Agencies

5.4.5 Risk of Account Compromise

Description: A benefit of cloud computing is easy accessibility. A customer can use cloud computing services anywhere they have Internet access. However, Internet threats such as phishing, pharming, and spyware are designed to steal usernames and passwords (credentials). Given this Internet security threat environment, customers adopting cloud computing need to understand how user accounts are protected from hijacking and misuse.

Importance: Account hijacking is not new, but the concern is heightened in the context of cloud computing because:

⁵ Guidelines for Secure Use of Cloud Computing by Federal Departments and Agencies

⁶ <http://www.cloudsecurityalliance.org/csaguide.pdf>

- There is additional attack surface exposure due to increased complexity and dynamic infrastructure allocation;
- New APIs/interfaces are emerging that are untested; and
- The customer's account, if hijacked, may be used to steal information, manipulate data, and defraud others, or to attack other tenants as an insider in the multi-tenancy environment.

Mitigation 1: Consumers need to take steps to verify and cloud service providers need to implement strong authentication mechanisms, including:

- Enforcement of strong passwords and periodic password changes;
- Multifactor authentication;
- Prompts to require users to enter passwords during sessions, and in response to suspicious events;
- Use of a white-listed address range to constraint logins; and
- Multifactor authentication through biometrics.

Mitigation 2: Consumers need to take steps to verify and cloud service providers need to apply encryption to credentials and credential exchanges, including:

- Provision of a dedicated VPN;
- Use of HTTPS and LDAPS;
- Measures to enable secure cookies; and
- Use of strong cryptographic PKI keys.

Mitigation 3: Secure APIs/interfaces

Consumers need to take steps to verify and cloud service providers need to provide common security models for cloud APIs/interfaces (e.g., WS*, WS-I, SAML for Web services).

Consumers need to take steps to verify and cloud providers need to protect application security using secure APIs/interfaces (e.g., input validation/escaping/encoding against injection exploits such as SQL injection and cross-site scripting).

References:

- Symantec Internet Security Threat Report, Trends for 2010, Volume 16, April 2011

5.4.6 Identity and Access Management (IAM) and Authorization

Description: Unauthorized access to sensitive information in public, private, and hybrid clouds is a major security concern. Even though identity and access management (IAM) has long been used to manage users and their access to IT resources, there is a need to specify IAMs in terms of identity proofing, strength of credentials, and access control mechanisms for effective federal cloud-based authentication and authorization.

Importance: High. The identity and access management (IAM) needs to be effective and scalable, and considered in the context of multiple clouds. To achieve effectiveness and scalability, seamless extension of controls from agencies to the cloud is needed. Establishing trust relationships between cloud customers and cloud providers and potentially identity, credential, and attribute providers is key.

Mitigation 1: Consumers and cloud service providers need to specify use of the provider's IAM for cloud-based services and use of agency IAM for internal systems and functions.

There is a need to not only consider the effort in creating user identities and account provisioning.

Mitigation 2: Consumers and cloud service providers need to specify the degree and method of integrating the agency's IAM with cloud-based services.

For example, cloud providers may accept agency-created identity credentials, verify attributes of users and objects through accepted techniques and enforce authentication and authorization policies in a context-aware fashion.

Mitigation 3: Consumers and cloud service providers need to consider and specify claim-based Federated Identity Management

In this example, a single sign-on (SSO) solution that relies on an external identity system to provide cloud services with information about the user (claims) along with cryptographic assurance (a security token) that the identity data comes from a trusted source (an issuing authority). Cloud services can then make authentication and authorization decisions based on these supplied claims. There are many types of issuing authorities, from domain controllers that issue Kerberos tickets, to certificate authorities (CAs) that issue X.509 certificates.

Consumers and cloud service providers also need to consider and may specify use of unifying standards such as SAML to exchange authentication and authorization decisions between security domains (for example, identity providers and service providers).

Mitigation 4: Digital Identity

Consumers and cloud service providers also need to consider and may specify emerging user-centric technologies such as Information Cards (for federal agencies, PIV cards) or OpenID. Rather than centering on a directory (domain-centric), digital identity is focused around the user, enabling users to apply their digital IDs to use of cloud services, with on-the-spot validation (similar in concept to the way driver's licenses are used in the real world to establish the identity of individuals). This solution is consistent with the scalability and flexibility requirements to support use of multiple and various cloud services.

Mitigation 5: Standards-based Access Control

No matter what access control model (discretionary access control, mandatory access control, role-based access control, or attribute-based access control) is used, consumers and cloud service providers also need to consider emerging standards such as XACML to express and enforce confidentiality and integrity requirements in a flexible and unifying way for a variety of cloud environments. The flexibility allows an agency to specify and deploy access control policies to match its mixture of assets and portfolio of business functions, and to plug in additional policies as business and infrastructure evolve. The unity is designed to express access control policies in a single language and format to support use of multiple and various cloud services.

