

Special Publication 500-293

**US Government Cloud Computing
Technology Roadmap
Volume I**

**High-Priority Requirements to
Further USG Agency Cloud
Computing Adoption**

Lee Badger, David Bernstein, Robert Bohn, Frederic de Vault, Mike Hogan, Michaela Iorga, Jian Mao, John Messina, Kevin Mills, Eric Simmon, Annie Sokol, Jin Tong, Fred Whiteside and Dawn Leaf

*This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.SP.500-293>*

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

This page left intentionally blank

NIST Special Publication 500-293

US Government Cloud Computing Technology
Roadmap Volume I

High-Priority Requirements to Further USG Agency
Cloud Computing Adoption

Lee Badger, David Bernstein, Robert Bohn, Frederic de
Vaulx, Mike Hogan, Michaela Iorga, Jian Mao, John
Messina, Kevin Mills, Eric Simmon, 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

This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.SP.500-293>

October 2014



U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Acting Under Secretary for Standards and Technology and Acting 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

Natl. Inst. Stand. Technol. Spec. Publ. 500-293, 40 pages (October 2014)

This publication is available free of charge from:

<http://dx.doi.org/10.6028/NIST.SP.500-293>

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, David Bernstein, Jian Mao, and Jin Tong, of Knowcean Consulting Incorporated (under contract through the Federal Integrated Product Team, Space & Naval Warfare [SPAWAR] Systems Center Atlantic), Frederic de Vault of Prometheus Computing LLC (under contract), Fred Whiteside of the Department of Commerce, and Lee Badger, Robert Bohn, Mike Hogan, Michaela Iorga, John Messina, Kevin Mills, Eric Simmon, Annie Sokol, 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 working groups.

We especially acknowledge Lisa Carnahan, Romaine Hines, Mary Saunders, Terry Schwarzhoff, and James St. Pierre of NIST for providing editorial review and support.

We especially acknowledge Earl Crane, of the Department of Homeland Security and Information Security and Identity Management Committee (ISIMC), whose advice and technical insight assisted this effort.

We also wish to acknowledge the members of the Federal Cloud Computing Standards and Technology Working Group and other interagency contributors, listed in Appendix A of this document.

This page left intentionally blank

Table of Contents

Executive Summary	ix
1 Purpose and Scope	1
1.1 USG Cloud Computing Technology Roadmap Purpose.....	1
1.2 Intended Audience and Use.....	1
1.3 Document Organization	2
2 USG Cloud Computing Technology Roadmap Requirements	4
2.1 Requirement 1: International Voluntary Consensus-Based Standards	5
2.2 Requirement 2: Solutions for High-priority Security Requirements which are technically de-coupled from organizational policy decisions	6
2.3 Requirement 3: Technical Specifications to Enable development of Service-Level Agreements	7
2.4 Requirement 4: Clear & Consistently Categorized Cloud Services.....	8
2.5 Requirement 5: Frameworks to Support Federated Community Clouds	9
2.6 Requirement 6: Updated Organization Policy that reflects the Cloud Computing Business and Technology model.....	10
2.7 Requirement 7: Defined Unique Government Requirements and Solutions.....	12
2.8 Requirement 8: Collaborative Parallel “future cloud” Development Initiatives	13
2.9 Requirement 9: Defined & Implemented Reliability Design Goals	14
2.10 Requirement 10: Defined & implemented Cloud Service Metrics.....	15
3 Other Considerations and Observations	16
3.1 Regarding Academia, Industry, Standards Organizations, and Government Collaboration	16
3.2 Interdependency with Cyber Security initiatives.....	17
3.3 Interdependency with emerging Big Data technology	17
3.4 Organizational Policy	18
3.5 Interdependency with Other National Priority Initiatives.....	18
3.6 Education of Technical Staff and Cloud Consumers.....	19
4. Progress and Next Steps	21

US Government Cloud Computing Technology Roadmap, Volume I

4.1 NIST Cloud Computing Program Future Phases 23

4.2 Summary of Time Line and Deliverables 26

This page left intentionally blank

Executive Summary

The National Institute of Standards and Technology (NIST), consistent with its mission,¹ has a technology leadership role in support of United States Government (USG) secure and effective adoption of the Cloud Computing model² to reduce costs and improve services. This role is described in the 2011 *Federal Cloud Computing Strategy*³ as “... a central one in defining and advancing standards, and collaborating with USG Agency CIOs, private sector experts, and international bodies to identify and reach consensus on cloud computing technology & standardization priorities.”

This NIST Cloud Computing program and initiative to develop a USG Cloud Computing Technology Roadmap is one of several complementary and parallel USG initiatives defined in the broader Federal Cloud Computing Strategy referenced above.

The *Federal Cloud Computing Strategy* characterizes cloud computing as a “*profound economic and technical shift (with) great potential to reduce the cost of federal Information Technology (IT) systems while ... improving IT capabilities and stimulating innovation in IT solutions.*”

In the technology vision of *Federal Cloud Computing Strategy* success, USG agencies will be able to easily locate desired IT services in a mature and competitive marketplace, rapidly procure access to these services, and use them to deliver innovative mission solutions. Cloud services will be secure, interoperable, and reliable. Agencies will be able to switch between providers easily and with minimal cost, and receive equal or superior services.

Decision makers contemplating cloud computing adoption face a number of challenges relating to policy, technology, guidance, security, and standards. Strategically, there is a need to augment standards and to establish additional security, interoperability, and portability standards to support the long-term advancement of the cloud computing technology and its implementation. Cloud computing is still in an early deployment stage, and standards are crucial to increased adoption. The urgency is driven by rapid deployment of cloud computing in response to financial incentives. Standards are critical to ensure cost-effective and easy migration, to ensure that mission-critical requirements can be met, and to reduce the risk that sizable investments may become prematurely technologically obsolete. Standards are key to ensuring a level playing field in the global marketplace.⁴

Recognizing the significance and breadth of the emerging cloud computing trend, NIST designed its program to support accelerated US government adoption, as well as leverage the strengths and resources of government, industry, academia, and standards organization stakeholders to support cloud computing technology innovation.

¹ This effort is consistent with the NIST role per the National Technology Transfer and Advancement Act (NTTAA) of 1995, which became law in March 1996.

² *NIST Definition of Cloud Computing*, Special Publication 800-145, September, 2011.

³ 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.

⁴ The roadmap primary focus on interoperability, portability, and security requirements does not preclude the need to address reliability, maintainability, performance, accessibility and other essential requirements.

Framing the Discussion -- underlying principles and assumptions:

The USG Cloud Computing Technology Roadmap is a mechanism to define, communicate, and recommend:

- ***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 need to be in place to satisfy these requirements; and,***
- ***Candidate Priority Action Plans (PAPs) which are recommended for voluntary self-tasking by the cloud computing stakeholder community to support standards, guidelines, and technology development.***

Following this executive summary, Volume I intentionally presents each requirement at a very basic level, and uses illustrative examples to explain in plain language why from at least one perspective these requirements are not considered to be fully met.

The intent is to lay the groundwork to more directly tackle a subset of cloud computing technology scope, consistent with the Federal Cloud Computing Strategy to accelerate USG cloud adoption. This does not imply an intent to prescribe a USG-centric view.

On the contrary, the “roadmap” is intended to foster a substantive discussion among cloud computing stakeholders in government and the private sector. In practical terms, the roadmap is a vehicle for NIST to fulfill its collaboration role and leverage input from the hundreds of organizations and individuals who have contributed to the NIST-led cloud computing working group analysis and discussions.

The requirements identified in the roadmap are common for the adoption of any emerging technology. Throughout the November 2010 – August 2012 time frame, NIST sought to verify the set that are highest priority for USG agencies. Through public comments provided in response to the draft of this document issued in November, 2011, NIST confirmed that the roadmap requirements are generally accepted to be priorities. Ideally, the roadmap will serve as a vehicle to continue to refine the requirements and identify relevant work which is under way.

Finally, the roadmap initiative is designed to help ensure that NIST’ technical standards, guidance, and research work is focused on the priorities that are most important, not only in the view of NIST computer scientists and researchers, but also in the eyes of those who are building and deploying cloud technology.

The basis for the following list of prioritized requirements is the work completed November 2010 through August 2012 as part of the NIST Cloud Computing program and collaborative *USG Cloud*

US Government Cloud Computing Technology Roadmap, Volume I

Computing Technology Roadmap effort, including disposition of comments received during the December 2011 public comment period.

The USG Cloud Computing Technology Roadmap requirements which are identified as high priorities to further USG Cloud Computing Technology Adoption are:

- Requirement 1:*** International voluntary consensus-based standards (interoperability, performance, portability, and security standards)
- Requirement 2:*** Solutions for High-priority Security Requirements, technically de-coupled from organizational policy decisions (security standards and technology)
- Requirement 3:*** Technical specifications to enable development of consistent, high-quality Service-Level Agreements (interoperability, performance, portability, and security standards and guidance)
- Requirement 4:*** Clearly and consistently categorized cloud services (interoperability and portability guidance and technology)
- Requirement 5:*** Frameworks to support seamless implementation of federated community cloud environments (interoperability and portability guidance and technology)
- Requirement 6:*** Updated Organization Policy that reflects the Cloud Computing Business and Technology model (security guidance)
- Requirement 7:*** Defined unique government regulatory requirements and solutions (accessibility, interoperability, performance, portability, and security technology)
- Requirement 8:*** Collaborative parallel strategic “future cloud” development initiatives (interoperability, portability, and security technology)
- Requirement 9:*** Defined and implemented reliability design goals (interoperability, performance, portability, and security technology)
- Requirement 10:*** Defined and implemented cloud service metrics (interoperability, performance, and portability standards)

Note: The order in which the requirements are listed does not imply relative importance.

US Government Cloud Computing Technology Roadmap, Volume I

These requirements as stated reflect refinement of the November 2011 draft version of this document. Specifically Requirements 2 and 6 have been modified. Requirement 2 now combines two aspects of solutions for high priority security requirements: the solutions must satisfy the USG identified requirements AND must be de-coupled from organizational policy decisions. Requirement 6 has been modified to separately and more clearly acknowledge the need for updated policy guidance that responds to the changes associated with the cloud computing business and technology model. Requirement 8 has been assessed to be an immature future requirement, and some argue not currently essential to further USG Cloud Computing Technology Adoption. Requirement 8 is treated as a “placeholder” or “stretch” requirement.

NIST Cloud Computing program work which supports the definition of these requirements, and the rationale for the assessment that the requirements are not fully met at present, is summarized in Volume II of the roadmap document. Volume II⁵: 1) describes a conceptual Cloud Computing Reference Architecture and Taxonomy, 2) presents USG Business Use Cases and technical cloud use cases, 3) identifies existing applicable standards and guidance, 4) specifies high-priority standards, guidance, and technology gaps, 5) summarizes work completed in the area of Service Level Agreements, and 6) provides insight into the rationale for the list of action plans which are recommended for voluntary self-tasking by government and private sector organizations.

The content of this document was developed by leveraging an open public process that engaged the broad spectrum of Cloud Computing stakeholder communities and the general public. Input to date has been provided through five public workshops held in May and November 2010, April and November 2011, and June 2012. More than 1,500 individuals representing hundreds of organizations participated in these events. NIST also consulted with stakeholders through extensive outreach efforts, including, five public working groups formed in November 2010, and the *Federal Cloud Computing Standards and Technology Working Group*. The latter body was formed under the auspices of the US Federal CIO Council to represent common US government interests. This report has been subjected to a 30-day public review and comment period. All comments received have been carefully reviewed and resolutions are incorporated in preparation of the final version of this report.

⁵ Updates to the November 2011 version of Volume II include reference architecture broker and security architecture elements, a new service level agreement section, and an updated standards assessment and summary.

1 Purpose and Scope

1.1 *USG Cloud Computing Technology Roadmap Purpose*

The collaborative NIST initiative to develop a *USG Cloud Computing Technology Roadmap* and the resulting multi-volume interagency NIST SP 500-293 document is designed to:

- Foster adoption of cloud computing by federal agencies and support the private sector;
- Reduce uncertainty by improving the information available to decision makers; and,
- Facilitate further development of the cloud computing model.

This document is intended to serve as:

- A vehicle to define and communicate high-priority strategic and tactical security, interoperability, and portability requirements; these must be met for USG agencies to further adopt the cloud computing model to meet the *Federal Cloud Computing Strategy* goals;
- A vehicle to define and communicate the relevant standards, guidance, and technology that must be in place to satisfy these requirements;
- A vehicle to define and communicate a list of candidate Priority Action Plans (PAPs) to be developed to support develop standards, guidance, and technology;
- The mechanism to integrate and present analysis, findings, and useful technical artifacts generated through the NIST Cloud Computing program public working groups, internal NIST Cloud Computing and related projects, and the NIST chaired *Federal Cloud Computing Standards and Technology Working Group*, along with referenced related and complementary work that was reviewed and considered in the roadmap generation process;
- The mechanism to focus discussion on the proposed “technology roadmap” steps to move federal IT from its current early-cloud state (“point A”) to a cloud-based foundation (“point B”) and to fully execute the US *Federal Cloud Computing Strategy*); and
- The basis to assess and plan the NIST Cloud Computing program and the *Federal Cloud Computing Standards and Technology Working Group* efforts going forward.

1.2 *Intended Audience and Use*

This publication is intended for a diverse audience:

- **US Policy Makers, US Federal CIO Council, and those with key roles identified in the Federal Cloud Computing Strategy** – as a technology-oriented reference to inform policy and planning;
- **USG Agencies** – as a tool in the context of the USG *Federal Cloud Computing Strategy* risk-based management “Decision Framework for Cloud Migration”; and
- **Cloud Computing Stakeholders (Academia, Government, Industry, Standards Developing Organizations)** – as a consolidated presentation of USG cloud computing technology perspectives, including a list of candidate Priority Action Plans which are recommended for voluntary self-tasking and which present opportunities to leverage stakeholder efforts to further cloud computing.

1.3 Document Organization

The *US Government Cloud Computing Technology Roadmap* is anticipated to evolve and be updated periodically.

This release of this document consists of three volumes. Consistent with the NIST Cloud Computing program strategy, the roadmap focuses on both strategic and tactical objectives related to cloud computing. 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 I, *High-Priority Requirements to Further USG Agency 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.

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*.

The remainder of Volume I is organized into several sections. Section 2 presents the USG Cloud Computing Technology Roadmap requirements. Section 3 presents other considerations which are related to, but out of the scope of, the roadmap initiative and document. Section 4 identifies the Next Steps, as currently planned for the NIST Cloud Computing program and its collaborative USG Cloud Computing Technology Roadmap initiative.

Volume II, *Useful Information for Cloud Adopters*, is designed to be a technical reference for those actively working on strategic and tactical cloud computing initiatives, including, but not limited to, US government cloud adopters. Volume II integrates and summarizes the work completed to date, and explains how these findings support the roadmap introduced in Volume I.

Volume III, *Technical Considerations for USG Cloud Computing Deployment Decisions*, is released as a draft volume. Volume III was developed with input from US Federal agencies and the Federal Cloud Computing Standards and Technology Working Group. 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 draft version defines and proposes a methodology and process, and proof-of-concept examples. Volume III was initiated in parallel, but is logically dependent on the technical work contained in Volume II, and therefore is a less mature part of the roadmap. Consistent with the precedent established in November 2011 for volumes I & II, the initial Volume III draft special publication is released for a 30-day public comment period.

US Government Cloud Computing Technology Roadmap, Volume I

All of these documents 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. See <http://www.nist.gov/itl/cloud/index.cfm>.

2 USG Cloud Computing Technology Roadmap Requirements

The requirements discussed in this section of the USG Cloud Computing Technology Roadmap are those which have been identified as high-priority strategic and tactical security, interoperability, portability, performance and related requirements that must be met for USG agencies to further adopt the cloud computing model to meet the objectives of the *Federal Cloud Computing Strategy*.

Throughout the November 2010 – August 2012 time frame, the NIST Cloud Computing program has sought to analyze, assess, and verify the set of requirements that are of highest priority for USG agencies.

