

# **NIST Cloud Computing Challenging Security Requirements for USG Adoption of Cloud Computing**

---

**National Institute of Standards and  
Technology**

---

Lee Badger  
Fred Whiteside  
Jian Mao  
Shilong Chu

**NIST Special Publication 500-XXX**

NIST Cloud Computing

Challenging Security Requirements for  
USG Adoption of Cloud Computing

*National Institute of Standards and  
Technology*

Lee Badger  
Fred Whiteside  
Jian Mao  
Shilong Chu

---

# Internetwork Technologies

---

Computer Security Division  
Information Technology Laboratory  
National Institute of Standards and Technology  
Gaithersburg, MD 20899-8930

September 2011



U. S. Department of Commerce

*Gary Locke, Secretary*

National Institute of Standards and Technology

*Patrick Gallagher, 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. ITL's responsibilities include the development of technical, physical, administrative, and management standards and guidelines for the cost-effective security and privacy of sensitive unclassified information in Federal computer systems. 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-XXX**

**Natl. Inst. Stand. Technol. Spec. Publ. 500-XXX, 56 pages (September 2011)**

### **DISCLAIMER**

This document has been prepared by the National Institute of Standards and Technology (NIST) and describes standards research in support of the NIST Cloud Computing Program.

Certain commercial entities, equipment, or material may be identified in this document in order to describe a 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 these entities, materials, or equipment are necessarily the best available for the purpose.

## **Acknowledgements**

This document emerged from a protracted period of discussions with a number of stakeholders including the NIST Cloud Computing Security Working Group, chaired by Lee Badger and Fred Whiteside of the US Department of Commerce, National Institute of Standards and Technology. Discussion with members of the NIST Cloud Computing Working Groups: Reference Architecture Working Group led by Robert Bohn and John Messina, Standards to Acceleration to Jumpstart the Adoption of Cloud Computing (SAJACC) led by Lee Badger, Standards Roadmap Working Group led by Michael Hogan and Annie Sokol, and Business Use Cases Working Group led by Fred Whiteside.

The authors acknowledge the contributions from participants of the Federal Cloud Computing Standards and Technology Working Group and the NIST Cloud Computing Security Working Group.

## Table of Contents

1	Executive Summary .....	1
2	Introduction .....	1
3	Template.....	2
4	Process Oriented.....	3
4.1	APPLICATION OF 800-53-STYLE CONTROLS AND COMPLIANCE.....	3
4.2	CLOUD AUDIT ASSURANCE AND LOG SENSITIVITY MANAGEMENT .....	6
4.3	CLOUD CERTIFICATION AND ACCREDITATION GUIDELINES.....	9
4.4	NEED CLEAR E-DISCOVERY GUIDELINES.....	11
4.5	CLOUD PRIVACY GUIDELINES .....	14
4.6	CLARITY ON SECURITY CONTROL ROLES AND RESPONSIBILITIES.....	19
4.7	TRUSTWORTHINESS OF CLOUD OPERATORS.....	21
4.8	BUSINESS CONTINUITY AND DISASTER RECOVERY .....	22
4.9	TECHNICAL CONTINUOUS MONITORING CAPABILITIES .....	25
5	Focused Technical.....	27
5.1	VISIBILITY FOR CUSTOMERS .....	27
5.2	CONTROL FOR CUSTOMERS .....	29
5.3	DATA PROTECTION CONCERNS .....	32
5.4	RISK OF ACCOUNT COMPROMISE.....	34
5.5	IDENTITY AND ACCESS MANAGEMENT (IAM) AND AUTHORIZATION .....	35
5.6	MULTI-TENANCY RISKS AND CONCERNS .....	38
5.7	CLOUD BASED DENIAL OF SERVICE .....	40
5.8	INCIDENT RESPONSE .....	43
6	References: .....	50

## 1 Executive Summary

The Federal CIO introduced the 25 Point Implementation Plan to Reform Federal Information Technology Management. The plans overall objective is to achieve greater efficiency in federal IT while reducing costs through the use of shared resources and greater collaboration and consolidation. A major highlight of the 25 point plan is the shift to “Cloud First” policy whereby; each agency will identify three “must move” services within three months, and move one of those services to the cloud within 12 month and the remaining two within 18 months<sup>1</sup>.