References:

- DHS Top Security Controls
- SAJACC Identity in the Cloud - Use Cases Version 1.0 OASIS
- SAJACC NIST Cloud Computing Use Cases
- Electronic Authentication Guideline. NIST Special Publication 800-63 Version 1.0.2

5.4.7 Multi-tenancy Risks and Concerns

Description: Cloud computing provides the potential to reduce costs through resource sharing. Different tenants use services provided on common cloud computing hardware and software simultaneously. The most common intuitive concerns are that:

- A tenant may access to other tenants' virtual machines, network traffic, actual/residual data, or other resources; and
- A tenant may impact the normal operation of other tenants, access their data or identities.

Importance: Although many network services and programs have simultaneously supported multiple tenants in the past, cloud computing elevates this concern because the resource sharing is pervasive, exposes many possibly vulnerable interfaces, and potentially occurs at a very large scale. Thus, this is a new challenge and federal agencies are not familiar with this kind of massive resource sharing and its security ramifications. The uncertainty may impede the adoption of cloud computing. The following mitigations address these concerns by ascertaining application separation and data encryption in cloud computing.

Mitigation 1: Consumers need to take steps to verify and cloud service providers need to apply data encryption, including the following aspects:

- Data in transit: Encrypt data using a one-time session key similar to how SSL/TLS works. Data at rest: Selectively encrypt sensitive data using NIST 140-2 validated algorithms;
- Manage keys separately from data with higher privileges and preferably make them accessible only through defined procedures/programs;
- Change keys periodically and ensure that data is unencrypted and re-encrypted with the new key; and
- Compile and/or wrap the encryption procedure/program to hide additional data transformation or padding to make it even harder for a snooper to get the key.

Mitigation 2: Consumers need to take steps to verify and cloud service providers need to apply Application Partitioning, including:

- Separate access control functionality from business processing functionality;
- Separate logic processing functionality from data access functionality;
- Separate user functionality from system management functionality; and
- Aggregate functionalities with similar security requirements to run in the same virtual environment and take advantage of modern compartmentalized data centers (vLANs/sub-network zones with varying levels of security controls).

Mitigation 3: Consumers need to take steps to verify and cloud service providers need to apply logical separation, including:

- Support holistic logical separation of the resources at all the layers: computing (virtualization), networking (vSwitches and vLANs), and storage (logical separation of files with access controls);
- Secure the virtualization server (hypervisor isolation settings to limit accesses);
- Secure the virtual network by working hand-in-hand with the physical network security, especially against man in the middle attacks (MAC spoofing and ARP poisoning); and
- Harden the VM so that the virtualization layer is not exposed to attack.

Mitigation 4: Consumers need to take steps to verify and cloud service providers need to apply physical separation, including:

- Special virtual environments with physical separation of the full-stack cloud infrastructure provisioning to customers with special security requirements; and
- Consider special virtual environments provisioning standardization to respond to increasing demands.

References:

- Draft Cloud Computing Synopsis and Recommendations - <http://csrc.nist.gov/publications/drafts/800-146/Draft-NIST-SP800-146.pdf>
- Proposed Security Assessment & Authorization for U.S. Government Cloud Computing - <http://www.cio.gov/pages.cfm/page/Federal-Risk-and-Authorization-Management-Program-FedRAMP>
- Security Guidance for Critical Areas of Focus in Cloud Computing V2.1 - <https://cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf>
- Top Threats to Cloud Computing V1.0 - <https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>
- SaaS, PaaS, and IaaS: a Security Checklist for Cloud Models - <http://www.csoonline.com/article/print/660065>
- Cloud – 10 Risks with Cloud IT Foundation Tier - https://www.owasp.org/index.php/Cloud-10_Risks_with_Cloud_IT_Foundation_Tier
- Edward L. Haletky, “VMware vSphere and Virtual Infrastructure Security”, Prentice Hall, 2009, ISBN-13: 978-0-137-15800-3.
- Cloud Computing and Security – A Natural Match - http://www.trustedcomputinggroup.org/files/resource_files/1F4DEE3D-1A4B-B294-D0AD0742BA449E07/Cloud%20Computing%20and%20Security%20Whitepaper_July29.2010.pdf.

5.4.8 Cloud-Based Denial of Service

Description: Because cloud customers depend on functional networks to access their resources, and because networks are often not under customer control, there is a perceived increase risk that services provided using the cloud model may not be available. Note: High latency on the cloud carrier network and operational errors that have been widely observed and reported over the last year may have the same net effect as a successful Denial of Service (DoS) attack.

Importance: In the case of cloud computing, the DOS attack surface can expand. Internally accessed applications which transition to remotely accessed cloud services are exposed to network-based DoS threats. Through multi-tenancy, DoS attacks can be launched by insiders through shared resources, as in the case of side channel attacks. Malicious users can theoretically initiate distributed DoS attacks at a new level of severity using the vast resources of cloud.