The analysis, assessment, and verification took several forms. This included the public academic, government, industry, and standards developing organization collaborative public working group and outreach activities described earlier. The analysis, assessment, and verification also included the objective, technical research and development activities, internal and collaborative, that are described and referenced in Volume II: *Useful Information for Cloud Adopters*.

Confirmation also included two 60-day review exercises through the *Federal Cloud Computing Standards and Technology Working Group*. This group includes representatives from approximately 30 U.S. government agencies. This review was deemed essential to ensure that the priorities reflect the viewpoint of those in the government who are directly responsible for ensuring that Information Technology resources are applied effectively and securely to support USG agency missions.

The descriptions of each requirement provide an explicit link between:

- Prioritized strategic and tactical requirements that must be met for USG agencies to further cloud adoption;
- Interoperability, portability, and security standards, guidelines, and technology to satisfy these requirements; and
- Recommended voluntary self-tasking Priority Action Plans (PAPs).

Volume 1 is designed to help ensure that NIST technical standards, guidance, and research work is focused on the priorities that are important to those who are deploying cloud technology.

Volume I of the technology roadmap intentionally presents the information related to each requirement at a very high level, and uses the illustrative examples to explain in plain language why these requirements are not fully met at present.

The order in which the requirements are listed does not imply relative importance.

2.1 Requirement 1: International Voluntary Consensus-Based Standards ⁶

Government, industry, and other stakeholders need to define priorities and requirements⁷, develop international voluntary consensus-based interoperability, portability, security, performance, and related standards, and implement them in products, processes and services.

Why: *Standards-based products, processes, and services are essential for USG agencies to ensure that: a) public investments do not become prematurely technologically obsolete, b) agencies are able to easily change cloud service providers to flexibly and cost-effectively support their mission, c) agencies can economically acquire commercial and develop private clouds using standards-based products, processes, and services, and d) the US government supports a level economic playing field for service providers.*

Illustrative example of why this requirement is not considered to be fully met at present:⁸ *While data, software, and infrastructure components that enable cloud computing (e.g., virtual machines) can be ported from selected providers to other providers, the process requires interim steps to move the data, software, and components to a non-cloud platform or conversion from one proprietary format to another.*

Rationale: *USG agencies have identified mission-related requirements that depend on technical interoperability, portability, security and other standards. The NIST public Cloud Computing Standards Working Group identified a small number of emerging standards that respond to requirements which are unique to cloud computing. The NIST initiated Standards Acceleration to Jumpstart the Adoption of Cloud Computing project found interoperability, portability, and security use cases to be tightly coupled, highlighting the need for integrated standards.*

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Develop international consensus-based standards.	<u>2012-2016</u>
Encourage test tool development to support cloud standards development.	<u>2012-2015</u>
Encourage standards conformity assessment practices (e.g. conformance and performance testing, test result validation, tester accreditation) through procurement.	<u>2012-2013</u>
Develop mutual recognition arrangements, to facilitate voluntary sharing and recognition of test results, so that test reports (first, second, or third party) can be used widely by providers to compete in global markets.	<u>2013-2014</u>
Develop additional technical use cases focusing on multi-cloud scenarios.	<u>2013-2014</u>

⁶ "Legislation, policy, and treaty obligations guide how the US government engages with the standards system. The Trade Agreements Act of 1979 and the National Technology Transfer and Advancement Act are two key pieces of US legislation affecting the use of standards developed in the private sector by Federal agencies. The Office of Management and Budget Circular A-119 Revised: Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities establishes policy. These laws and policy require Federal agencies to use international, voluntary consensus standards in procurement and regulatory activities..."

⁷ NIST SP 500-291 Version 2, NIST Cloud Computing Standards Roadmap, Chapter 9, *USG Priorities to fill Cloud Computing Standards Gaps*

⁸ NIST SP 500-291 Version 2, NIST Cloud Computing Standards Roadmap, Section 6.4, *Cloud Computing Standards for Interoperability and Portability*

2.2 Requirement 2: Solutions for High-priority Security Requirements which are technically de-coupled from organizational policy decisions

There are two aspects of this requirement. Solutions need to be defined to address USG security requirements⁹. Equally important, industry needs to develop technical solutions which are abstracted from (and therefore able to support) diverse sovereign, legal, business, or other authoritative policy rules.

***Why:** Federal decision makers need more transparent and effectively demonstrated cloud services' security to inspire confidence to a degree where security is not perceived to be an impediment, and to support risk-based management decisions to migrate additional IT services to the cloud model. Traditionally, IT security has relied on logical and physical system boundaries. The inherent characteristics of Cloud Computing make these boundaries more complex and render traditional security mechanisms less effective. Moreover, the ability to bridge policy differences and policy evolution is essential. Mechanisms must be developed to allow differing policies to co-exist and be implemented with a high degree of confidence, irrespective of geographical location and sovereignty. De-coupling the technical implementation of cloud security controls from the policy of their application will foster cloud adoption because consumers will be able to agree on defined security controls and the methods for their assessment, without having to agree on when it is appropriate to apply them.*

***Illustrative example of why this requirement is not considered to be fully met at present:** While cloud computing security requirements are not separate from general IT security or unique in their entirety, the cloud computing environment presents unique security challenges. Security controls need to be reexamined in the context of cloud architecture, scale, reliance on networking, outsourcing, and shared resources. For example, multi-tenancy is an inherent cloud characteristic that intuitively raises concern that one consumer may impact the operations or access data of other tenants running on the same cloud.*

***Rationale:** This assessment is based on Security Requirements identified by public working groups and USG forums, including the Federal Cloud Computing Standards and Technology Working Group, Information Security and Identity Management Committee, Federal Risk and Authorization Management Program, and private sector publications. A related example that illustrates the need to decouple technical solutions from policy is the September 2012 EC announcement of plans to develop a legal framework that ensures that EU data protection standards will be applied for EU consumers regardless of where the data or service provider is based, and a broader cloud computing related regulatory and legislative environment.*

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Continue to identify Cloud Consumer Priority Security Requirements.	quarterly
Identify and assess the extent to which risk can be mitigated through existing and emerging security controls and guidance.	periodically
Identify gaps and modify existing controls and monitoring capabilities	periodically
Develop neutral cloud security profiles, technical security attributes, and test criteria .	2012 – 2014
Define an international standards-based conformity assessment system approach.	2013 – 2014

⁹ “Security Requirements” refers to the high-priority USG security requirements, summarized in Volume II of the November 2011 draft version of this document.

2.3 Requirement 3: Technical Specifications to Enable development of Service-Level Agreements

Industry and government need to develop and adopt consistent technical specifications, of high quality and completeness, to enable the creation and practical evaluation of Service-Level Agreements (SLAs) between customers and cloud providers.

Why: *Cloud SLAs specify the services that will be provided to a consumer, and represent part of a negotiated service contract between two parties. This requirement must be met to: a) ensure that key cloud service elements (warranties, guarantees, reliability and performance) are defined and enforceable, b) develop common SLA terms and definitions and avoid misunderstandings between parties, and c) create an environment which allows consumers to objectively compare services.*

In utility industries, the notion of units of measurement is fundamental to buying and selling service. This contrasts with the traditional approach in computing operations to benchmark performance of system components such as hardware, operating systems, database and Web servers. Cloud computing service delivery uses a utility model; IT resources are supplied as abstracted services, such as Infrastructure or Platform as a Service. Consumers pay for a metered “quantity” and a “quality” of the service. There is a need for clear and consistent technical specifications to precisely and predictably specify cloud services.

Illustrative example of why this requirement is not considered to be fully met at present: *The concept of reliability is a key cloud computing element addressed by providers’ SLAs. However, the definition of what is being measured, and associated guarantees vary widely. Customers are faced with evaluating SLAs from cloud providers which define reliability using different terms (uptime, resilience, or availability), cover different resources (servers, HVAC systems, data storage, customer support), cover different time periods (hours, days, years), and use different guarantees (response time versus resolution time). SLA and measurement ambiguities leave the customer at risk.*

Rationale: *In creating a Reference Architecture, the NIST public working group identified cloud SLAs as an important gap that needs clarification (scope) and refinement (structure). A survey of publicly available cloud SLAs showed that an industry-wide accepted standard SLA form for cloud services does not exist. Disparities in cloud providers’ SLAs, and issues related to high-profile cloud failures support the conclusion that SLAs are inadequate. Government agencies have specific requirements (e.g. FISMA policy) which require SLA modifications. In 2010 and 2011, the NIST-led public Cloud Computing USG Target Business Use Case, SAJACC, and Security working groups, and in 2012, the Federal Cloud Computing Standards and Technology Working Group, independently confirmed this requirement.*

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Develop a controlled and standardized vocabulary and set of cloud SLA terms and definitions.	2012 ¹⁰ – update periodically
Ensure consistent guidance and policy regarding SLA relevant terms and definitions.	2013 – update periodically

¹⁰ N.b. The NIST Cloud Computing Definition, Reference Architecture and Taxonomy referenced above are among submissions which are currently being worked through international standards bodies.

Develop a cloud SLA Taxonomy to ensure the complete specification of key cloud computing elements that need to appear in an SLA.	2012 – update periodically
--	----------------------------

2.4 Requirement 4: Clear & Consistently Categorized Cloud Services

Industry needs to clearly and consistently categorize cloud services.

Why: *This requirement must be met to ensure that: a) customers will understand the intricacies of different types of cloud services and will be better able to select cloud services suitable to meet their business objectives, b) customers will be able to objectively evaluate, compare, and select between products from cloud vendors, and c) providers will have clear guidance where interoperability and portability must exist within similar categories of cloud services.*

Illustrative example of why this requirement is not considered to be fully met at present: *The NIST cloud computing definition has identified three distinct categories of cloud service models: Software as a Service, Platform as a Service, and Infrastructure as a Service. Currently, consumers must seek to understand cloud services through a customized and product specific view presented by each service provider (understandably intended to differentiate products in the marketplace). Moreover, many vendors seek to establish categories of “cloud” services in addition to the three listed above, however it is not clear that proposed categories are unique and not already covered in the three primary services defined to date. Examples of proposed additions include Data as a Service, Network as a Service, Service as a Service and others. The result is a confusing landscape of possible cloud services that make it difficult for consumers to compare cloud services from an “apples to apples” perspective.*

Rationale: *In 2010, a NIST cloud computing reference architecture project team surveyed 11 existing cloud computing reference models and services proposed by cloud organizations, vendors, and federal agencies to see if there was any clear industry consensus. Analysis showed a wide disparity. A neutral common understanding and model is needed by customers in order to clearly and consistently understand how cloud services compare (i.e., apples to apples.) In November 2010, a NIST-hosted public working group explored proposed recommendation, and synthesized and leveraged this work through consensus to define a single neutral reference architecture. The reference architecture and taxonomy focus on the “what” as opposed to the “how” of implementation, and are not tied to a specific vendor implementation. In 2011 and 2012, industry participants validated the model by mapping it to their cloud services.*

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Encourage adoption of the NIST Reference Architecture by ISO/IEC JTC1, or any alternate neutral reference architecture through an international consensus-based standards body.	2012 ¹¹ – 2013
Categorize products using the NIST Reference Architecture ¹² to provide a consistent view of cloud services to USG agencies.	2012 – update periodically

¹¹ In 2012, the NIST reference architecture was accepted as an expert submission and is being used as a basis for the standards process.

2.5 Requirement 5: Frameworks to Support Federated Community Clouds

Industry and the USG need to develop frameworks to support seamless implementation of federated community cloud environments.

Why: *In Community Cloud deployment, infrastructure is shared by organizations that have common interests (e.g., mission, security requirements, and policy). In the case where a Community Cloud deployment model is not implemented in one (private cloud or public) environment which accommodates the entire community of interest, there is a need to clearly define and implement mechanisms to support the governance and processes which enable federation and interoperability between different cloud service provider environments to form a general or mission-specific federated Community Cloud.*

Illustrated Example of why this requirement is not considered to be fully met at present: *In the case of a Community Cloud deployed by a single Cloud Provider, the cloud PaaS layer can be used by developers to create applications. If developers establish common technical policies and credentials within that Community Cloud, they can use tools and management systems from different vendors, and connect applications to others using common PaaS facilities. However, in a federated multi-cloud environment with diverse cloud implementations and policies, the modules may need manual intervention to function together. Technical policies, credentials, namespaces, and trust infrastructure must be harmonized to support a Community Cloud that spans multiple service providers' physical environments.*

Rationale: *The importance of the Community cloud was clearly identified in the NIST-hosted Reference Architecture public working group. The architecture anticipated potential multi-cloud configurations such as Hybrid cloud or those topologies involving a Cloud Broker. It did not address the generalized notion of a federated Cloud Community. USG agencies, the National Security Telecommunications Advisory Committee, and the Open Grid Forum are examples of potential cloud adopters which have identified this as a high priority. The concept has been developed in earlier IT models such as the "GRID," where public and private sector research labs and universities make up a community of High-Performance Computing scientists. Federation techniques have been applied across GRIDs, data centers, and countries to create a "multi-GRID community logical GRID."*

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Define federated Community cloud requirements and scenarios.	2012 ¹³ – 2014
Identify how Hybrid Cloud and Cloud Broker elements described in the cloud Reference Architecture can be leveraged and harmonized.	2012 – 2013
Present analysis of GRID communities' applicability to federated cloud communities, including technology, trust infrastructure, & governance.	2012 - 2013
All stakeholders -- assess Intercloud efforts (e.g., Standards Developing	2012 - 2013

¹² NIST Special Publication 500-292, [NIST Cloud Computing Reference Architecture](#)

¹³ USG agencies have initiated development of federated clouds; these examples are being leveraged in the draft Volume III of this document which is currently under development.

Organizations) for applicability.	
-----------------------------------	--

2.6 Requirement 6: Updated Organization Policy that reflects the Cloud Computing Business and Technology model

Organizations need to review, revise, and develop policy in the context of the global business and technical model enabled by cloud computing and other enabling technologies.

Why: *While it may be developed in parallel, clearly defined organization policy is a prerequisite to security guidance and technical controls. Technology is appropriately used to implement policy, not to create it. However, if policy is not explicitly defined or does not reflect current and realistic data access capabilities, for example, the roles become reversed. Technology limits become (inappropriately) the default creator of policy as opposed an implementation tool. In the case of cloud, updated and transparent policy, which can be interpreted to a limited set of defined guidance and technical levels is essential. This is necessary and complementary to the development of technical solutions which satisfy security requirements AND allow differing policies to coexist side by side in a global environment irrespective of geographical location and sovereignty (requirement 3.) Organizations need to define policy that recognizes that reliance on the ability to enforce a legal framework and policy through physical location is insufficient for IT services delivered using the cloud model. In the absence of defined policy, organizations seek to informally achieve policy objectives through technical standard and product definitions. A possible end result is one where service providers are driven to artificially differentiate technological products and standards, resulting in technology stagnation as opposed to innovation. At a minimum the full potential of technology to foster universal world-wide quality of life improvements and a level international economic playing field will be stymied.*

Illustrative example of why this requirement is not considered to be fully met at present: *Historically differences in organization values, including but not limited to those at the sovereign nation level, related to such subjects as privacy and the free flow of information have been resolved by relying on physical data location in the context of geographical borders. In some cases organizations have responded with updated policy such as that categorized as “safe harbor”¹⁴ However, these models rely on point-to-point agreements and individual organization certifications which are overwhelmed by the volume and distribution of cloud service providers and global service options, as well as the rapid pace of technological change driven capabilities. Policies have not been developed which respond to the “anywhere anytime,” co-tenancy, unplanned demand levels, and utility characteristics of cloud.*

Rationale: *The TechAmerica Foundation issued recommendations in July 2011 that called for a “technology-neutral privacy framework...,” “Security and Assurance Frameworks... which are international...” and “...U.S. government ...willingness to trust cloud computing environments in other countries for appropriate government workloads.” The Business Software Alliance issued a global cloud readiness assessment in July 2012 that cited policy as a broad inhibitor of cloud adoption.*

¹⁴ E.g. a) 1998 On-Line Copyright Infringement Liability Limitation Act (OCILLA) effort to protect service providers on the Internet from liability for the activities of its users. Codified as [section 512 of the Digital Millennium Copyright Act \(DMCA\)](#); b) U.S.-EU Safe Harbor Framework, final documents issued by the United States and published in the Federal Register on July 24, 2000 and September 19, 2000, and by the European Commission on July 28, 2000.