NIST, along with other agencies was tasked with a key role and specific activities aimed at acceleration the adoption of cloud computing, including the publication of special publications and the delivery of the NIST Technology Roadmap. Accordingly, NIST has developed cloud computing security guidance to include NIST SP800-137 Information Security Continuous Monitoring for Federal Information Systems and Organizations, NIST SP800-144 Guidelines on Security and Privacy in Public Cloud Computing and NIST SP800-146, Cloud Computing Synopsis and Recommendations.

In order to achieve widespread acceptance, and ultimately adoption of cloud computing, it is necessary to address the fear and uncertainty surrounding cloud security which is well documented and discussed throughout the industry. Undoubtedly security ranks near the top of most any list of concerns or fears related to cloud migrations. This is likely due to the fact that there are few documented details that directly address how to achieve some aspects of security in a cloud environment.

This document lists perceived security impediments to the federal adoption of cloud computing, and provides descriptions of existing mitigations, or if no mitigation exists, descriptions of ongoing efforts that support mitigation are provided. If there are no existing efforts ongoing, recommendations for mitigation are provided.

## 2 Introduction

The NIST Cloud Computing Security Working group was created to achieve broad collaboration between federal and private stakeholders in efforts to review the security related issues expressed by federal managers. Through its research, the Working Group identified a list of primary concerns that are impediments to the adoption of Cloud computing. This document is focused on describing those security impediments and identifying mitigations to each if they

---

<sup>1</sup> <http://www.cio.gov/documents/25-Point-Implementation-Plan-to-Reform-Federal%20IT.pdf>

exist, and where no mitigation exists, provide description of activities that will serve to mitigate if not reduce the risk created by the issue to an acceptable level. The issues listed here have been added by members of the NIST security working group, or by NIST, and do not necessarily represent consensus by the group. The working group's charter and meeting notes can be found at: <http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/CloudSecurity>

This list is not intended to be a comprehensive list of the highest risks to federal data in a cloud environment. It is instead a practical look at the concerns expressed by federal managers and documented by a number of sources within government and private industry. The identified priorities in this document are grouped into two categories, Process Oriented and Focused Technical. Some of the impediments may have technical mitigations that may exist but require guidance or standards in place to ensure effective application. Others may have policies or mandates requiring implementation but lack the technical application in the cloud environment. The impediments discussed in this document are not listed by priority or importance. The priority and/or importance can vary widely based on environment and implementation details.

### 3 Template

**<Impediment-name>** *(as short as possible while being descriptive)*

**Description:** *A paragraph-sized (or less) description of the impediment.*

**Importance:** *A sentence or paragraph-sized rationale for why we should care: why does this impediment matter?*

**Mitigation 1:** **<mitigation-name>**

*Text here should describe a mitigation of the impediment and how sufficient the mitigation is. The first part should take one of two forms:*

*The text can be an English-language narrative of how to mitigate the impediment. This narrative might be similar in form to the text in a success scenario for a NIST SAJACC technical use case. E.g., see: <http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/UseCaseOpenAccount>.*

OR

*The text can be a short synopsis of what others have done for this mitigation (perhaps a paragraph), along with a URL pointing to their public work. If the document identified is large, the text should include some navigation information, like page numbers or section names to help the reader quickly find the relevant material.*

**Sufficiency Comment:** *Informally characterize the effectiveness of this mitigation, and whether it depends on other mitigations.*

**Mitigation 2:** ...

...

**Mitigation n:** ...

**References:** We anticipate that different mitigations will share references, so this field will ultimately exist at the end of the document in its own section but may initially exist per-section.

## 4 Process Oriented

The process-oriented requirements rely on human-centered processes, procedures, and guidance for mitigation.