Mitigation 1: The cloud consumer may adopt a hybrid approach to contract with two or more cloud providers. This improves the probability that an outage experienced by one cloud provider will not result in total loss of cloud consumer access to cloud-based data unless cloud provider two also experiences an outage or share a common vulnerability (e.g., exposure to a national emergency or critical infrastructure).

Mitigation 2: The cloud consumer may contract with a cloud carrier (or cloud broker) for diverse network access from customer site(s). Cloud consumer site(s) access diversity can take the form of ingress/egress, route, switch, serving wire center and interconnection points.

Mitigation 3: The cloud consumer may contract a cloud carrier, or cloud broker, to supply redundant customer premises equipment (CPE) with failover (FO) capability to provide high-availability network access to complement diverse network access to cloud provider network. The cloud carrier, through its transport agent, for example, may provide required equipment as part of the cloud-based service contract.

References:

- CSA Cloud Control Matrix
- Federal Risk and Authorization Management Program (FedRAMP)
- NIST Cloud Computing Reference Architecture (30 March 2011), Version 1
- NIST SP 800-53: Recommended Security Controls for Federal Information Systems and Organizations

5.4.9 Incident Response

Description: Incident response and computer forensics in a cloud environment require different tools, techniques, and training to accurately assess a situation and capture appropriate evidence when conducting an incident response that follows federal incident response guidelines. The response plan should address the possibility that incidents, including privacy breaches and classified spills, may impact the cloud and shared cloud customers.

Importance: This requirement highlights the need to update guidance and procedures to comply with federal incident response and reporting requirements and mission operational needs in a cloud environment.

Mitigation: Cloud providers should develop and provide a documented incident response plan that is consistent with existing federal guidance and supports the robust NIST four-phase incident handling guide that is implemented within the federal government. This incident response life cycle consist of Preparation, Detection and Analysis, Containment, Eradication, and Recovery, and Post-Incident Activity.

References:

- NIST SP 800-61: Computer Security Incident Handling Guide

6 Other Related Work

This section focuses on relevant issues to the cloud computing model that arose during the November 2010 – September 2011 course of study, including the NIST-chaired public working groups. Discussions on these topics are well suited to and will continue to be studied by subgroups.

6.1 Cloud Data Issues

Germane to the study of computing is the manipulation, processing, storage, and transmission of data. Identifying the importance of data, common data functions and data-intensive implementations as they relate to cloud computing are a key underpinning. The data functions can be categorized in two tiers, one as an underlying operational tier and the second as a higher-level informational tier. The distinctions between the two tiers are important because of the functions that these two data types provide and are made clearer when considered in terms of the primary users. Data at the operational tier is more likely to be used by the Cloud Provider, Cloud Auditor, Cloud Broker, and Cloud Carrier. In some cases, the Cloud Consumer may need to use this type of data as well. Operational data functions support the manipulation, extraction, and presentation of meaningful results to end users. For the informational data type, the Cloud Consumer is considered the chief user; however, other actors in the cloud computing environment may use this as well.

6.1.1 Operational Data Functions

The following is a list of typical data services functions that are associated with data in the cloud.

- Analytics Services - Reporting and Business Intelligence Services
- Change Control/Tracking - Track User Versions of Files, View/Restore of prior versions
- Common Functions - Data Delete, eDiscovery, Data Fusion, Data Visualization, Data filtering/reduction
- Data Integrity Services - Data Replication for Disaster Recovery and Business Continuity, Data Recovery objectives (i.e., time and point), Data authenticity, Media Sanitization
- Data Maintenance - Backup/Restore, Retention/Hold
- Data Portability - File Portability, Archive Portability, Meta Data Portability, Database Portability, Document and Record Portability
- Data Security - Identity and Privilege Management, Data sensitivity and protection, User Access/Role Controls, Forensic Analysis tools
- Data Storage and Archive - Data Archive and Restore, Application storage, Internet “Drive” secondary storage, “Scale out” storage, Compression, Encryption, Latency, Throughput, Long Term and temporary retention and preservation, Database/Data Warehouse/Business Intelligence, Video Library, Disk-Archive management
- Data Translation - Data Locality
- Data Transport - Data Presentation – Streaming and feeds, Cloud Data Exchange / Synchronization, Common file sharing (e.g., Wikis etc.), Bulk data transfers, Geographic Placement
- File Management - Create/Modify/Delete files, Distribute files
- Policy Management - Common standard Management Framework and interface, Quota Management, Archive Policy Management, Exception Management, Data locality policy administration, Geographic restriction on data location, Disclosure Policy and implementation review, Security

Policy compliance assessment (FISMA, DoD, etc.), Privacy Policy compliance review, Support for Multiple Data Policies (GAAP, HIPPA, etc.)