US Government Cloud Computing Technology Roadmap, Volume I

<i>Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)</i>	<i>Proposed Target Date</i>
Define transparent policies that can be translated to specific levels of cloud computing security, privacy, and service criteria.	2012 - 2014

2.7 Requirement 7: Defined Unique Government Requirements and Solutions

The federal government needs to identify mandated requirements which are not clearly met in commercial cloud services, assess the extent to which the requirements are met, and define and communicate the gaps in technology and service offerings to industry.

Why: *In addition to the US federal policy related to cloud services adoption, USG agencies are also subject to other policy and regulatory requirements which are unique to government agencies. Government agencies must ensure that cloud services and products meet these policy and compliance requirements as well satisfy mission functionality requirements. Although agencies use commercial services to complete key elements of their mission, USG agencies cannot delegate inherently governmental federal authorities and public trust responsibilities to the private sector. USG institutions cannot mitigate risk through commercial means (e.g., financial penalties, insurance, litigation) to the same degree as private sector organizations. Failure to recognize and address government constraints may slow the adoption of cloud services.*

Illustrative example of why this requirement is not considered to be fully met at present: *OMB memo M-11-11¹⁵ reaffirmed the importance of the implementation of Homeland Security Presidential Directive (HSPD)-12¹⁶ and the need to move quickly to an authentication and access control mechanism which is defined and used government-wide. USG agency systems that are not “national security systems” as defined by 44 U.S.C 3542(b) (2) should be required to use Personal Identification Verification (PIV) cards as a way of authentication.¹⁷ This is an example of a requirement where it is necessary to identify and address technology gaps in order for USG agencies to authorize use of cloud services. However, equally important, and often overlooked, the US federal government has requirements for accessibility¹⁸, which are often not met in commercial cloud service offerings.*

Rationale: *Target USG Business Use Cases have identified cases where government requirement constraints can affect the way the services are designed and implemented and introduce the need for additional features. To expand USG use of cloud computing services, it is necessary to explicitly and objectively identify requirements not currently met in commercial cloud technologies and services, and to formulate strategies to supply missing functionality.*

<i>Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)</i>	<i>Proposed Target Date</i>
Identify regulatory factors that could affect cloud requirements, those which if unmet will prevent adoption by USG agencies, and cloud-based system features that satisfy these regulatory requirements.	2012 – ongoing
Develop technology and products to fill the gaps.	2012 – ongoing

¹⁵ <http://www.whitehouse.gov/sites/default/files/omb/memoranda/2011/m11-11.pdf>

¹⁶ http://www.dhs.gov/xabout/laws/gc_1217616624097.shtm

¹⁷ <http://csrc.nist.gov/publications/fips/fips201-1/FIPS-201-1-chng1.pdf>

¹⁸ Section 508 of the Rehabilitation Act (29 U.S.C. 794d), as amended by the Workforce Investment Act of 1998 (P.L. 105-220), August 7, 1998....REQUIREMENTS FOR FEDERAL DEPARTMENTS AND AGENCIES.-- ... (1) ACCESSIBILITY

2.8 Requirement 8: Collaborative Parallel “future cloud” Development Initiatives

Academia, industry, and the US and international governments need to define and begin work on “future cloud” development initiatives.

Why: *To date, innovation and technology for deploying Web-scale (nation-scale) clouds has been developed by industry. Much of the construction know-how is therefore not available in the public domain; the technology is considered to be intellectual property. However, government agencies have legislated, and public trust authorities and responsibilities that cannot be outsourced to private companies, including but not limited to responsibilities for ensuring that high security impact systems and data are protected, and that emergency and critical infrastructure public services are provided on a massive scale.¹⁹ Development of a demonstrable and practical technology knowledge base focused on state-of-the-art, nation-size clouds which are scalable and capable, and development of accessible standards and technologies, is needed to solve these nation-scale challenges. A focused set of cloud services and research would more rapidly lead to world-class cloud advancements to support critical national priorities and citizen services.*

Illustrative example of why this requirement is not considered to be fully met at present: *There are two aspects to this requirement. One is a long-standing set of inherently governmental functions and responsibilities. The other aspect is the private sector role in developing emerging technology. The core premise of this requirement is that there are cases where government requirements can and should not be privatized, and intellectual property and capabilities are appropriately maintained in the private sector. This highlights the need for partnership in applying emerging technology in the interest of the public good. For example, in the case of cloud construction and operation, one could envision a need to support more than 100,000 servers, spanning multiple data centers, and new challenges in network design.*

Rationale:

Government target business use cases have identified examples where cloud service providers could help to support applications of great benefit to the public. USG agencies see a need to provide geospatial data for public use in emergencies. A real-life proof-of-concept precedent was established through Japan’s response to the earthquake and tsunami that struck the Greater Tohoku region in March 2011.²⁰ In 2012, the government of Japan defined an objective to apply cloud computing to support emergency response. The Japanese agency NISC, NIST, and others are collaborating on the development of a cyber-physical cloud concept to combine cloud computing and physical device control to respond to emergencies such that resources, including robotic and automated mechanisms, could be rapidly deployed.

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Define scenarios to support testing state-of-the-art, interoperable, nation-size clouds.	2012 – 2016
Define project concepts. Identify likely technical and standards challenges.	2012 – 2017
Define conceptual research strategy.	2012 - 2015

¹⁹ National Security Telecommunications Advisory Committee, **NSTAC Report to the President on Cloud Computing**, May 15, 2012.

²⁰ *Responding to the Greater Tohoku Disaster, The Role of the Internet and Cloud Computing in Economic Recovery and Renewal*, Internet Economy Task Force, 2011.

2.9 Requirement 9: Defined & Implemented Reliability Design Goals

Industry needs to define and implement reliability design goals, best practices, and related measurement and reporting processes. (*interoperability, portability, and security technology*)

Why: *As USG agencies increase their use of cloud computing to provide essential public services, it is essential that industry be able to ensure that design flaws do not result in catastrophic failures or significant outages over large regions or for extended periods of time.*

Illustrative examples of why this requirement is not considered to be fully met at present: *Cloud Builders create mechanisms to compensate for component failures and deliver High Availability, but the news has highlighted major cloud provider outages. In several cases, cloud providers suffered failures or design flaws which affected the accessibility of cloud-based services for many subscribers. In April 2011, an erroneous network reconfiguration triggered a failure, followed by a cascade of recovery events and subsequent failures, and a lengthy outage. In May 2011, a sequence of cloud outages and software errors led to email delays. During June and July 2011, the same cloud provider suffered outages that disabled services. In August 2011, an intense lightning storm overloaded a power transformer; cloud services were unavailable for hours. In August 2011, a cleanup software bug resulted in customers losing backup data.*

Rationale: *Cloud Computing exemplifies reliability scenarios that are not found in traditional computing and communications architectures. In traditional computing architectures, there is an affinity between the application and the specific hardware on which it runs; high-availability strategies are implemented per-platform, usually through hardware redundancy. In cloud computing, the application and the hardware have less affinity because of virtualization. The economics of hardware redundancy are different in cloud environments in that redundancy within a cloud must be supported by a cloud provider (because users cannot reliably know workload-hardware bindings), and cross-cloud redundancy can trigger additional usage fees. Due to scale, a statistically rare failure event may be a common occurrence in a cloud; clouds compensate with redundancy implemented by cloud software.*

Working with industry and academia, government researchers have identified needs to model, understand, and predict global behavior and ensure reliability in large distributed systems, such as the Internet and computational grids. USG researchers²¹ uncovered design flaws in open-source cloud software that could result in significant resource leakage when systems operating that software are exposed to simple malicious attacks.

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Formulate and publish best practices on achieving reliability.	2012 – 2014
Develop a consensus process to measure and report industry-wide cloud reliability information to assess current and future cloud reliability.	2012 – 2017
Define research methods for real-time measurement and monitoring to predict onset of catastrophic failure in cloud systems, and tools to identify failure vulnerabilities.	2012 - 2015

²¹ Dabrowski, C., and K. Mills. "VM Leakage and Orphan Control in Open-Source Clouds." Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on. IEEE, 2011.

2.10 Requirement 10: Defined & implemented Cloud Service Metrics

Industry needs to establish Cloud Service Metrics, including Standardized Units of Measurement for Cloud Resources.

Why: In utility industries, the notion of units of measurement is fundamental to buying and selling service. Benchmarking is used in traditional computing system operations to determine the performance of system infrastructure such as hardware and operating systems, and for key application platform elements such as database servers and Web servers. However, in the case of cloud computing service delivery, which uses a utility model, IT resources are supplied as abstracted services, often characterized as Infrastructure as a Service or Platform as a Service. For example, networking and storage are often provided as abstracted services. Abstracted services can be set to run fast or slow, to be small or large, and to be as reliable as desired (subject to underlying technology constraints). Service consumers pay for a “quantity” and a “quality” of the service, which is metered by a cloud computing system. Consumers need to be able to precisely specify and receive services.

Illustrative example of why this requirement is not fully met at present: In contrast to the precision with which we categorize units of measurement in electricity, light, or fuels, cloud computing measurements are relatively imprecise. Furthermore, there is no common collection of vendor agreed-upon specific terms. For example, while one provider uses an informal “Elastic Compute Unit,” it is imprecise and does not account for workload mix or speed to memory. The characteristics of storage and access to storage over a network vary. Service providers have not defined and applied standardized units of measurement that can be specified in Service-Level Agreements and interoperability exchanges. Therefore, consumers cannot determine and request cloud services as a utility with a high degree of predictability, and cannot achieve maximum cost-effectiveness in cloud computing service application.

Rationale: The USG Target Business Use Case, Reference Architecture, and the public security working groups have all identified this requirement. IaaS services include processing, memory, network, and storage. Considering only storage, for example, a Gigabyte is not the only unit of measurement. There are several “flavors” of storage services: structured and unstructured, replicated and non-replicated, fast-access and slow-access. Furthermore, IaaS attributes have additional dimensionality, such as variation in access speed or processor speed. In other utility industries, the notion of units of measurement is fundamental to creating an economy. This requirement will yield a portfolio of formal Standards for units of measurement in cloud computing, which will be used in a number of ways, from SLA specifications to interoperability exchanges.

Recommended Priority Action Plans (candidates for voluntary self-tasking by cloud computing community stakeholders)	Proposed Target Date
Specify and Standardize the Units of Measurement for cloud services, seeking public comment and collaboration.	2012 – 2013
In parallel, incorporate Cloud Service Units of Measurement consistently in Service-Level Agreements.	2012 - 2013

3 Other Considerations and Observations

The following is a small subset of subjects with which the scope of cloud computing has a Venn Diagram-like relationship. Cloud computing is not a subset or a superset of these topics. More specifically, while these topics inform the NIST collaborative initiative to build a *USG Cloud Computing Technology Roadmap*, in their entirety they are outside of the scope of this effort. The topics are listed here to make the point that work in these areas is recognized as being highly interdependent with and essential for overall effectiveness of the roadmap effort.

3.1 Regarding Academia, Industry, Standards Organizations, and Government Collaboration

While the last several years have seen an increase in cloud deployment and benefit, and there are a large number of cloud community stakeholders accomplishing valuable work in advancing cloud computing standards, guidance, and technology, the rapid pace of cloud computing evolution (which has been characterized as “building the plane while we are flying it”) is still such that the community needs to work even harder to explicitly leverage our efforts and get ahead of the curve.

For example, there are many approaches to cloud computing standards. In some cases, standards are being developed in consensus-driven working groups, but are not being applied in implementations. In other cases, non-standardized implementations evolve in parallel, but do not transition to the point where the work is leveraged through formal Standards Developing Organizations. One example of a general benefit that would ensue from aggressively pursuing cloud computing standards is that US government agencies procuring services would be positioned to specify standards, as opposed to specific cloud provider services or products. This would improve cost-effectiveness for the taxpayer and level the playing field for the private sector consumers and service providers.

Collaboration is a productive, but unstructured process that is often driven from the bottom up in the sense that developers and adopters have individual mission, schedule, and resource objectives and constraints. Despite these differences, it is clear that there is much convergence in principle. International technical exchanges²² and reports²³ illustrate this point. Priorities defined explicitly through international conferences hosted by the governments of Canada, China, and the European Commission and standards organizations,²⁴ but not exclusively there, include: standards, a level playing field that supports technical innovation, interoperability and open interfaces, a desire to harness the power of cloud to improve public services, a need for improved understanding of cloud computing by policy makers, guidance to architects and engineers, and conformity assessments and testing. An example of a practical collaboration is a mapping exercise that was initially completed by the EC Standards and Interoperability for eInfrastructure implementation initiative (SIENA) project to look for commonality and synergism between the NIST

²² U.S.-Japan Economic Harmonization Initiative, ICT –IPR Working Group, Washington, D.C., July 2011.

²³ *Exploring the Future of Cloud Computing: Riding the Next Wave of Technology Driven Transformation*, World Economic Forum, 2010.

²⁴ EC-ETSI workshop “Standards in the Cloud: a transatlantic mindshare”, Sophia-Antipolis, France, September 28-29, 2011.

technical use cases and Cloud Usage Scenarios with European eScience developments.²⁵ In 2012, the effort continues with technical exchange between the government of Japan, ETSI in coordination with the EC, and NIST to validate the inventory of standards relevant to cloud computing.

3.2 Interdependency with Cyber Security initiatives

As mentioned in Section 2.2, while cloud computing security requirements are not unique in their entirety or separate from general IT security requirements, the cloud computing environment presents certain unique security challenges resulting from the cloud's very high degree of outsourcing, dependence on networks, sharing (multi-tenancy), and scale. Several initiatives that relate to these challenges are:

The Department of Homeland Security Continuous Asset Evaluation, Situational Awareness, and Risk Scoring (CAESARS) project, which is providing an architecture for dynamic system monitoring and reporting;

The Security Content Automation Protocol (SCAP) initiative at NIST, which provides specifications for expressing security configurations and events, event management, and incident handling;

The National Science Foundation Future Internet Architectures initiative which is developing Internet architectures to provide advanced security and reliability in the context of emerging Internet usage patterns; and

The Federal Information Security Management Act (FISMA). In accordance with the Act, Federal Information Processing Standards (FIPS) 200 and NIST Special Publication 800-53 (periodically updated) provide baseline security controls and guidance for federal information systems.

The Federal Risk and Authorization Management Program (FedRAMP) is an internationally recognized effort originally conceived by the US CIO Council sponsored Cloud Security Working group in 2010. NIST serves as a technical advisor to the General Services Administration which executes the FedRAMP program under the cognizance of the Office of Management and Budget, United States Chief Information Officer and staff, and the US Federal CIO Council. FedRAMP became operational in 2012.

Security requirements are tightly coupled with interoperability and portability, reliability, and maintainability, which also include considerations which are specific to the cloud computing model. One example of general security work that directly relates to security requirements in the cloud environment is the ability to securely migrate virtual machines between dissimilar organizations or hardware/software environments. In other words, such work aims to provide confidence that before a virtual machine is created in a new physical environment, that environment satisfies the technical policy controls specific to the application and data. Other general areas include authentication techniques such as multifactor authentication with tokens, applied cryptography, and software assurance techniques (e.g., testing and analysis) needed to build confidence that logical boundaries implemented in cloud systems are sufficiently strong to provide security.

3.3 Interdependency with emerging Big Data technology

Big data has emerged as a technology term and trend that is complementary to and considered to be equally as transformational as the cloud computing model. Cloud Computing subject matter experts

²⁵ OASIS International Cloud Symposium, October 2011, 2012.

consider cloud to be an enabler of big data capture, storage, analysis, sharing and management. Big data subject matter experts commonly refer to cloud computing as being indistinguishable from big data.

Just as cloud computing struggled with definition early in its adoption, and similarly was represented as an “old” or “new” capability depending on the perspective of those defining it, big data as a concept is the focus of definition and framing discussions. In 2012, the US federal government identified a Big Data Research and Development Initiative to explore how big data can be used to address government requirements. In planning its Cloud Computing Forum & Workshop Outreach event (January 2013), NIST expanded the agenda to explore the convergence of cloud computing and big data, with the expectation of informing its respective planning and program efforts.

3.4 Organizational Policy

The perspective presented in this document is that technology can be used to inform organization policy, and can be used to help implement organization policy, but is not one and the same as organization policy. As highlighted in Section 2.6, it is necessary to have technical solutions which allow differing policies to coexist side by side in a global environment irrespective of geographical location and sovereignty. If not, the benefits of large-scale interoperability and portability for cloud workloads will not be realized. Moreover, the ability to bridge policy differences is essential for maintaining service while policies evolve.

This capability of abstracting technical solutions, so they can be used to implement sovereign policy decisions, but are not prescriptively constrained by specific policy decisions, is essential to universal implementation of the security requirements and associated controls which are critical to ensuring privacy rights and global Ecommerce. This same capability is essential in the development of common commercial application terms of Service-Level Agreements, including commonality of pricing unit definitions, customer protective contract terms, liability ownerships, audit rights, exit provisions, and business continuity.

3.5 Interdependency with Other National Priority Initiatives

The Cloud Computing model is clearly an enabler of national priority initiatives such as Health IT and Smart Grid, and is enabled by programs such as National Strategy for Trusted Identities in Cyberspace (NSTIC). There are tremendous win-win opportunities if we can quickly move toward integrated development of consensus-based cloud computing standards.

An intuitive illustrative target case is the application of cloud computing as an enabler to improve health care for veterans. There is great focus on government security requirements, but other government requirements, such as Section 508²⁶ compliance, are often overlooked. One of the strengths of the cloud model is the anytime/anywhere deployment on a broad variety of end-devices. This would be a key advantage in addressing disability access. Physical disabilities, post-traumatic stress disorder, or depression can make downloading 508-compliant profiles to individual devices challenging. One way to help address this is to deploy Health IT systems using a common profile that defines a preference specific to each individual. However, the benefit of achieving this scenario applies much more broadly than satisfying a government requirement or supporting a specific interest group. The same solution could be

²⁶ Section 508 of the Rehabilitation Act (29 U.S.C. 794d), as amended by the Workforce Investment Act of 1998 (P.L. 105-220), August 7, 1998....REQUIREMENTS FOR FEDERAL DEPARTMENTS AND AGENCIES.-- ... (1) ACCESSIBILITY

applied to improve the ability of parents, educators, and other responsible parties to screen Internet-accessible content by minors. Over and above these specialized requirements, the same capability could be leveraged to improve convenience and ease of use for all cloud service users. These requirements can be met without applying the cloud model – the cloud computing model is simply an enabler that has the potential to accelerate this capability. This concept intuitively demonstrates the relationship between roadmap requirements and practical implementation. A simple test of the capabilities described above would require integrated security standards to secure the data and protect the privacy of the profile as well as the data, data portability, and interoperability at the software, platform, and infrastructure levels of cloud.

Information security is naturally a critical factor for widespread adoption of Cloud Computing. For government users, particularly early adopters, security fears are front and center. In addition to data confidentiality, integrity, and availability, the need for trusted identities and secure and efficient management of these identities while users' privacy is protected is a key element for the successful adoption of any cloud solution. Augmenting security technologies and best practices, NSTIC could enhance security and privacy for cloud services. NSTIC defines a mean to create a secure, trusted Identity Ecosystem that is capable of establishing a user-centric privacy protection for any Cloud Ecosystem. It is generally acknowledged that the use of passwords does not provide optimal security or assurance. The NSTIC Strategy²⁷ calls for the development of interoperable technology standards and policies – the “Identity Ecosystem”²⁸ – where individuals, organizations, and underlying infrastructure – such as routers and servers – can be authoritatively authenticated. The mechanisms employed by an Identity Ecosystem are structured in a robust framework comprised of the overarching set of interoperability standards, risk models, privacy and liability policies, requirements, and accountability mechanisms. Individuals will be able to validate their identities and then securely access the Identity Ecosystem. Within NSTIC’s trusted framework of defined security requirements based on risk and sensitivity, Cloud services will be more securely supported. The objective is more than lowering cost and increasing access; it also supports interoperability, portability, and security.

3.6 Education of Technical Staff and Cloud Consumers

Major transformation using Cloud Computing technology requires business and technical stakeholders to work together. One of the major impediments to cloud computing adoption is lack of a common understanding of business and technical benefits. Standards and interoperability make inter-agency integration possible resulting in better business outcomes. Business and technical objections to cloud computing adoption can be overcome by educating technical staff on the business value of cloud computing and business users of the technical capabilities now possible that were not available before.

²⁷ NSTIC Strategy, “Enhancing Online Choice, Efficiency, Security, and Privacy”, The White House, Washington, April 2011, http://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf. The notion of an “Identity Ecosystem” is drawn from the above reference. There is no intent in the USG Cloud Computing Technology Roadmap to endorse or advocate the establishment of the “Identity Ecosystem”. Comments regarding the “Identity Ecosystem” should be referred to the NIST NSTIC Program: <http://www.nist.gov/nstic/about-netic.html>

²⁸ NSTIC Strategy, “Enhancing Online Choice, Efficiency, Security, and Privacy”, The White House, Washington, April 2011, http://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf. The notion of an “Identity Ecosystem” is drawn from the above reference. There is no intent in the USG Cloud Computing Technology Roadmap to endorse or advocate the establishment of the “Identity Ecosystem”. Comments regarding the “Identity Ecosystem” should be referred to the NIST NSTIC Program: <http://www.nist.gov/nstic/about-netic.html>

US Government Cloud Computing Technology Roadmap, Volume I

With the technology continuing to mature, there needs to be an effort to constantly keep interested parties up to date on technology changes and legal issues to lower barriers to cloud adoption.

4. Progress and Next Steps

This document marks the completion of the second phase of the NIST Cloud Computing Program.

The first phase of the NIST Cloud Computing program and initiative to collaboratively build a *USG Cloud Computing Technology Roadmap* completed in November 2011, and marked by the draft release of SP 500-293 US Cloud Computing Technology Roadmap, volumes I and II.

Over the past year, NIST incorporated over 200 public comments it received in response to the November 2011 draft, and re-issued volume I as this final special publication. In working with its public and private sector partners from academia, industry, standards organizations, US federal, state and local government agencies, and the international community, NIST was able to achieve major Phase 2 objectives, including:

- Validating the Phase 1 Reference Architecture (SP 500-292) through cloud service provider examples of categorized services, and working with cloud stakeholders to establish a repository of the mapped vendor services to support USG and others in comparing cloud service offerings; and through formal standards organization activities, including but not limited to ISO and I-TUT working groups;
- Continuing to identify high-priority interoperability, portability, and security requirements which must be met for USG agencies to accelerate the adoption of the cloud computing model; continuing to assess standards, guidance, and technology that must be in place to meet these requirements, and recommending Priority Action Plans (PAPs) for voluntary self-tasking by the cloud stakeholder community, to support standards, guidance, and technology advancement;
- Working with cloud stakeholders to identify efforts which satisfy the objectives of the PAPs, assessing and communicating the extent to which the requirements are satisfied, and defining processes to leverage these efforts to support the USG adoption of cloud computing;
- Identifying the subset of PAP objectives which are consistent with NIST core mission standards, guidance and research activities; developing NIST PAP plans, and executing those plans; and communicating the progress accomplished by the PAP projects toward the USG roadmap requirements through the June 2012 NIST Cloud Computing Forum & Workshop event;
- Integrating these strategic activities with NIST tactical program projects and working groups; continuing to deliver special publications, technical guidance, and support collaborative Web-based tools to support these tactical efforts;
- Defining and tracking measures and metrics to assess program effectiveness, the most significant being the federal balanced scorecard objectives for the reference architecture and FedRAMP technical advisory functions;
- Continuing outreach activities including the NIST Cloud Computing Forum & Workshop series to calibrate and leverage NIST efforts with the broader stakeholder community; and
- Analyzing and assessing the technical work completed through these efforts, and applying this analysis to revise the *USG Cloud Computing Technology Roadmap* on a periodic basis.

The first two phases of the program executed and achieved results consistent with the program strategy initially defined May through October 2010, and the program time line presented in November 2010.

US Government Cloud Computing Technology Roadmap, Volume I

As intended, the roadmap document has served as a practical mechanism to integrate and present analysis, findings, and useful technical artifacts generated through the NIST Cloud Computing program public and federal working groups, and internal NIST projects.

However, the roadmap, and PAP project efforts related to the roadmap, also provided an opportunity over a two year period to assess progress and effectiveness. The NIST assessment is that the collaborative approach has been effective, and the initiative has met the goal of technically advancing the cloud computing model – particularly in its target area of interoperability, portability, and security standards, guidance, and technology requirements.

The NIST assessment is based on a 30-month continued level of engagement with the cloud community in public working groups and NIST Cloud Computing Forum and Workshop events. Hundreds of individuals and organizations are registered working group members, and the NIST-hosted cloud forum events have been registered to capacity. In terms of results, the “useful information for cloud adopters” available publically on the NIST Cloud Computing Web site and special publications produced from NIST projects and working groups are widely referenced and used. These are summarized in Volume II of the roadmap document. The most widely recognized work, after the *NIST Cloud Computing Definition*, SP 800-145(Draft), is the *NIST Cloud Computing Reference Architecture*, SP 500-292, which was first issued in September 2011, and continues to be refined and used as the basis for developing a standardized reference architecture by international standards bodies, and by US government agencies and industry for its intended purpose of categorizing cloud services so that government agencies and others can compare cloud services from different providers more easily. The major work completed in 2012 was in the area of service level agreements and refinement of the security components of the reference architecture.

NIST also bases its assessment on the review and support for the *US Government Cloud Computing Technology Roadmap Volume 1, Release 1.0, (DRAFT) High-Priority Requirements to Further USG Agency Cloud Computing Adoption* by the representatives designated by the US Federal CIO Council to participate in the Federal Cloud Computing Standards and Technology Working Group. Confirmation of the priorities presented here by a broad sample of representatives of USG organizations who are responsible for deploying IT to support agency missions reinforces the conclusion that “we” – NIST and its cloud community collaboration partners – “got it right” in terms of the objectives for the effort and the roadmap. Equally important, of the 200 public comments received in response to the roadmap, there were no disagreements in principle or challenges to the overall value of the work – comments were refinements, additions, and in some cases, requests that scope be added to the effort. Finally, one of the most basic measures of the value of work is whether it is used – NIST is happy to report that the roadmap and technical work completed in partnership with the cloud community is widely used and referenced, not only in the US, but broadly in the international community.

Given this assessment, the following section presents the current thinking, strategy, and plan for the NIST Cloud Computing program and USG Cloud Computing Technology Roadmap initiative to continue to leverage and assess the roadmap progress going forward, in support of the overall goal of supporting USG agencies in the secure and effective deployment of cloud computing.

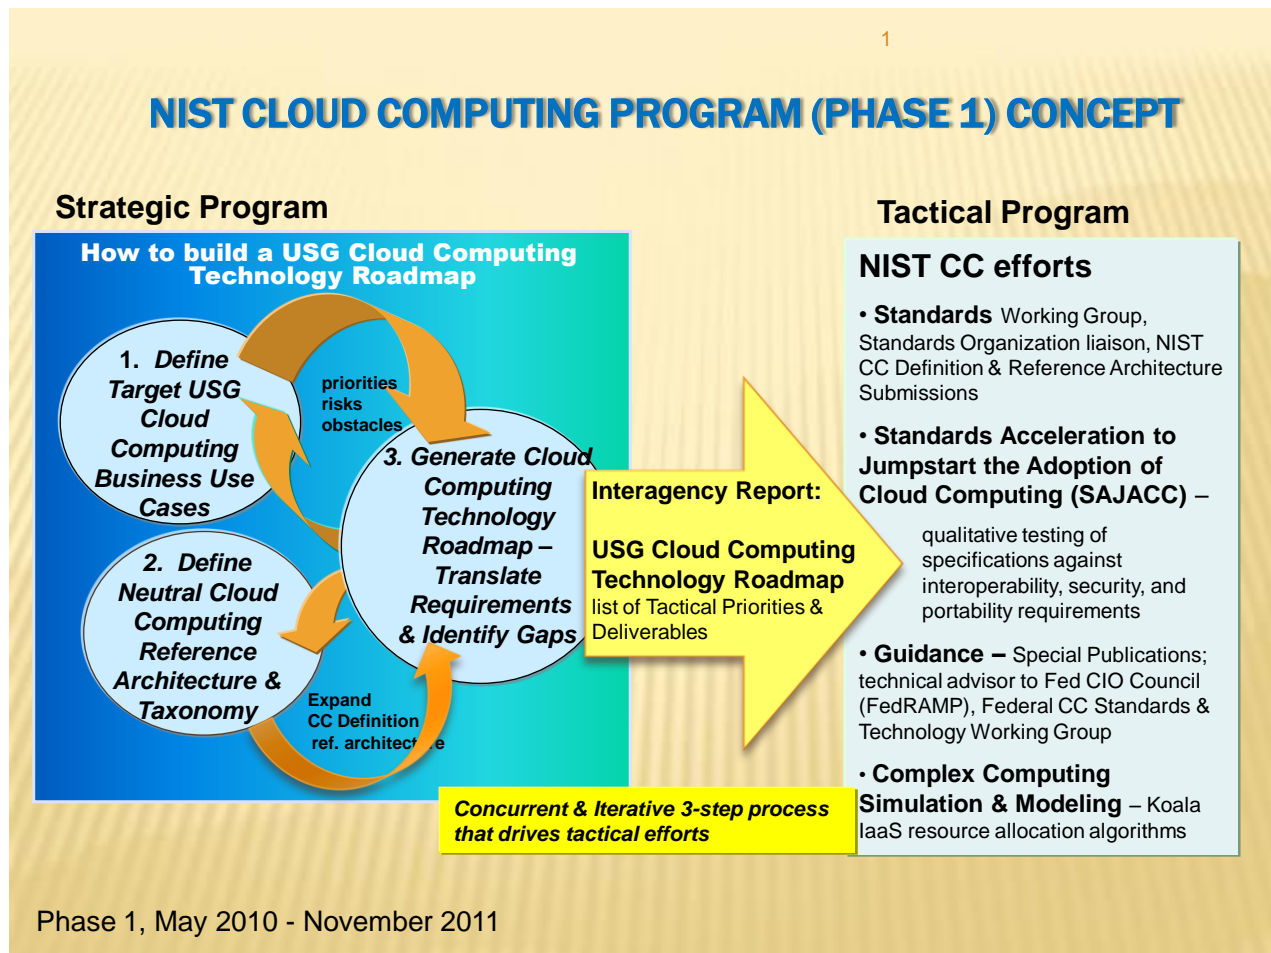
The expectation is that the NIST Cloud Computing program and USG Cloud Computing Technology Roadmap initiative has established a baseline of consensus requirements that must be met to accelerate

cloud adoption, and that the focus will continue to be the work of NIST' external partners as it relates to each of the respective 10 priority requirements.