- Reporting Services - Power Utilization tracking and optimization, Administrative Reporting, Notification requests and management (e.g., notify when a reference document updated), Power Consumption tracking, Provider SLA reporting which including performance not accessible to general users, Activity Review, Quota management
- Search - File Name and Content Search, Advanced Search (owner, creation date, modification date, accessed by)
- Others - Database Operations Services, Published reference files, Forms, Training (student materials, videos, testing), Data interoperability

With this list, the above operational data functions can now be mapped to distinct sections of the RA. Security and privacy for the operational data functions are cross-cutting issues for all of the tabulated items as well:

	Service Layer			Cloud Service Orchestration	Resource Abstraction	Physical Resource
	SaaS	PaaS	IaaS			
Analytics Services	x	x	x			
Change Control/Tracking	x	x				
Common Functions		x				
Data Integrity Services	x	x				
Data Maintenance			x			
Data Portability			x			
Data Security	x	x	x			
Data Storage and Archive			x			x
Data Translation	x	x				
Data Transport	x	x	x			x
File Management	x	x				
Policy Management				x		
Reporting Services (administrative, SLA, data movement, etc.)				x		

Search	x	x				
VM Instance Management					x	

6.1.2 Informational Data and Data Services

Besides the operational data functions identified above, informational data and their associated services play important roles in the cloud computing landscape. Data services are not new computing concepts. With the use of cloud computing where the aggregation or the mash-up of multiple data sources, located in data centers across the globe, into a correlated purposeful data set needs to be identified in the Cloud Computing Reference Architecture.

Data services can be defined as a set of computing services exposing informational data in a way that adhere to cloud computing reference architecture – stand-alone or within a system of systems. There are many prominent examples, which with Application Program Interface (API) provide end users with human-readable meaningful results. These services are useful to end users because of the standardized format and methodologies that allow them to work seamlessly.

Data services that are derived from informational data, depending on their usage, can be categorized as a part of Software as a Service (SaaS) or as a part of Platform as a Service (PaaS). In SaaS or PaaS, to leverage the data and their associated metadata, software applications or standard Web interfaces are needed to extract the intended information from disparate data sets. The *NIST Cloud Computing Standard Roadmap* document defined data functions within the SaaS and PaaS environments.

SaaS

The varieties of the SaaS applications determine what can be consumed by the SaaS consumer. There are varying degrees of functional standardization. SaaS applications are mostly consumed using a Web browser, and some are consumed as a Web service using other application clients, such as stand-alone desktop applications and mobile applications.

For example, standard metadata format and APIs are needed to describe and generate eDiscovery metadata for emails, document management systems, financial account systems, etc., that will help government consumers to leverage commercial off-the-shelf (COTS) and government off-the-shelf (GOTS) software products to meet eDiscovery requirements. This is especially important when email messaging systems, content management systems, or Enterprise Resource Planning (ERP) and financial systems are migrated to a SaaS model.

PaaS

PaaS functional interfaces encompass the runtime environment with supporting libraries and system components for developers to develop and deploy SaaS applications. Standard-based APIs are often part of a PaaS offering to begin with (such that the PaaS provider can lure existing development away to cloud-based hosting environment).

6.2 Service-Level Agreement Taxonomy

Highlights: Through the procedure of defining the cloud computing reference architecture, the NIST-led cloud computing reference architecture working group also identified cloud SLAs as an important gap that needs further clarification.

In April 2011, the SLA subgroup was formed and a survey of the publicly available cloud SLAs was conducted.

The study showed the disparities and ambiguities in cloud providers' SLAs, which confirms the necessity for industry and USG agencies to develop **“Technical Specifications to Enable Consistent, High-Quality Service-Level Agreements” - NIST USG Cloud Computing Technology Roadmap Vol.1, Requirement 3.**

Note: NIST has provided the SLA Taxonomy to the General Services Administration for reference in its development of cloud computing procurement guidance.

At the completion of version 1.0 of the Reference Architecture (RA) the Taxonomy subgroup was asked to identify additional areas of cloud computing that could be better defined through the development of appropriate taxonomies. The group reached immediate consensus that cloud Service-Level Agreements would be an ideal area for an additional taxonomy. (The SLA is a contract between a cloud service provider and a cloud service consumer that specifies, in measurable terms, what services and guarantees the cloud provider will provide.)

A survey of publicly available SLAs showed that while numerous cloud SLAs exist, there is little harmonization between the different types, key elements, and vocabulary. With no universally accepted cloud SLA format, no clear guidance on how required policies can be mapped to a SLA, and differing terminology, it was clear that the area of cloud SLAs could be enhanced through the development of a suitable taxonomy. Creating a SLA taxonomy would establish both a SLA classification system (identifying key elements

that should exist within a given SLA) as well as a controlled vocabulary of terms and definitions (which would facilitate meaningful communication). With this clear need identified, the group then proceeded to work on a draft cloud SLA taxonomy.

The first issue encountered was identifying the proper level at which to start the taxonomy. The natural inclination is to start with cloud Service-Level Agreements, but it is apparent that starting one level of abstraction higher (at what is often referred to as the Master Term of Service-level) provided a better grounding for establishing the common understanding of the domain. This also helped separate many of the traditional elements of a SLA (non-cloud specific) to be dealt with at the higher level. This was an important distinction since SLAs have existed for some time, and this would allow the group to focus its efforts on cloud specific elements of the SLAs.

After the starting point was established, the resources identified by the group were then reviewed to identify common elements that should appear within a SLA. These elements were then organized into two mindmaps (pictorial representations of taxonomies) that reflect the planned separation into the master terms of service and the cloud Service-Level Agreements. Within the master term of service mindmap, a sub child of the top element was then identified as the cloud Service-Level Agreement (CSLA), which would then hold the cloud-specific SLA elements.

The two mindmaps generated by this exercise are listed below:




Table 9: Service-Level Agreement Generic Concepts Mindmap




Table 10: Cloud-Specific SLA Concepts Mindmap

In the CSLA mindmap, several interesting items were identified. First was the fact that within the CLSA, there was a split between elements that support business-level objectives and service-level objectives. Second, an enforceable SLA requires measurable cloud service metrics, which supports the concept of a “resource” which is only implied in the main RA documentation. In the exercise, it was notable that in many cases, the objectives could be mapped to the NIST CC RA which provides additional support to the RA structure.

This exercise was valuable in that it helped perform a survey of the key elements that should appear within a cloud-focused SLA.

6.3 Reliability Research in Cloud-based Complex Systems

Cloud computing systems are complex, encompassing enormous scale and capability. This complexity implies that

- 1) Failures in such systems can emerge from event sequences that are difficult to predict; and
- 2) The consequences of those failures, which typically require substantial time to diagnose and repair, can prove quite costly.

These factors, along with numerous and continuing failures in cloud computing systems, led NIST to identify the need:

- To formulate and publish best practices on achieving reliability;
- To develop a consensus process to measure and report industry-wide cloud reliability information;
- To develop methods for measurement and monitoring to predict onset of catastrophic failure in cloud systems; and
- To investigate tools to identify failure vulnerabilities in designs and deployments.

NIST researchers are pioneering methods to model, analyze, and predict global behavior in complex information systems, such as the Internet and computational grids and clouds.

With respect to cloud systems, these modeling and analysis methods have been used to compare resource-allocation algorithms and to discover potential virtual machine leakage vulnerabilities in open-source IaaS clouds.⁷⁸ Future NIST research will focus on adapting modeling and analysis tools from the physical sciences to identify failure vulnerabilities in designs and deployments of IaaS cloud systems and related cloud applications. Success in this research will enable designers and providers of cloud systems to identify potential reliability vulnerabilities and to develop designs and deployment strategies to mitigate those vulnerabilities, leading to increased cloud reliability, and reducing the costs associated with extensive cloud failures.

NIST researchers are currently planning to investigate measurement and monitoring regimes that can predict the onset of catastrophic failure in cloud systems. Success on this latter research can improve the effectiveness of monitoring and measurement regimes designed and deployed by cloud providers.

⁷ *Koala*: A Discrete-Event Simulation Model of Infrastructure Clouds, K. Mills, J. Filliben and C. Dabrowski

⁸ C. Dabrowski and K. Mills, VM Leakage and Orphan Control in Open-Source Clouds

7 Summary and Next Steps

The *USG Cloud Computing Technology Roadmap*, including Volume II, *Useful Information for Cloud Adopters*, and the work used as the basis for drafting it, was completed in one year. Volume II is not an exhaustive or complete reference of technical work in the subject areas of cloud computing reference architecture and taxonomy, business and technical use cases, standards, and security.

However, Volume II is a key component of the first Draft USG Cloud Computing Technology Roadmap which is intended to be a first step toward a two-fold objective:

- Strategic – Volume II presents technical work which objectively communicates the rationale for the high-priority USG Cloud Computing Requirements; and
- Tactical – Volume II supports adopters in the interim period while the cloud model and implementation is maturing by providing information to reduce uncertainty.

There is a need to explicitly agree on strategic priorities. This seems basic, yet prior to the roadmap effort there wasn't a manageable consolidated list of USG interoperability, portability, and security high-priority requirements related to standards, guidance and technology. There are many potential sources – numerous publications by academic, standards, and industry organizations, and government agencies. However, when consolidated, they yield hundreds of requirements.

The roadmap process assessed and synthesized the inputs from a broad set of collaborators and sources, and applied some level of research and analysis to determine the highest priorities and recommended Priority Action Plans, candidates for self-tasking by the cloud community presented in Volume I. The Volume I high-priority requirements will be refined and met over a multi-year time frame, inherent in the paradigm of emerging technology development.

In addition to summarizing the work that supports these priorities, Volume II provides the basis for defining immediate actions in the short term that can be completed in parallel with developing and executing the PAPs. Each major area of work represented in Volume II is supported by a tactical collaborative process which is under way. Activities which can continue to immediately go forward with cloud computing community participation include:

- Use of the Reference Architecture and Taxonomy by cloud service providers to consistently categorize services so that USG agencies can compare services and products more easily; (SP 500-292), applied to Service-Level Agreement specifications;
- Continued identification and development of Cloud Computing interoperability, portability, and security standards, including USG involvement, and starting with the current list identified in the *NIST Cloud Computing Standards Roadmap* (SP 500-291);
- Development and exchange of additional USG Target Business Use Cases and their SAJAAC technical counterparts; leverage the SAJACC process and portal to continue the qualitative test process that was demonstrated through proof-of-concept;
- Assessment of existing IT security management and technical controls and solutions in the context of the high-priority security requirement challenges, and development of the mitigation solutions; and
- Additional application of complex computing research to the Cloud Computing model.