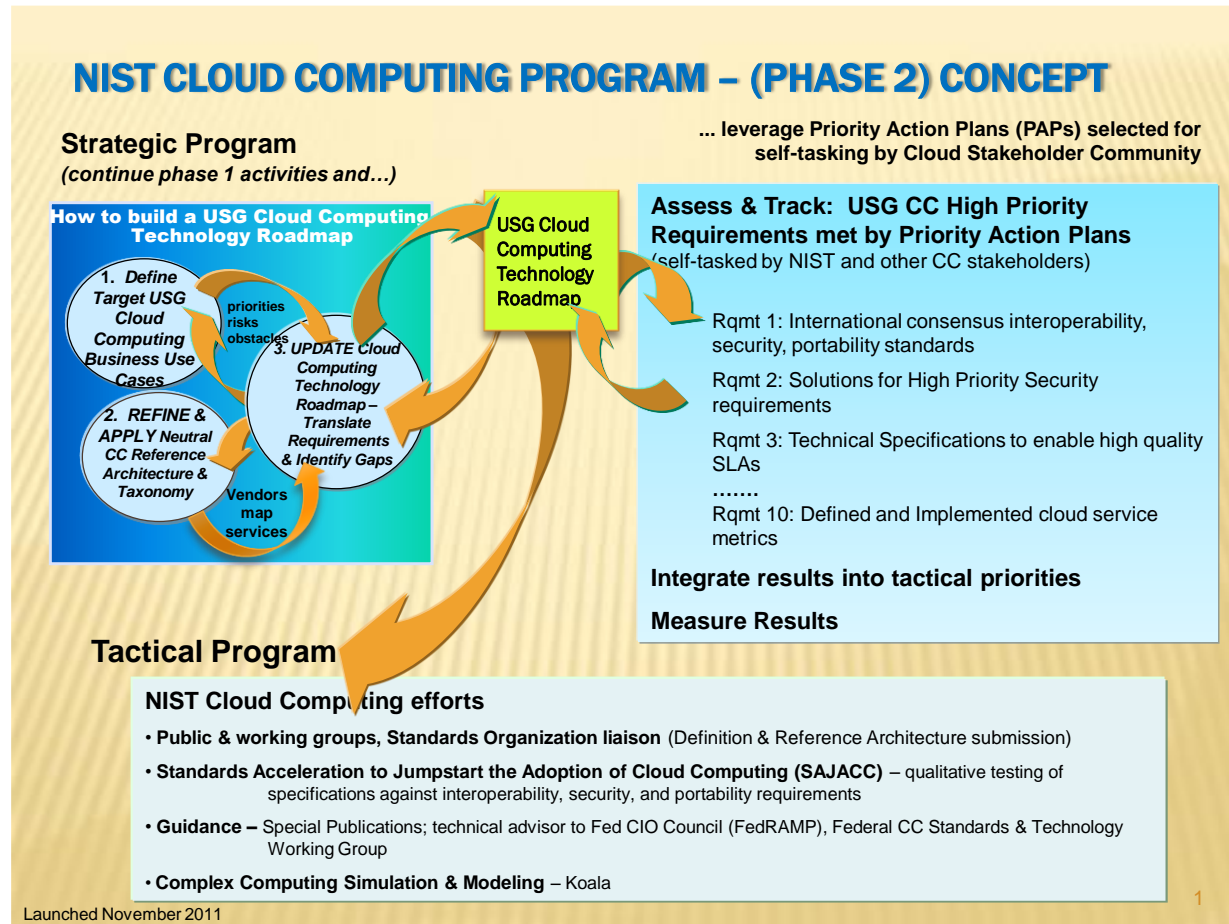
The expectation is that the program will continue its presence through outreach activities and interactions with other USG and international stakeholders, and track progress towards the priorities presented in this document.

4.1 NIST Cloud Computing Program Future Phases

For context, the following diagram revisits Phase 1 of the NIST Cloud Computing program. Phase I effectively established and integrated three strategic processes, public working groups, and NIST cloud efforts to develop the first USG Cloud Computing Technology Roadmap, and help NIST to prioritize its internal projects.



Phase 2 of the NIST Cloud Computing program continued the Phase 1 scope and activities, but shifted focus to introduce new strategic activities to leverage the *USG Cloud Computing Technology Roadmap* produced in Phase 1.



Future phases of the program are planned to include:

- Applying the USG Cloud Computing Business Use Case template to support USG development of agency mission use cases;
- Leveraging this effort to complete and issue the roadmap Volume III: *Technical Considerations for USG Cloud Computing Deployment Decisions*;
- Continuing to validating the Reference Architecture (SP 500-292) through cloud service provider examples of categorized services, and working with cloud stakeholders to establish a repository of the mapped vendor services to support USG and others in comparing cloud service offerings;
- Continuing to identify high-priority interoperability, portability, and security requirements which must be met for USG agencies to accelerate the adoption of the cloud computing model;

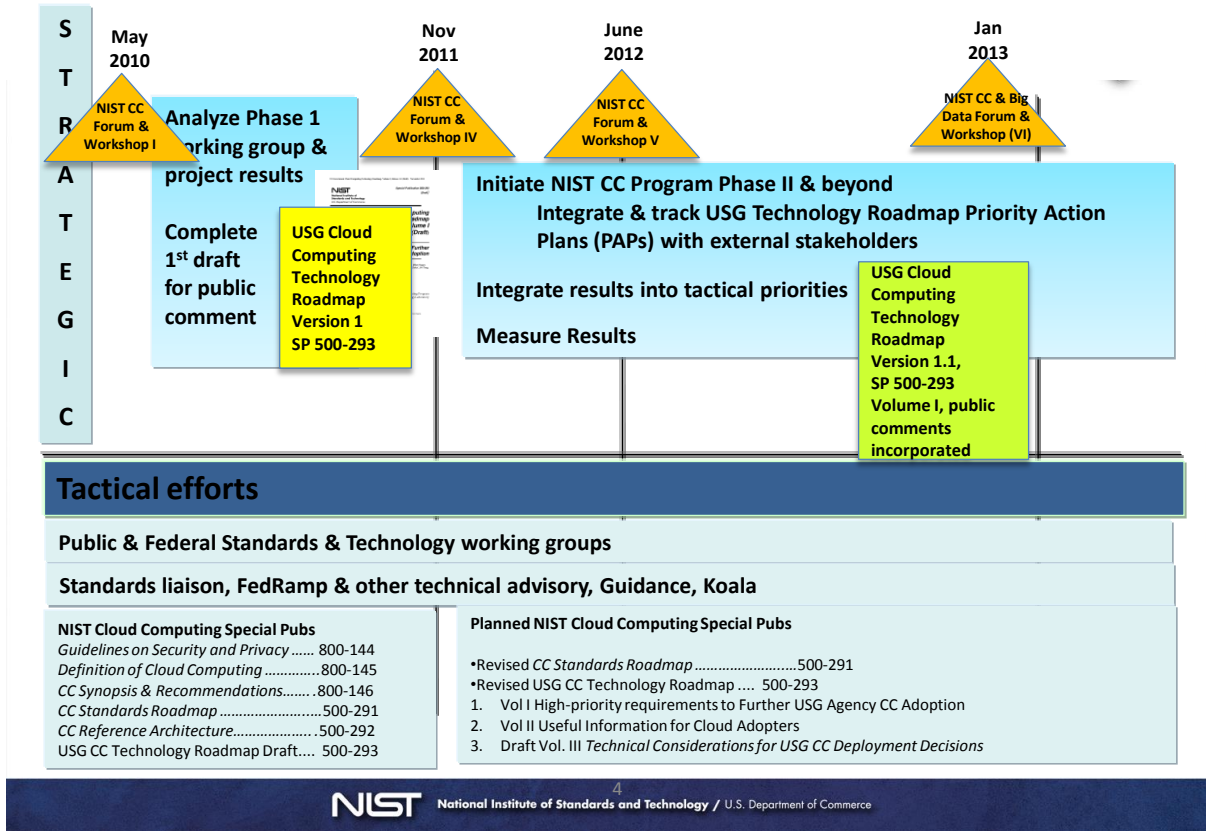
US Government Cloud Computing Technology Roadmap, Volume I

continuing to assess standards, guidance, and technology that must be in place to meet these requirements, and recommending Priority Action Plans (PAPs) for voluntary self-tasking by the cloud stakeholder community, to support standards, guidance, and technology advancement;

- Working with cloud stakeholders to identify efforts which satisfy the objectives of the PAPs, assessing and communicating the extent to which the requirements are satisfied, and defining processes to leverage these efforts to support the USG adoption of cloud computing;
- Identifying the subset of PAP objectives which are consistent with NIST core mission standards, guidance and research activities; developing NIST PAP plans, and executing those plans;
- Integrating these strategic activities with NIST tactical program projects and working groups; continuing to deliver special publications, technical guidance, and support collaborative Web-based tools to support these tactical efforts;
- Defining and tracking measures and metrics to assess program effectiveness;
- Continuing outreach activities including the NIST Cloud Computing Forum & Workshop series to calibrate and leverage NIST efforts with the broader stakeholder community; and
- Analyzing and assessing the technical work completed through these efforts, and applying this analysis to revise the *USG Cloud Computing Technology Roadmap* on a periodic basis.

4.2 Summary of Time Line and Deliverables

NIST COMPUTING PROGRAM TIMELINE



Appendix A: USG Interagency partners and contributors

Bruce Beckwith, Department of Energy

Kathy Conrad, Principal Deputy Associate Administrator, General Services Administration

Earl Crane, Department of Homeland Security, Information Security and Identity Management Committee (ISIMC)

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

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

Audrey M. Hogan, Tennessee Valley Authority

Dr. Prabha N Kumar, Special Assistant, Department of Defense, OCIO

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

Mr. James Ramskill, Office of the Director of National Intelligence

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

Lew Sanford Jr., DCS-OESAE, Social Security Administration (with other SSA participants)

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

Param Soni, Environmental Protection Agency

Gerald L. Smith, Department of Defense and OASIS

Peter Tseronis, Chief Technology Officer, Department of Energy

Special Publication 500-293

**US Government Cloud Computing
Technology Roadmap
Volume II**

Useful Information for Cloud Adopters

*Lee Badger, Robert Bohn, Shilong Chu, Frederic de Vault,
Mike Hogan, Michaela Iorga, Viktor Kauffman, Fang Liu, Jian Mao,
John Messina, Kevin Mills, Eric Simmon, Annie Sokol, Jin Tong,
Fred Whiteside and Dawn Leaf*

*This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.SP.500-293>*

NIST

**National Institute of
Standards and Technology**

U.S. Department of Commerce

US Government Cloud Computing Technology Roadmap, Volume II
Useful Information for Cloud Adopters

This page left intentionally blank

US Government Cloud Computing Technology Roadmap, Volume II

Useful Information for Cloud Adopters

NIST Special Publication 500-293

US Government Cloud Computing Technology
Roadmap Volume II

Useful Information for Cloud Adopters

Lee Badger, Robert Bohn, Shilong Chu, Frederic de
Vaulx, Mike Hogan, Michaela Iorga, Viktor Kauffman,
Fang Liu, Jian Mao, John Messina, Kevin Mills,
Eric Simmon, 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

This publication is available free of charge from:
<http://dx.doi.org/10.6028/NIST.SP.500-293>

October 2014



U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Acting Under Secretary for Standards and Technology and Acting Director

US Government Cloud Computing Technology Roadmap, Volume II

This page left intentionally blank

US Government Cloud Computing Technology Roadmap, Volume II

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

Natl. Inst. Stand. Technol. Spec. Publ. 500-293, 98 pages (October 2014)

This publication is available free of charge from:

<http://dx.doi.org/10.6028/NIST.SP.500-293>

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.

US Government Cloud Computing Technology Roadmap, Volume II

Acknowledgements

The authors, David Bernstein, 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, Michaela Iorga, John Messina, Kevin Mills, Eric Simmon, 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, and Peter Mell of NIST for providing technical input and detailed editorial support.

Table of Contents

1	INTRODUCTION.....	1
1.1	NIST CLOUD COMPUTING PROGRAM BACKGROUND	1
1.2	NIST CLOUD COMPUTING PROGRAM VISION	2
1.3	INTENDED AUDIENCE AND USE.....	2
1.4	DOCUMENT ORGANIZATION	3
2	NIST CLOUD COMPUTING DEFINITION AND REFERENCE ARCHITECTURE	5
2.1	REVISITING THE DEFINITION	5
2.2	NIST CLOUD COMPUTING REFERENCE ARCHITECTURE	7
2.2.1	CONCEPTUAL MODEL.....	8
2.2.2	CLOUD COMPUTING ACTORS.....	9
2.2.3	ARCHITECTURE COMPONENTS	15
2.3	NIST CLOUD COMPUTING TAXONOMY	17
3	CLOUD COMPUTING USE CASES AND REQUIREMENTS	18
3.1	TARGET BUSINESS USE CASE AND HIGH-LEVEL REQUIREMENTS	18
3.1.1	BUSINESS USE CASE TEMPLATE	20
3.1.2	BUSINESS USE CASE SUMMARIES.....	20
3.1.3	BUSINESS USE CASE ANALYSIS	24
3.2	SAJACC USE CASES AND TECHNICAL REQUIREMENT	41
4	CLOUD COMPUTING STANDARDS AND GAP ANALYSIS.....	43
4.1	CLOUD COMPUTING STANDARDS.....	44
4.2	CLOUD COMPUTING STANDARDS GAPS AND USG PRIORITIES	44
4.3	ACCELERATING THE DEVELOPMENT THE USE OF CLOUD COMPUTING STANDARDS	46
5	HIGH-PRIORITY SECURITY REQUIREMENTS	48
5.1	UNDERSTANDING SECURITY IN THE CLOUD CONTEXT.....	49
5.1.1	CLOUD SERVICE MODEL PERSPECTIVES.....	49
5.1.2	IMPLICATIONS OF CLOUD DEPLOYMENT MODELS	49
5.1.3	SHARED SECURITY RESPONSIBILITY.....	49
5.1.4	DEVELOPING SECURITY ARCHITECTURE FOR CLOUD SYSTEMS.....	50
5.2	CHALLENGING SECURITY REQUIREMENTS AND RISK MITIGATIONS	50
5.3	PROCESS-ORIENTED REQUIREMENTS.....	51
5.3.1	NIST SP 800-53 SECURITY CONTROLS FOR CLOUD-BASED INFORMATION SYSTEMS.....	51
5.3.2	CLOUD AUDIT ASSURANCE AND LOG SENSITIVITY MANAGEMENT.....	52
5.3.3	CLOUD CERTIFICATION AND ACCREDITATION.....	56

US Government Cloud Computing Technology Roadmap, Volume II

5.3.4 *NEEDED ELECTRONIC DISCOVERY GUIDELINES*.....56

5.3.5 *NEEDED CLOUD PRIVACY GUIDELINES*57

5.3.6 *CLARITY ON CLOUD ACTORS SECURITY ROLES AND RESPONSIBILITIES*59

5.3.7 *TRUSTWORTHINESS OF CLOUD OPERATORS*.....60

5.3.8 *BUSINESS CONTINUITY AND DISASTER RECOVERY*.....61

5.3.9 *TECHNICAL CONTINUOUS MONITORING CAPABILITIES*.....62

5.4 FOCUSED TECHNICAL REQUIREMENTS64

5.4.1 *VISIBILITY FOR CONSUMERS*.....64

5.4.2 *CONTROL FOR CONSUMERS*65

5.4.3 *DATA SECURITY*.....67

5.4.4 *RISK OF ACCOUNT COMPROMISE*.....69

5.4.5 *IDENTITY CREDENTIAL AND ACCESS MANAGEMENT (ICAM) AND AUTHORIZATION*70

5.4.6 *MULTI-TENANCY RISKS AND CONCERNS*.....72

5.4.7 *CLOUD-BASED DENIAL OF SERVICE*74

5.4.8 *INCIDENT RESPONSE*.....75

6 SUMMARY AND NEXT STEPS..... 76

7 APPENDIX A – SERVICE LEVEL AGREEMENT (SLA) TAXONOMY AND METRICS 78

8 APPENDIX B – RELIABILITY RESEARCH IN CLOUD-BASED COMPLEX SYSTEMS 82

9 APPENDIX C – USEFUL REFERENCES..... 84

List of Figures

Figure 1: The Conceptual Reference Model	9
Figure 2: Cloud Broker Interactions.....	13
Figure 3: Intermediary Cloud Provider Brokerage Example	14
Figure 4: Scope of Controls between Provider and Consumer	15
Figure 5: Cloud Provider – Service Orchestration	16
Figure 6: Challenging Security Requirements to Mitigation Mapping	50
Figure 7: Information Life Cycle Management Phases	69
Figure 8: Service-Level Agreement Generic Concepts Mindmap	79
Figure 9: Cloud-Specific SLA Concepts Mindmap	80
Figure 10: High Level View of Metric Concept Model.....	81

List of Tables

TABLE 1: RELATIONSHIP BETWEEN VOLUME I REQUIREMENTS AND WORK PRESENTED IN VOLUME II	
TABLE 2: MISSION REQUIREMENTS FROM TARGET BUSINESS USE CASES.....	26
TABLE 3: BUSINESS USE CASES AND MISSION REQUIREMENTS	27
TABLE 4: CROSS-CUTTING SECURITY SYSTEM REQUIREMENTS	32
TABLE 5: CROSS-CUTTING INTEROPERABILITY SYSTEM REQUIREMENTS	35
TABLE 6: CROSS-CUTTING PORTABILITY REQUIREMENTS	39
TABLE 7: MAPPING SYSTEM REQUIREMENTS TO MISSION REQUIREMENTS.....	40
TABLE 8: SAJACC USE CASES	42
TABLE 9: AREA OF STANDARDIZATION GAPS AND STANDARDIZATION PRIORITIES	46

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 Table 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.