Appendix A. USG Federal Cloud Computing Standards Working Group Members and Interagency, Academic, Standards Organizations, and Industry Contributors

The views and opinions of the following contributors are entirely their own and do not necessarily represent their employer or affiliated organization's positions, strategies, or opinions. Inclusion of affiliation does not imply endorsement of the contributor's views and opinions by the contributor's employer or affiliated organization.

USG Federal Cloud Computing Standards Working Group Members

Bruce Beckwith, Department of Energy

Earl Crane, Department of Homeland Security

Dominic Gomes, Office of the Chief Information Officer, Department of Health and Human Services

Lon D. Gowen, Ph.D., National Aeronautics and Space Administration, Goddard Space Flight Center

David Harrity, Department of Education

Audrey M. Hogan, Tennessee Valley Authority

Thomas Kireilis, General Services Administration

Dr. Prabha N Kumar, Special Assistant, DoD CIO

Stefan Leeb, Program Manager, NOAA

Hamilton Miller, Department of Justice

Festus C. Onyegbula, Office of Information Technology, National Institute of Food and Agriculture, U.S. Department of Agriculture

Charles Santangelo, Senior IT Budget Manager, Capital Planning and Governance, Office of the CIO, NASA

Lew Sanford Jr., DCS-OESAE, Social Security Administration (with input from SSA staff)

Robert C. Seay, General Services Administration

Gerald L. Smith, Department of Defense and OASIS

Param Soni, Environmental Protection Agency

Vincent Sritapan, DHS HQ, Cybersecurity Strategy

James Ramskill, Office of the Director of National Intelligence

David Raw, Office of the Chief Information Officer, Department of Homeland Security

Peter Tseronis, Chief Technology Officer, Department of Energy

Interagency, Academic, Standards Organizations, and Industry Contributors

Shin Adachi, GICTF- Global Inter-Cloud Technology Forum, NTT DATA Agilent, L.L.C.

Gabriel Akisanmi, KPMG LLP

Leslie Anderson, Raytheon Company

Gary Ardito, NetIQ

Scott Armstrong, Symantec Corporation

Kapil Bakshi, Cisco Systems Inc.

Jeffrey S. Bardin, Treadstone 71 Utica College

Roger Bass, Traxian, OASIS

Bill Becker, SafeNet, Inc.

Bhavesh C. Bhagat, Cloud Security Alliance DC, ConfidentGovernance.com, EnCrisp LLC

Corey Bidne, USDA

Michael Binko, kloudtrack, Software and Information Industry Association

Dr. Alan H. Blair, Defense Engineering Inc.

Mark Bohannon, Red Hat, Inc.

Robert Borochoff, Administrative Office of the US Courts

David W. Boyd, Data Tactics Corporation, Lorenz Research Corp.

Richard Brackney, Microsoft

Nadeem Bukhari, Kinamik Data Integrity

Winston Bumpus, DMTF, VMware, Inc.

William (Bill) Butler, Capitol College

Kevin Call, Booz Allen Hamilton

Karen Luigard Caraway, The MITRE Corporation

Mark Carlson, SNIA, DMTF, Oracle Corporation

Peggy Canale, Avocent Products and Services, Emerson Network Power

Saravana R. Chandran, Strategy and Technology Direction

Te-An Chang, Compuwright Solutions

Gene Cartier, SRA International

Eric Charlesworth, Cisco Systems, Inc.

Arunava Chatterjee, Deloitte Consulting LLP

G. Hussain Chinoy, USDA NRCS

Augusto Ciuffoletti, Università di Pisa, Italy

John Crandall, Brocade

John Crout, United States Coast Guard Auxiliary

Cory Dell, Coupa Software

Yuri Demchenko, University of Amsterdam

Frederic de Vault, Prometheus Computing, LLC

Michele Drgon, DataProbit

Josiah Dykstra, UMBC

Carlo Espiritu, Triple Point Security

Christopher Ferris, IBM

Omar Fink, SAIC

L. Bruce Finn, Federal Deposit Insurance Corporation

David A. Foley, SNHU former student

Harry J. Foxwell, PhD, Oracle Corporation

Barry Garman, The Mercator Group

Parisa Ghodous, University of Lyon I

Richard Gordon, Jr., RICHMAR & Associates

Nedim S. Goren, U.S. Census Bureau

Dr. Nancy W. Grady, SAIC
Jay Greenberg, IEEE-USA
Mateen Greenway, HP
A. Larry Gurule, CSC, FCP
Daneyon Hansen, Cisco Systems
Doug A. Hansen, Department of Homeland Security
David Harper, Johns Hopkins University
Thor Henning Hetland, Webstep AS, Cantara AS
Jenny Huang, AT&T Inc.
David P. Hunter, VMware, Inc.
Istian Islam, GTSI Corp.
Anthony Jackson, Army Contracting Command - National Capital Region
Kevin L. Jackson, NJVC, LLC
Babak Jahromi, Microsoft Co.
Karuna P. Joshi, University of Maryland, Baltimore County
Harun Kazaz, Booz Allen Hamilton
Ravi Kalaputapu, Ph.D, Converge Networks Corporation
Shrikanth Kashyap, The Open Group, Wipro Technologies
Lawrence Kelly, Kelly Technology Enterprises, Inc.
Dean Kemp, Independent Consultant
Joe Keochinda, Livanta
Jerry Kickenson, SWIFT
Paul Krein, Federal Office of the CTO
David Kye, Deloitte Consulting, LLP
Donald Lamb, Booz Allen Hamilton
Cary Landis, Virtual Global, Inc.
Nancy M. Landreville, PhD, Department of Veterans Affairs, University of MD
Michel Landry, Systec LLC
Margaret Leary, Avaya Government Solutions
David LeDuc, Software & Information Industry Association
Cheng-Yin Lee, Independent Consultant
Dr. Craig A. Lee, Open Grid Forum, The Aerospace Corporation
Keutlwile Leso, Molemi Global
Bob Linehan, CoreMax
Dorothy Lorenz, Unisys Corporation
Eugene Luster
Robert M. Mack, SunGard Availability Services LP
Shamun Mahmud, DLT Solutions, Incorporated
Robert Marcus, ET-Strategies, Cloud Standards Customer Council
John T. McDonald, Raytheon
Mike McGee, Coalfire Systems, Inc.
Steven McGee, SAW Concepts LLC
Matthew Metheny, One Enterprise Consulting Group, LLC
T.S. Mohan, PhD, INFOSYS LIMITED
Felicia Moore, Department of Transportation
Shawn P Myers, Coventry Health Care
Stacey Myers, The MITRE Corporation
Felix N. Njeh, CSC
Michael P. O'Doherty, Lockheed Martin
Anthony Pagano, Community Health Network of CT

Lilia R. García Perellada, Instituto Superior Politécnico José Antonio Echeverría
William Perlowitz, URS-Apptis
Rodney Pieper, HP
Tom Plunkett, Oracle
James M. Poffel, SunGard Availability Services LP
Ioannis Polyzos, Glasgow Caledonian University
Donita Prakash, Acumen Solutions
Sundararajan (Sundar) Ramanathan, Capgemini
Ryan K. Rees, SABRE SYSTEMS, INC
John W. Rogers, Data-Tactics Corporation, US Navy Maritime ISR Cloud Environment and the
Intelligence Community Cloud Computing Integration Environment Working Group
Matthew Rogers, Booz Allen Hamilton
Dr. Ken Roberts
Tom Rutt, Fujitsu America
Regina Ryan, MITRE Corporation
Paul Sand, IP3, Inc., Chicago InfraGard Members Association, Illinois Terrorism Task Force
Richard Santalesa, Esq., InfoLawGroup, LLP
Dr. Hasan Sayani, University of Maryland University College
Andrey Sazonov, Coalfire Systems
Naresh Sehgal, Intel Corp
Paul Sforza, U.S. Department of the Treasury
Sean Sherman
Alan Sill, Ph.D., Texas Tech University, Open Grid Forum
Charles Spence, Healthland
Ken E. Stavinoha, Cisco Systems
Mike Stewart, Department of Navy
Andrew Strear, The New York Times Company
Richard Tychansky, Lockheed Martin Corporation
Emmet J. Tydings, AB&T Telecom
Mark Underwood, TransitCenter Inc.
David Vidal, Polytechnic University of Madrid
Brian Vosburgh, Stonesoft Inc.
Bryan Ward, Serco
Steven Woodward, Cloud Perspectives
David L. Woolfenden, eVectis Technologies LLC
James Yaple, Department of Veterans Affairs
Gwen Young, Office of Natural Resources Revenue
Michael Young, Esri
Robert Zimmerman, Inforistec Group
Joel P. Zysman, University of Miami Center for Computational Science

Appendix B. Useful References

The following sources may be useful for further reference.

NIST Special Publications and Drafts

[NIST Special Publication 800-53](#), Rev. 3, *Recommended Security Controls for Federal Information Systems and Organizations*.

[NIST Special Publication 800-61](#), Rev.1, *Computer Security Incident Handling Guide*.

[NIST Special Publications 800-144](#), *Draft Guidelines on Security and Privacy Issues in Public Cloud Computing*.

[NIST Special Publication 800-145](#), *The NIST Definition of Cloud Computing*.