US Government Cloud Computing Technology Roadmap, Volume II, Release 2.0

	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 Case Working Group	Standards Acceleration to Jumpstart the Adoption of Cloud Computing	NIST Special Publications: 800-125, 800-144, 800-146	NIST Complex Information System Measurement Project: Koala, IaaS Computing Simulation Model
Requirement 1: International voluntary consensus-based standards (<i>interoperability, performance, portability, and security standards</i>)	X				X	X	
Requirement 2: Solutions for High-priority Security Requirements, technically de-coupled from organizational policy decisions (<i>security standards and technology</i>)	X		X	X	X	X	
Requirement 3: Technical specifications to enable development of consistent, high-quality Service-Level Agreements (<i>interoperability, performance, portability, and security standards and guidance</i>)		X	X			X	
Requirement 4: Clearly and consistently categorized cloud services (<i>interoperability and portability guidance and technology</i>)		X					
Requirement 5: Frameworks to support seamless implementation of federated community cloud environments (<i>interoperability and portability guidance and technology</i>)		X		X		X	
Requirement 6: Updated Organization Policy that reflects the Cloud Computing Business and Technology model (<i>security guidance</i>)			X				
Requirement 7: Defined unique government regulatory requirements and solutions (<i>accessibility, interoperability, performance, portability, and security technology</i>)			X	X		X	
Requirement 8: Collaborative parallel strategic “future cloud” development initiatives (<i>interoperability, portability, and security technology</i>)				X			
Requirement 9: Defined and implemented reliability design goals (<i>interoperability, performance, portability, and security technology</i>)			X	X		X	X
Requirement 10: Defined and implemented cloud service metrics (<i>interoperability, performance, and portability standards</i>)		X	X	X			X

Table 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 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.

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.

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.

US Government Cloud Computing Technology Roadmap, Volume II, Release 2.0

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

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 provides 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 Define and implemented cloud service metrics – Volume I, Requirement 10.”

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.

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 Development 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 and highlights the changes brought upon during the final editing process for this document.

2.2.1 Conceptual Model

Figure 1 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.

The reference architecture displayed in Figure 1 is an updated version based on additional public comments received in the revision process for this document. Through the RA/TAX Public Working Group process, this new model has been verified and approved by its members. The principal difference between the original reference architecture (found in NIST 500-292) and the one in this document is the change in the position of the "Security" and "Privacy" components. Security and Privacy were originally identified as cross-cutting concerns and items that are shared responsibilities for each cloud computing actor, therefore the placement of Security and Privacy as a backplane to the cloud computing reference architecture is an appropriate change to the model.

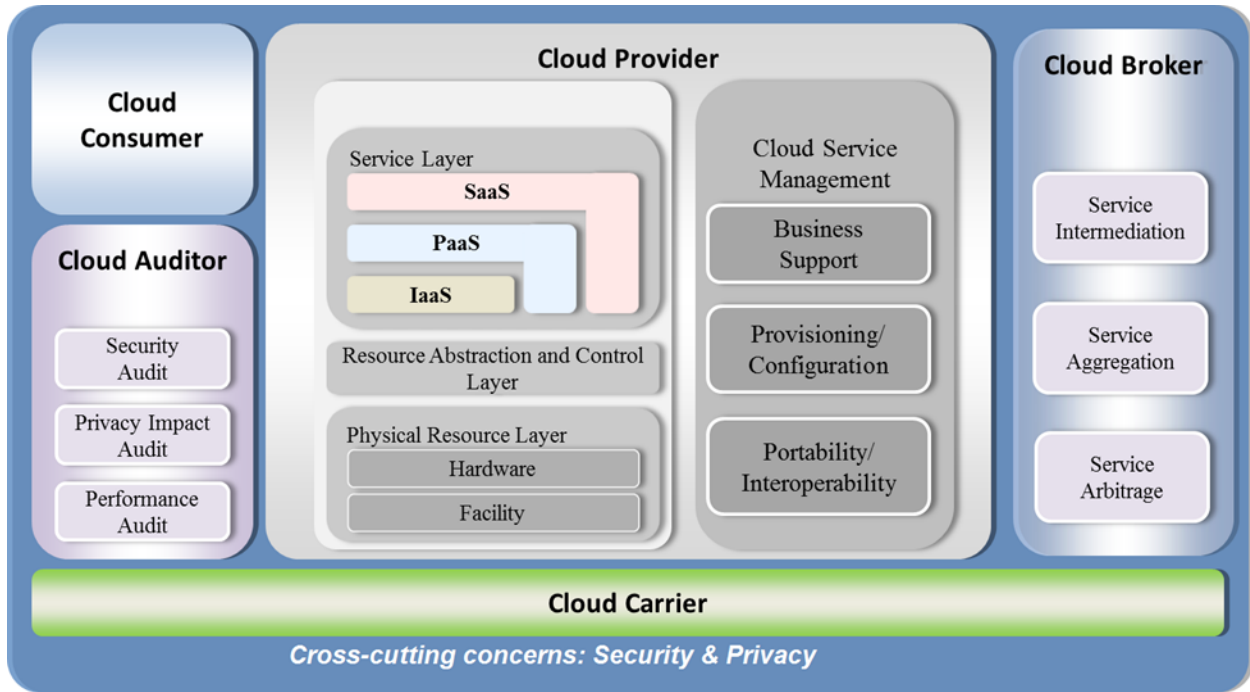


Figure 1: The Conceptual Reference Model

2.2.2 Cloud Computing Actors

As shown in Figure 1, 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 for the cloud computing service. 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 need SLAs to specify the technical performance requirements fulfilled by a cloud-provider. SLAs can cover terms regarding the quality of service, security, and remedies for performance failures.

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, and performance. There are many reasons an organization (government or not) may have aspects of privacy evaluated by an auditor.

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 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.

A Cloud Broker may provide services in two separate domains:

- Business and relationship support services (business intermediation).
- Technical support service (aggregation, arbitrage and technical intermediation), with a key focus on handling interoperability issues among multiple providers.

Cloud Brokers may behave as Business brokers in some cases, Technical brokers in others or may take on both roles.

A Business Cloud Broker is an entity that offers Cloud Consumers business and relationship services to evaluate and select Cloud Providers based upon the consumer's requirements. Business brokerage does not offer technical broker-related capabilities to interact with Cloud Consumer data in Cloud Provider environments. Business brokerage can be combined with or operate independently of technical Brokerage services. They do not have any contact with the consumer's data migrated to the cloud, consumer operational processes in the cloud or consumer-based cloud artifacts such as images, volumes or firewalls.

A Technical Cloud Broker is an entity that offers Cloud Consumers the capability to consistently interact with the consumer's operational processes, cloud artifacts and/or data residing in Cloud Providers environments by aggregating services from multiple providers and adding a layer of technical functionality that addresses consistent interface and interoperability issues. Technical Brokerage does not offer business broker-related capabilities to evaluate and select Cloud Providers. Technical brokerage can be combined with or operate independently of business brokerage services. This individual will interact with the consumer's operational processes, cloud artifacts and/or consumer data by aggregating services from multiple providers and adding a layer of technical functionality by addressing single-point-of-entry and interoperability issues.

A Technical Cloud Broker has two defining qualities which distinguish it from a Cloud Provider:

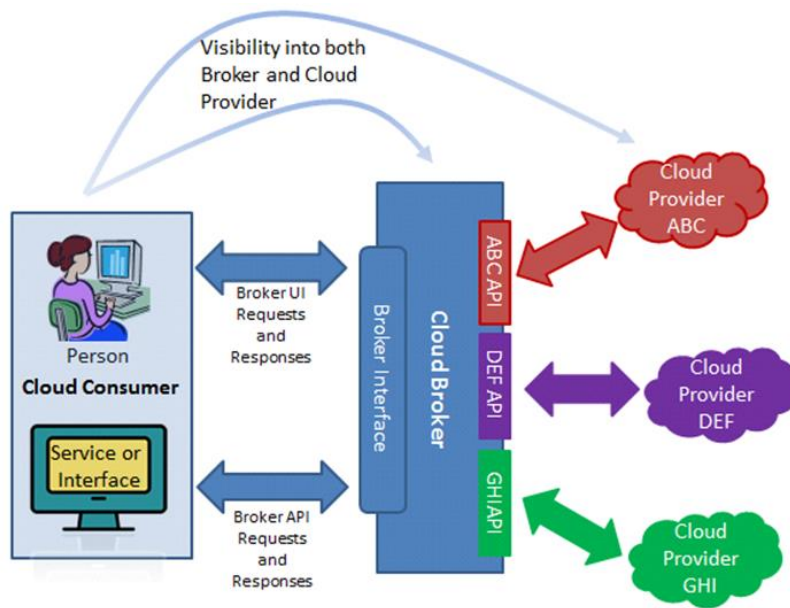
1. The Cloud Broker provides a single point of entry for managing multiple cloud services. The key defining features that distinguish a Cloud Broker from a Cloud Service Provider are the ability to provide a single consistent interface to multiple differing providers, whether the interface is for business or technical purposes and to provide the Cloud Consumer with complete transparency into the identity of the supporting Cloud Service Providers.
2. The Cloud Broker provides transparency to the Cloud Consumer on the identity of the Cloud Providers in the background. A Cloud Broker will always allow the cloud Consumer a particular level of transparency into the identity of the target Cloud Providers. An entity that provides additional layers of functionality to only one Cloud Provider or does not allow transparency into their underlying Cloud Providers will be considered an Intermediary Cloud Provider.

In both cases, the Cloud Brokers provide a single interface across multiple cloud service provider targets, for business or technical services respectively. Combinations of technical and business brokerage can be carried out by the same entity. Cloud Providers who operate in a Cloud Broker role can provide brokerage services as well.

There are two classes of Cloud Providers - primary and intermediary. A Primary Cloud Provider is an entity that provides the cloud consumer with a full stack of cloud services without relying on other entities to provision particular functional service layers, whereas an Intermediary Cloud Provider is a cloud provider that relies on one or more other cloud service providers to deliver services to the cloud consumer, but does not provide the consumer with visibility into the underlying cloud provider services in use. They may also provide additional layers of functionality to only one Cloud Provider.

Basic Broker Model

Figure 2 shows a basic example of cloud brokerage. Depending upon the broker services rendered, the brokerage can be business oriented, technically oriented or a combination of the two. Cross-provider business services might include service catalogue lookups, subscription handling, customer relation management, etc. and cross-provider technical services might include orchestration, load management and cloud-bursting, integrated identity and authorization management, security Brokerage and integrated security management, metrics retrieval, unified billing, cost and usage reporting, etc.



Note that two key characteristics of brokerage are fulfilled:

- The Cloud Consumer uses a single broker interface to engage with multiple providers.
- The Cloud Consumer retains visibility into the cloud service providers they use through the broker, either through the broker, directly or both.

Figure 2: Cloud Broker Interactions

Intermediary Cloud Provider Example

Figure 3 below shows a simple example of an intermediary cloud provider interaction with a cloud consumer. Both the broker and the intermediary provider have the capability to interact with multiple cloud providers. However, the intermediary cloud provider presents only its service interface to the cloud consumer and does not offer visibility into or control over any additional cloud or non-cloud providers used in the creation of the service behind the scenes. Instead, the intermediary cloud uses the additional providers as invisible components of its own service, which is presented to the customer as an integrated offering.

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.

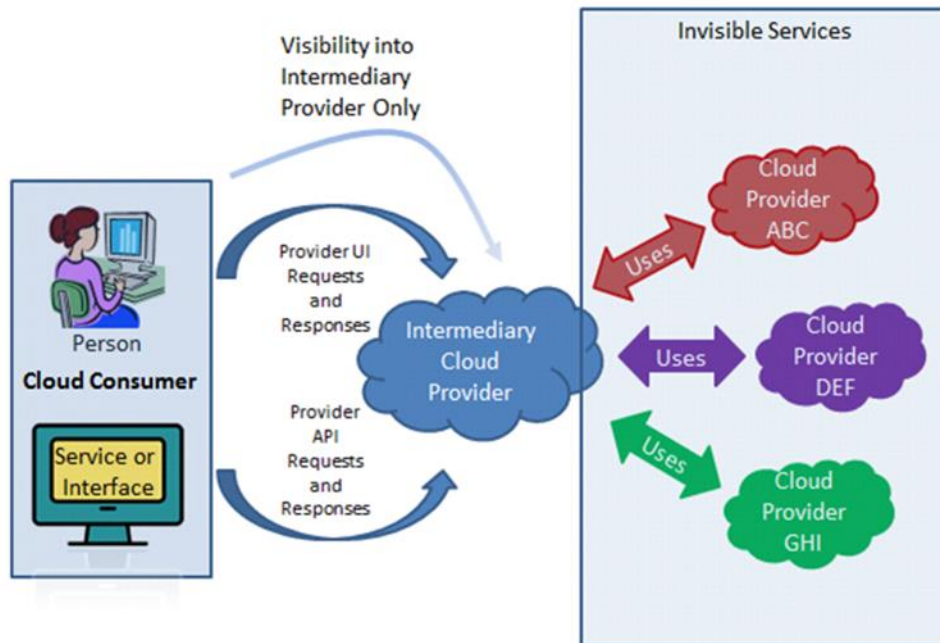


Figure 3: Intermediary Cloud Provider Brokerage Example

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. The figure 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.

- 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.




Figure 4: Scope of Controls between Provider and Consumer

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 5 shows a generic stack diagram of this composition that underlies the provisioning of cloud services.

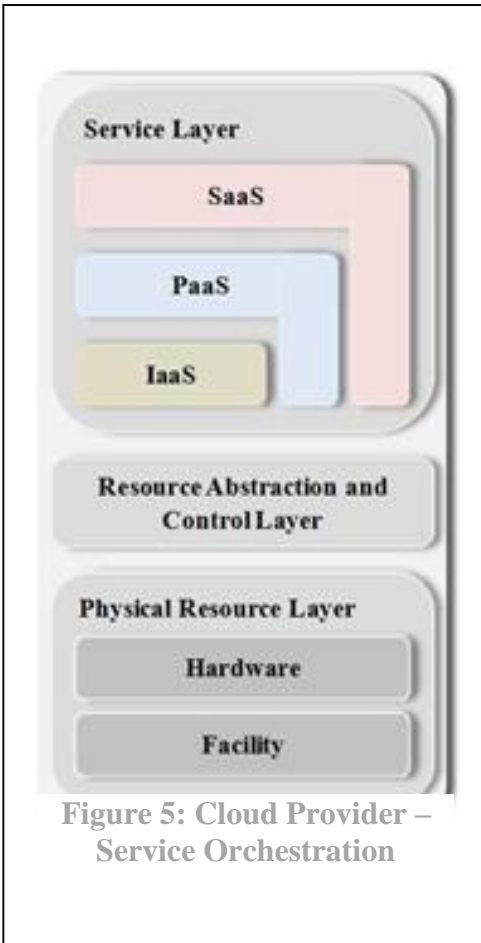


Figure 5: Cloud Provider – Service Orchestration

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 5, 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.
- 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

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 1 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;
- “Defined and implemented cloud service metrics” – Requirement 10.

use cases are developed, 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.