[NIST Special Publication 800-146](#), *Draft NIST Cloud Computing Synopsis and Recommendations*.

[NIST Cloud Computing Use Cases](#).

[NIST IR-7756](#), *DRAFT CAESARS Framework Extension: An Enterprise Continuous Monitoring Technical Reference Architecture*.

Other Sources

Apache, *LibCloud*, <http://incubator.apache.org/libcloud/>.

Charlton, Stuart. [Cloud Computing and the Next Generation of Enterprise Architecture](#), Sys-Con Cloud Computing Expo. San Jose, CA, 2008.

Chief Information Officers Council, [Privacy Recommendations for the Use of Cloud Computing by Federal Departments and Agencies](#). 19 August 2010.

CISCO, [Cisco Cloud Computing - Data Center Strategy, Architecture, and Solutions: Points of View White Paper for U.S. Public Sector](#), 1st edition. 2009.

Cloud Security Alliance (CSA), [Security Guidance for Critical Areas of Focus in Cloud Computing V2.1](#), December 2009.

Cloud Security Alliance (CSA), [Top Threats to Cloud Computing V1.0](#), March 2010.

Cockburn, Alistair, [Writing Effective Use Cases](#), Addison-Wesley, 2001.

CSO Security and Risk Online, [SaaS, PaaS, and IaaS: a Security Checklist for Cloud Models](#). 31 January 2011.

Department of Homeland Security, [National Strategy for Trusted Identities in Cyberspace: Creating Options for Enhanced Online Security and Privacy, DRAFT](#), 25 June 2010.

Distributed Management Task Force, Inc. (DMTF), [*Interoperable Clouds: A White Paper from the Open Cloud Standards Incubator V1.0.0*](#), DSP-IS0101, 11 November 2009.

Distributed Management Task Force, Inc. (DMTF), [*Architecture for Managing Clouds: A White Paper from the Open Cloud Standards Incubator V1.0.0*](#), DSP-IS0102, 18 June 2010.

Distributed Management Task Force, Inc. (DMTF), [*Use Cases and Interactions for Managing Clouds: A White Paper from the Open Cloud Standards Incubator V1.0.0*](#), DSP-IS0103, 18 June 2010.

Federal CIO Council, [*Proposed Security Assessment & Authorization for U.S. Government Cloud Computing*](#). Draft version 0.96, 2 November 2010.

[*Federal Information Security Management Act of 2002*](#) (FISMA), December 2002.

Federal Standard 1037C, [*Telecommunications: Glossary of Telecommunications Terms*](#), 7 August 1996.

Gartner, [*Gartner Says Cloud Consumers Need Brokerages to Unlock the Potential of Cloud Services*](#). 9 July 2009.

Gasser, Morrie. [*Building a Secure Computer System*](#), Van Nostrand Reinhold Co., 1988.

Global Inter-Cloud Technology Forum (GICTF), [*Use Cases and Functional Requirements for Inter-Cloud Computing White Paper*](#), 9 August 2010.

GSA, [*Cloud Computing Initiative Vision and Strategy Document \(DRAFT\)*](#), February 2010.

Haletky, Edward L. [*VMware vSphere and Virtual Infrastructure Security*](#), Prentice Hall, 2009.

IBM, [*Introducing the IBM Security Framework and IBM Security Blueprint to Realize BusinessDriven Security*](#), 5 November 2010.

IBM, [*Cloud Computing Reference Architecture 2.0*](#), February 2011.

Juniper Networks, [*Cloud-ready Data Center Reference Architecture*](#), February 2011.

“[*Non-repudiation*](#).” IBM WebSphere MQ Information Center, 3 May 2011.

OASIS, [*OASIS Privacy Management Reference Model Technical Committee Charter*](#), http://www.oasis-open.org/committees/pmr/reference_model/charter.php.

Office of Management and Budget (OMB), [*Federal Cloud Computing Strategy*](#). 8 February 2011.

Office of Management and Budget, Memorandum 07-16, [*Safeguarding Against and Responding to the Breach of Personally Identifiable Information*](#). 22 May 2007.

The Open Group Architecture Framework (TOGAF), [*Section 21.3*](#).

Open Security Architecture (OSA), *SP-011: Cloud Computing Patterns*,
<http://www.opensecurityarchitecture.org/cms/library/patternlandscape/251-pattern-cloudcomputing>.

The Open Web Application Security Project, *Cloud – 10 Risks with Cloud IT Foundation Tier*. 26 July 2009.

OpenCrowd, *Cloud Taxonomy*. <http://cloudtaxonomy.opencrowd.com/>.

Storage Network Industry Association (SNIA), *Cloud Storage for Cloud Computing*, September 2009.

Storage Network Industry Association (SNIA), *Cloud Storage Use Cases*, 8 June 2009.

Symantec, Internet Security Threat Report, *Trends for 2010*, Volume 16, April 2011.

“*Taxonomy*.” Webopedia.com, 2011.

Trusted Computing Group, *Cloud Computing and Security- A Natural Match*, April 2010.