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

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 Census Virtual Desktop Infrastructure

Delivery Model: Private Cloud

Service Model: SaaS

Agency: United States Bureau of the Census

FISMA Impact Level: Moderate

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

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 either data storage or processing needs that appear to make them more cost-effective in a community cloud setting. As a subsidiary use case, 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 E-Discovery

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 to 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

The 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 usually 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 and adoption is expected to increase.

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. Those 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. Further 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
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.
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.
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.
Workforce Mobility	Provide mobile access to all IT services, enabling secure access from any device and any place where there is sufficient network bandwidth.
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.
Common Geospatial Platform	Provide agencies with the ability to create and deploy geospatial applications rapidly and efficiently.
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.

Requirement	Description
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.
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 2: 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 3: Business Use Cases and Mission 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:

US Government Cloud Computing Technology Roadmap, Volume II, Release 2.0

- Identify and provide solutions for high-priority security requirements (Requirement 2, 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 and Credential Access 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.

Requirement/Details	Description
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 Format and Exchange</i>	Both high-level and detailed security audit data must be provided in a standards-based format so that cloud consumers could analyze the 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 consumer 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 consumer at that cloud provider.

Requirement/Details	Description
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.

Requirement/Details	Description
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 obtain certification that the service being provided meets the requirements for the stated FISMA data classification. The Office of Management and Budget (OMB) issued the FedRAMP Policy memo, establishing FedRAMP and its role in the Assessment and Authorization (formerly known as the Certification and Authorization) process for cloud systems. The FedRAMP Program Management Office (PMO) issued a security control baseline at the Low and Moderate impact levels and a Concept of Operations (CONOPs). The FedRAMP A&A process is based on the NIST Risk Management Framework and leverages the NIST SP 800-53 (Revision 3) security controls for information and information systems and provides a methodology for assessing and authorizing cloud offerings to Federal Cloud consumers, providing a cost-effective, risk-based approach for assessing and authorizing cloud offerings to Federal cloud consumers.

Table 4: Cross-Cutting Security System 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.

Requirement/Details	Description
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 Tickets</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 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.

Requirement/Details	Description
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 5: Cross-Cutting Interoperability System 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.
Data Export	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.
Data Import	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.

Requirement/Details	Description
2 <u>Data Deletion</u>	<p>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).</p>
Deletion of Business Data	<p>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.</p>
Deletion of Logs, Usage Data, and Audit Data	<p>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.</p>
Code Escrow	<p>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.</p>

Requirement/Details	Description
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.
Portability of Responsive Electronically Stored Information	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 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.
Portability of Metadata Required for eDiscovery and FOIA	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.
Export of Electronically Stored Information for eDiscovery and FOIA	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.

Requirement/Details	Description
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.
Moving Virtual Desktops Between Vendors	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.
Migration of Virtual Desktops	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.
Accessibility of Virtual Desktops from Heterogeneous Devices	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.
Virtualization of Legacy Software	The cloud provider shall provide a means for virtualizing legacy 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 6: 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 7: Mapping System Requirements to Mission Requirements

3.2 SAJACC Use Cases and Technical Requirement

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 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.

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.**

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 on-going; however, the first set of published SAJACC use cases includes three categories: management, interoperability, and security, as shown in Table 8 below.

Management	Interoperability	Security
<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • 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 8: 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 requirement expressed in these use cases. Furthermore, the SAJACC project has developed a generic testing framework and implemented test drivers for an initial set of use cases using identified known system interfaces. This testing mechanism has demonstrated how cloud consumers’ technical requirements can now be implemented using existing public interfaces and also helped to surface issues and gaps in known system interfaces. The set of use cases, test drivers, and testing results will provide concrete data to SDOs in developing high-quality standards in addressing portability, interoperability, and security concerns expressed by the consumers. The SAJACC project will continue the efforts to maintain and develop technical use cases to update existing use cases with community input and address future looking technical requirements. The project furthermore will develop demonstrable test drivers to show how existing system interfaces can be used to implement these requirements and what issues and gaps exist to feed into SDOs 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

Highlights:

To support US government requirements for accessibility, interoperability, performance, portability, and security in cloud computing, the NIST public Cloud Computing Standards Roadmap Working Group has surveyed the existing standards landscape for accessibility, performance, 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
Version 2, *NIST Cloud Computing Standards Roadmap*

leveraging this existing, publicly available work, plus the work of the other NIST working groups, to identify standards, standards gaps, and standardization priorities.

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 accessibility, performance, 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>).

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.

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 9 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	High Priorities for Standardization Based On USG Requirements
<p>SaaS Functional Interfaces (Section 9.1.1 of SP 500-291 V2), e.g.,</p> <ul style="list-style-type: none"> - Data format and interface standards for email and office productivity - Metadata format and interface standards for e-discovery 	<p>High standardization priorities on:</p> <ul style="list-style-type: none"> - SaaS application specific data and metadata format standards to support interoperability and portability requirement when migrating high-value, low-risk applications to SaaS (Section 9.2.3 of SP 500-291 V2).
<p>SaaS Self-Service Management Interfaces (Section 9.1.2 of SP 500-291 V2), e.g.,</p> <ul style="list-style-type: none"> - Interface standards related to user account and credential management 	<p>Not a high standardization priority at this time</p>
<p>PaaS Functional Interfaces (Section 9.1.3 of SP 500-291 V2), e.g.,</p> <ul style="list-style-type: none"> - Standards of data format to support database serialization and de-serialization 	<p>Not a high standardization priority at this time</p>
<p>Business Support, Provisioning and Configuration (Section 9.1.4 of SP 500-291 V2), e.g.,</p> <ul style="list-style-type: none"> - 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 	<p>High standardization priorities on:</p> <ul style="list-style-type: none"> - Resource description and discovery standards to support data center consolidation using private and community IaaS cloud systems (Section 9.2.4 of SP 500-291 V2)

Area of Standardization Gaps	High Priorities for Standardization Based On USG Requirements
<p>Security (Section 9.1.5 of SP 500-291 V2), e.g.,</p> <ul style="list-style-type: none"> - 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 	<p>High standardization priorities on:</p> <ul style="list-style-type: none"> - Security auditing and compliance standards to support secure deployment, assess, and accreditation process for cloud-specific deployment (Section 9.2.1 of SP 500-291 V2) - Identity and access management standards to support secure integration of cloud systems into existing enterprise security infrastructure (Section 9.2.2 of SP 500-291 V2)
<p>Accessibility (Section 9.1.6 of SP 500-291 V2), e.g.</p> <ul style="list-style-type: none"> - Standardized “framework” for exchanging an individual’s accessibility requirements 	<p>Not a high standardization priority at this time</p>

Table 9: Area of Standardization Gaps and Standardization Priorities

4.3 Accelerating the Development the Use of Cloud Computing Standards

Existing and new standards need to be applied or developed in support of agency's requirements, such as for interoperability, portability, security, performance, and accessibility, for cloud computing services. There is already 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 coordinate and contribute clear and comprehensive user requirements for cloud computing standards projects.

Recommendation 2 – Participate in Standards Development

Agencies should actively participate and coordinate in cloud computing standards development projects that are of high priority to their agency missions. The January 17, 2012, White House Memorandum, [M-12-08](#), lists five fundamental strategic objectives for federal government agencies whenever engaging in standards development:

- Produce timely, effective standards and efficient conformity assessment schemes that are essential to addressing an identified need;
- Achieve cost-efficient, timely, and effective solutions to legitimate regulatory, procurement, and policy objectives;
- Promote standards and standardization systems that promote and sustain innovation and foster competition;
- Enhance U.S. growth and competitiveness and ensure non-discrimination, consistent with international obligations; and
- Facilitate international trade and avoid the creation of unnecessary obstacles to trade.

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. Agencies should also include consideration of conformity assessment approaches currently in place that take account of elements from international systems, to minimize duplicative testing and encourage private sector support.

Recommendation 4 – Specify Cloud Computing Standards

Agencies should specify cloud computing standards in their procurements and grant guidance when multiple vendors offer standards-based implementations and there is evidence of successful interoperability testing.

Recommendation 5 – USG-Wide Use of Cloud Computing Standards

To support USG requirements for accessibility, interoperability, performance, portability, and security in cloud computing, the Federal Cloud Computing Standards and Technology Working Group, in coordination with the Federal CIO Council Cloud Computing Executive Steering Committee (CCESC) and the Cloud First Task Force, should recommend specific cloud computing standards and best practices for USG-wide use.

5 High-Priority Security Requirements

Highlights: 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;
- “Updated Organization Policy that reflects the Cloud Computing Business and Technology model”- 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*.

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:

- **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 may introduce new security architectures and solutions, resulting in services that are 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 Responsibility

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.

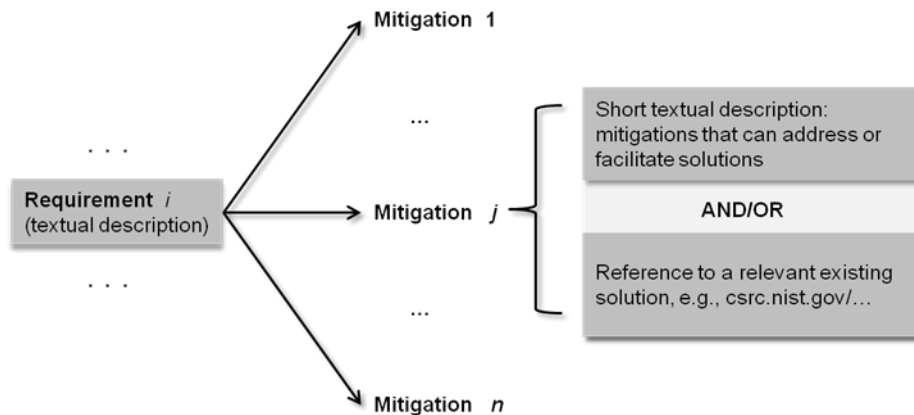


Figure 6: Challenging Security Requirements to Mitigation Mapping

Figure 6: Challenging Security Requirements to Mitigation Mapping 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 NIST SP 800-53 Security Controls for Cloud-Based Information Systems

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 actor, the Cloud 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 consumer 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 consumer of responsibility to monitor and audit aspects of the information system that the consumer operates or manage.

Importance: Cloud Auditing and continuous monitoring are requirements 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 on how 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://cloudaudit.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

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 Needed Electronic Discovery Guidelines

Description: Mechanism to provide access to data in response to lawful authority while protecting consumer 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, consumers must gain an understanding of how the cloud provider processes electronic discovery and litigation holds. The consumer 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: Consumers 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 consumer. 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 e-discovery. 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 Needed Cloud Privacy Guidelines

Description: This requirement addresses the need to build predictability and confidence in the degree to which cloud solutions provide privacy of 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: 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 Consumer and the Cloud Provider, given that the Cloud Consumer is still responsible for the protection of the data.

Mitigation 2: Establish and maintain the confidence of those individuals for and about whom federal agencies manage personal data. Cloud Consumers (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 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 Clarity on Cloud Actors Security Roles and Responsibilities

Description: Mechanism to define who (among cloud actors such as consumer 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 consumer) 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 consumer 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-consumer 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 consumers 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 consumer. 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 consumer contingency plan for individual application and/or system into the cloud provider's overall contingency plan;
- Establish management succession and escalation procedures between cloud provider and consumer; 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 consumer'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)

5.3.9 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 consumer'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 consumers that the cloud infrastructure is secure. Also important is the ability to provide consumer 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

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

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

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

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://csrc.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 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 Visibility for Consumers

Description: Mechanism to define how cloud subscribers (consumers) 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 providers 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 loss of business that results.

Mitigation 1: Agreement and cooperation between providers and consumers to implement customized controls based on consumer-specific requirements and to provide transparency to their implementation and use.

As pointed out in the FedRAMP's Considerations for Federal Cloud Computing Audit and Risk Assessment, SLAs should identify customer-specific requirements and clearly state who is responsible for what monitoring and audit task (to prevent visibility gap between provider and customer) and who is informed of the results. Additionally, standards and methods should be specified for customers to instruct the cloud as to what to monitor and to be alerted about.

Mitigation 2: Effective Monitoring

Consumers can achieve greater visibility with an effective monitoring system that includes:

- Packet-based, strategically deployed among physical and virtual machines, real-time monitoring and historical trending metrics;
- End-to-end monitoring and measurement handled by cloud applications and embedded into the cloud architecture using cloud application programming interfaces (APIs);

- Centralized monitoring and analysis system of configuration files and log files, with automatic alert capability; and
- SCAP-compliant monitoring tools. SCAP is an alert format standard mandated by the U.S. Government and which can help providers push alerts to consumers in a standard format.

Mitigation 3: Audit

CloudAudit.org is a Cloud Security Alliance (CSA) standardization initiative to provide a common interface and namespace (mostly through mapping) that allows providers to automate the Audit, Assertion, Assessment, and Assurance (A6) of their cloud environments and allows customers to do likewise via open, extensible, and secure APIs.

Mitigation 4: Unified monitoring and management tools

- Unified and centralized tools that monitor and manage both physical and virtual environments and can be accessed by both administrators and customers (e.g., EMC Ionix and VMware vCenter); and
- Tools that push monitoring to customers and allow customers to configure what is interesting to them (e.g., Amazon CloudWatch).

References:

- Cloud Security Alliance, www.cloudaudit.org
- Security Content Automation Protocol (SCAP), <http://scap.nist.gov>
- <http://aws.amazon.com/cloudwatch/>
- <http://symmetrix.com/products/detail/software/ionix-unified-infrastructure-manager.htm>
- NIST SP 800-144, Guidelines on Security and Privacy in Public Cloud Computing
- NIST SP 800-146, Draft Cloud Computing Synopsis and Recommendations
- FedRAMP's Considerations for Federal Cloud Computing Audit and Risk Assessment

5.4.2 Control for Consumers

Description: The assessment is that consumers have limited control over security policies enforced by cloud providers on their behalf. There is little automation available to help consumers to implement technical controls (policies) in their applications which are deployed in cloud models. A mechanism is needed to allow cloud consumers 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 consumers 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. Consumers 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

Consumers are responsible for the selection and use of appropriate cloud models. Through the selection process, consumers can ensure that they gain adequate visibility. When selecting the appropriate cloud model, consumers should research and understand:

- Public, hybrid, community, and private cloud models with increasingly greater customer control over tenants;
- SaaS, PaaS and IaaS service models with increasingly greater customer control over infrastructure;
- Externally hosted and internally hosted models with increasingly greater customer control over location; and
- External provider operated, outsourced, and internally operated with increasingly greater customer control over personnel.

Mitigation 2: Control Data Objects

Access control over data objects is a widely used and mature function. Consumers 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, Applications and Networks

Consumers can ascertain the correct implementation of the security controls better when they have of control of the VMs and existing applications. This process ensures consumers can:

- 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.
- Apply similar protection mechanisms of physical network (for example, firewall, IDS and antivirus) to intra-host virtual network (vSwitches/vLANs); and

- Make traffic in virtual network visible to security and monitoring devices on physical network.

Mitigation 5: Control of Geographic Location

Federal cloud consumers may, through SLAs and contract requirements, restrict the geographic location of data due to the potential variances in privacy and security regulations of some jurisdictions. Such restrictions could impose additional burden on providers and potentially impact cost and efficiency.

References:

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

5.4.3 Data Security

Description: The loss of confidentiality, integrity, or availability of consumer data results in a variety of impacts. Cloud consumers 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 consumer 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 consumer wishes to upload into a cloud must be protected in transit; similarly, data that a consumer 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 consumer and needs to be protected against access attempts by all other entities, while preserving the availability for the originating consumer;
- Passive stored with selective shared access: Data should be accessed only by entities that have been authorized by the originating consumer 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 consumers;
- 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 consumers;
- 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 consumer data.

Importance: High. If cloud services do not offer robust protection of consumer 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:

- 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 instead of Public ISP access;
- Use a secure (encrypted) communication when accessed from a mobile wireless devices;
- Use of intranet, cross-agency or cross-department
- Protect data using encryption for confidentiality and hashing or signatures for integrity.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 7. The security requirements in this life cycle are defined based on the types of data.

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

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



Figure 7: 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

5.4.4 Risk of Account Compromise

Description: A benefit of cloud computing is easy accessibility. A consumer 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, consumers 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:

- There is additional attack surface exposure due to increased complexity and dynamic infrastructure allocation;
- New APIs/interfaces are emerging that are untested; and
- The consumer'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

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: Use the National Strategy for Trusted Identities in Cyberspace (NSTIC) mechanisms to efficiently manage the identities while users' privacy is protected.

NSTIC provides the means of creating a secure, trusted Identity Ecosystem that is capable of establishing a user-centric privacy protection for any Cloud Ecosystem. The mechanisms employed by an Identity Ecosystem are structured in a robust framework composed of the overarching set of interoperability standards, risk models, privacy and liability policies, requirements, and accountability mechanisms.

Mitigation 4: 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:

- National Strategy for Trusted Identities in Cyberspace (NSTIC)
- Symantec Internet Security Threat Report, Trends for 2010, Volume 16, April 2011

5.4.5 Identity Credential and Access Management (ICAM) and Authorization

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

Importance: High. The identity credential and access management (ICAM) 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 consumers 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 ICAM for cloud-based services and use of agency ICAM 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 ICAM 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.

Mitigation 6: Use the National Strategy for Trusted Identities in Cyberspace (NSTIC) mechanisms to efficiently manage the identities while users' privacy is protected.

NSTIC provides the means of creating a secure, trusted Identity Ecosystem that is capable of establishing a user-centric privacy protection for any Cloud Ecosystem. The mechanisms employed by an Identity Ecosystem are structured in a robust framework composed of the overarching set of interoperability standards, risk models, privacy and liability policies, requirements, and accountability mechanisms.

References:

- XACML <http://www.oasis-open.org/committees/download.php/2406/oasis-xacml-1.0.pdf>
- 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
- National Strategy for Trusted Identities in Cyberspace (NSTIC)

5.4.6 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:

- 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: The risks associated with multi-tenancy could also be mitigated through physical separation which can be provisioned to consumers with special security requirements and which implies the use of special virtual environments with physical separation of the full-stack cloud infrastructure.

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.

- Edward L. Haletky, “VMware vSphere and Virtual Infrastructure Security”, Prentice Hall, 2009, ISBN-13: 978-0-137-15800-3.

5.4.7 Cloud-based Denial of Service

Description: Because cloud consumers depend on functional networks to access their resources, and because networks are often not under consumer 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: DoS attacks are not new, but cloud computing has increased the attack surface. Internally accessed applications become remotely accessible when provided as cloud services, and 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 using the vast resources of cloud at a new level of severity.

Mitigation 1: The cloud consumer may adopt a hybrid approach, potentially through a cloud broker, 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 consumer 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 consumer 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 with appropriate SLAs.

References:

- Cloud Security Alliance, The Cloud Control Matrix
- Federal Risk and Authorization Management Program (FedRAMP)
- NIST SP 500-291, Cloud Computing Standards Roadmap
- NIST SP 500-292, Cloud Computing Reference Architecture
- NIST SP 500-293, US Government Cloud Computing Technology Roadmap Volume 1,
- High-Priority Requirements to Further USG Agency Cloud Computing Adoption
- NIST SP 500-293, US Government Cloud Computing Technology Roadmap Volume III,
- Technical Considerations for USG Cloud Computing Deployment Decisions (Draft)

- NIST SP 800-53, Recommended Security Controls for Federal Information Systems and Organizations, Revision 3

5.4.8 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 consumers.

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.⁷

⁷ NIST SP 800-61, Computer Security Incident Handling Guide

6 Summary and Next Steps

The *USG Cloud Computing Technology Roadmap*, including Volume II, Useful Information for Cloud Adopters, and the work that was used as the basis to draft it, was completed in less than 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.

It is intended to be a first step toward a two-fold objective:

- Strategic – to support the identification and communication of high-priority USG Cloud Computing Requirements by providing an explanation of the objective rationale for the assertion that these are not currently met to the extent needed; and
- Tactical – to support adopters in the interim period while the cloud model and implementation is maturing by providing information to help make informed decisions, in this case, through a consolidated guide to existing NIST collaborative and projects work and conclusions.

To achieve the goals for USG Cloud Computing Adoption, it is necessary to work on both levels.

To make progress, there is a need to explicitly agree on what the strategic priorities are. This seems very basic in the context that at various cloud symposiums, the same subjects are discussed, but there wasn't an existing confirmed list. There are in some cases alternative lists, and many sources for elements to be included in a consolidated list – numerous publications by academic, standards, and industry organizations, and government agencies. However, these are “centric” to and developed from the perspective of the organization that drafted them. Moreover, 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 priorities and Priority Action Plans, identified as candidates for self-tasking by the cloud community presented in Volume I. These are presented for comment, and the expectation is that the Volume I high-priority requirements will be refined and satisfied over a multiyear time frame, consistent with technology development cycles.

Volume II summarizes the work that not only supports these priorities, but provides some level of information, help, and tools for the short term. Each major area of work has a very specific tactical collaborative process which is under way, which relate to the PAPs listed in Volume I, and which can immediately go forward with cloud computing community participation:

- 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;

US Government Cloud Computing Technology Roadmap, Volume II, Release 2.0

- 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.

7 Appendix A – Service Level Agreement (SLA) Taxonomy and Metrics

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.

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:

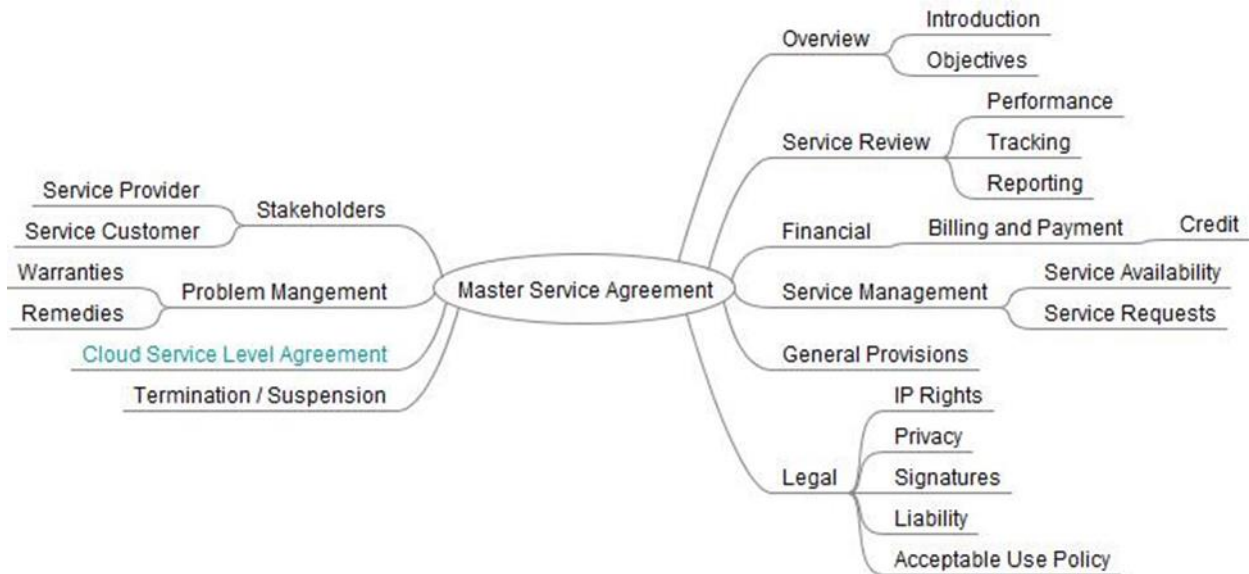


Figure 8: Service-Level Agreement Generic Concepts Mindmap

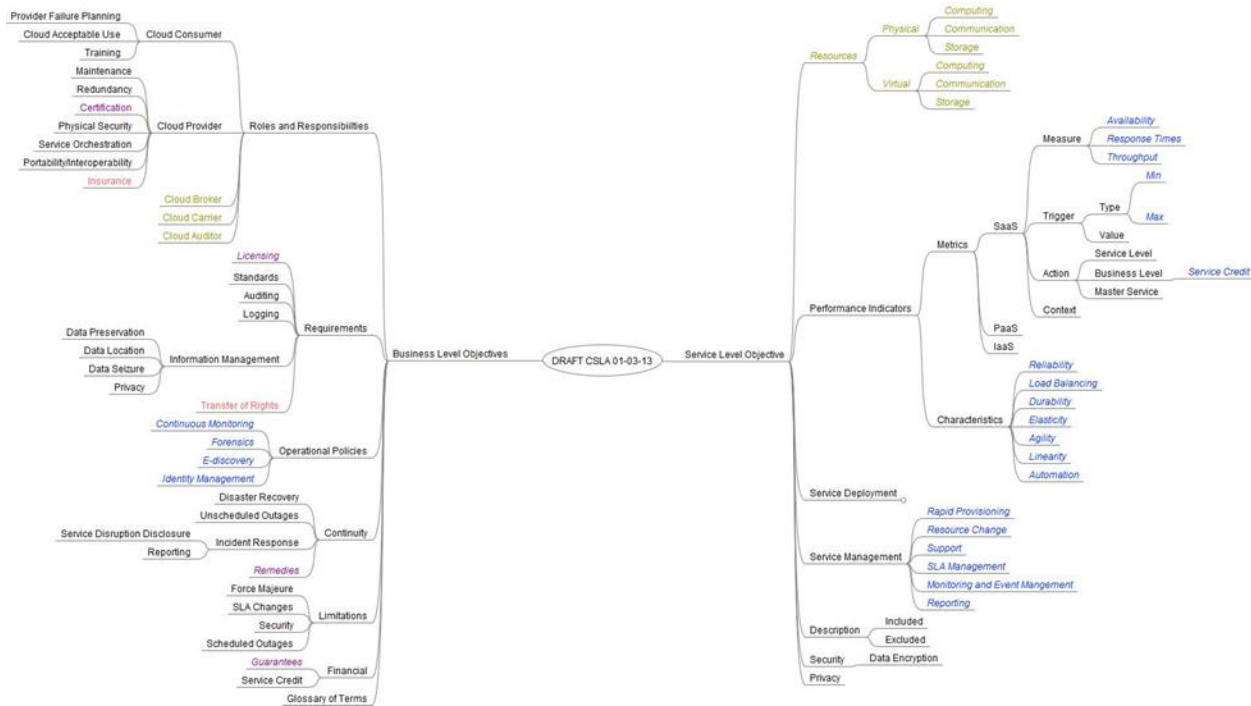


Figure 9: Cloud-Specific SLA Concepts Mindmap

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.

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

In January 2012, the Cloud Metrics subgroup was formed and a survey of the domain of metrics related to cloud service was conducted.

The study showed the disparities and ambiguities in cloud metrics, which confirms the necessity for industry and USG agencies to develop “**Defined and implemented cloud service metrics**” - NIST USG Cloud Computing Technology Roadmap Vol.1, Requirement 10.

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

CLOUD SERVICE METRICS

At the completion of version 1.0 of the Reference Architecture (RA) the Reference Architecture working group was asked to identify additional areas of cloud computing that would affect interoperability, portability and security. The group identified the specification of the metrics in the context of cloud services as critical to the development and use of efficient inter-connected cloud computing services.

Furthermore NIST identified in its definition of cloud computing a “Measured Service” as being one of the five essential characteristics of the cloud computing model.

The Reference Architecture working formed the Cloud Metrics subgroup with the mission to tackle this problem.

A survey of publically available documents referencing service oriented metrics; metrics usage etc. showed that as of today the measurement space is not necessarily well defined. Common terminologies (i.e. measures, metrics) or sets of measurement artifacts (i.e. units of measurements, metrics) often have several definitions, which makes it very difficult for the consumer to compare services or rely on third party tools to monitor the health of the service. It can also make it difficult for the provider to show that the service is performing correctly or to allow its service to enter into complex cloud service chain or federation.

The group also surveyed work done by organizations on metrics with the intent to gather as much knowledge as possible on existing metrics to either reuse them, adapt them to cloud services or identify gaps. Some organizations have already published documents defining valuable metrics. Other organizations have work on describing possible frameworks or measurement methods or have started efforts on the domain of cloud measurement.

This effort led to the creation of a metric concept model whose purpose is to identify and characterize the information and relationships needed to efficiently and consistently define and use measures, metrics and plans in the context of cloud services.

The Figure 10 below represents the high level view of the concept model. The content of the grey boxes define the three primary entities. The content of the respective colored boxes show some characteristics of these entities. For instance the Metric entity relies on a Measure entity and can have minimum and maximum limits that if reached will trigger an action.

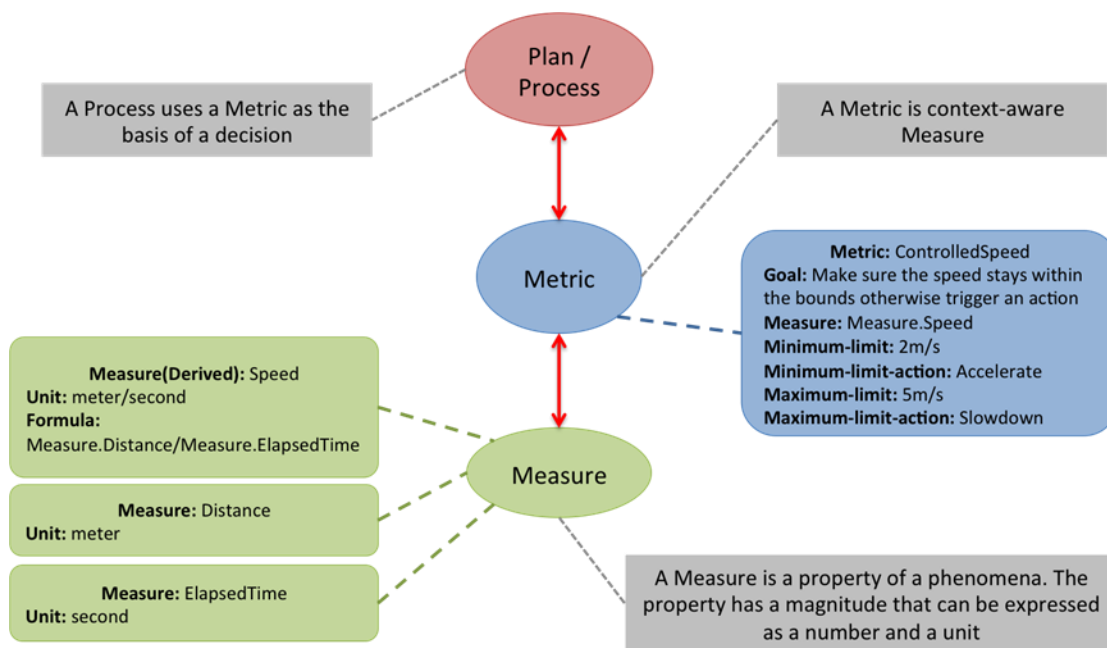


Figure 10: High Level View of Metric Concept Model

The ability to organize the information relevant to cloud service metrics and measures and their connections to plans like SLAs is very important as it allows stakeholders to better collect, assess and compare different aspects of the cloud services they intend to manipulate.

8 Appendix B – 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.⁸⁹ 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.

⁸ *Koala*: A Discrete-Event Simulation Model of Infrastructure Clouds, K. Mills, J. Filliben and C. Dabrowski

⁹ C. Dabrowski and K. Mills, VM Leakage and Orphan Control in Open-Source Clouds

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.

9 Appendix C – Useful References

The following sources may be useful for further reference.

NIST Special Publications and Drafts

- [NIST Special Publication 800-53](#), 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](#), Guidelines on Security and Privacy Issues in Public Cloud Computing.
- [NIST Special Publication 800-145](#), A NIST Definition of Cloud Computing.
- [NIST Special Publication 800-146](#), 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*](#)

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*](#).

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

OpenCrowd, [*Cloud Taxonomy*](#).

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