### 4.1 Application of 800-53-style Controls and Compliance

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

**Technical Considerations:** The catalog of NIST 800-53 controls should be applied and tested in the A&A process for cloud systems in much the same as they would in any IT system or environment. The introduction of and extensive use of virtualization technology in the cloud environment add to the complexity to the application of some security controls. The NIST SP800-125 *Guide to Security for Full Virtualization Technologies* highlights security considerations of particular interest for virtualization solutions. These considerations are not intended to be comprehensive, nor is there any implication that security elements not listed here are unimportant or unnecessary. In addition to following the security recommendations presented in this publication, organizations implementing full virtualization solutions should also follow the recommendations from NIST SP 800-53, Recommended Security Controls for Federal Information Systems and Organizations, which defines minimum recommended management, operational, and technical controls for information systems based on impact categories.<sup>2</sup>

---

<sup>2</sup> <http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf>++

**Practical Example:** NIST 800-53 controls are applicable to securing all systems including those deployed in a cloud environment. The FedRAMP A&A process leverages the NIST security controls to assess and authorize cloud computing systems. In some migration implementations there may be components of system architectures that may be migrated to a cloud environment while some components may continue to be hosted or managed within the agency. The Risk Management Framework NIST 800-37 should be used as guidance to secure such systems and the use of cloud technology has little impact on the RMF process as it is applied to those systems beyond documenting interfaces and referencing the applicable cloud A&A artifacts.

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

**Solution Maturity:** Mature. 800-53 controls have been in use for a number of years and have been widely accepted by much of industry and required for use when securing federal systems.

**Mitigation:** NIST Risk Management Framework

FISMA and 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 step, which remains an inherent federal responsibility that is directly linked to the management of risk related to the use of cloud services.

Organizations should 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 will 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 in order that cloud

service providers can deliver appropriate evidence of control effectiveness. This evidence will be used by organizations for their 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).

FedRAMP, a GSA-led cloud security program, will satisfy many of the requirements above including defining minimum security and privacy requirements for cloud-based information systems. The security and privacy controls from NIST SP 800-53 are defined for low- and moderate-impact information being 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.

**Sufficiency Comment:** The NIST 800-53 guidelines have been broadly developed from a technical perspective to complement similar guidelines for national security systems and may be used for such systems with the approval of appropriate federal officials exercising policy authority over such systems.<sup>13</sup> State, local, and tribal governments, as well as private sector organizations are encouraged to consider using these guidelines, as appropriate.<sup>3</sup>

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

---

<sup>3</sup> [http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final\\_updated-errata\\_05-01-2010.pdf](http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf)

## 4.2 Cloud Audit Assurance and Log Sensitivity Management

**Description:** How to gain assurance that:

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

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

- Who is doing any particular monitoring or auditing task?
- Who is informed of the results of any 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. Later in this document, section 5.9 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 are any number of methods for deployment of encryption to protect the data while ensure 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 and/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 much more dynamically. In such a private cloud scenario, monitoring sensors could be deployed without the concerns iterative relocating or modifying of sensors. In a public cloud, multi-tenancy concerns could emerge depending on the data monitored and/or

captured and if data is moved dynamically providers and subscribers could 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 CSP, the customer should specify measurable monitoring and reporting standards. The contract should spell out the measures to be taken if any SLA requirements are not met. A periodic review of the SLAs and their parameters should be spelled out in the contract. Monitoring tasks also do not absolve the customer of its responsibility to monitor and audit aspects of the information system that the customer may operate or manage.

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

**Solution Maturity:** Immature. While effective monitoring solutions have been in use for quite some time the high mobility inherent to the cloud computing environment and multi-tenancy provide unique challenges on how specific data can be monitored.

**Mitigation 1:** Risk management framework

The NIST Risk Management Framework (SP800-37) 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 SP800-37'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 the 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.

**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 twenty “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 to provides for:

- “locking down” on-line systems to minimize vulnerabilities
- capturing incident information for subsequent analysis when harmful incidents occur
- discovering and exchanging related information with some degree of assurance

The CYBEX Model includes:

- structuring cyber security information for exchange purposes
- identifying and discovering cyber security information and entities
- establishment of trust and policy agreement between exchanging entities
- requesting and responding with cyber security
- 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. The need for RID and related standards in cloud computing is that organizations will need to be able 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 - cloud is global) 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 it may be for an investigation that is currently underway with a request to mitigate or stop the incident. Incidents may also be benign and require quick reporting and mitigation methods, such as configuration issues or availability issues due to IT 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](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

CSA CCM SA-14 – Audit Logging / Intrusion Detection

CYBEX Overview - [http://www.itu.int/dms\\_pub/itu-t/oth/0A/0D/T0A0D00001D0004PDFE.pdf](http://www.itu.int/dms_pub/itu-t/oth/0A/0D/T0A0D00001D0004PDFE.pdf)

### 4.3 Cloud Certification and Accreditation Guidelines

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

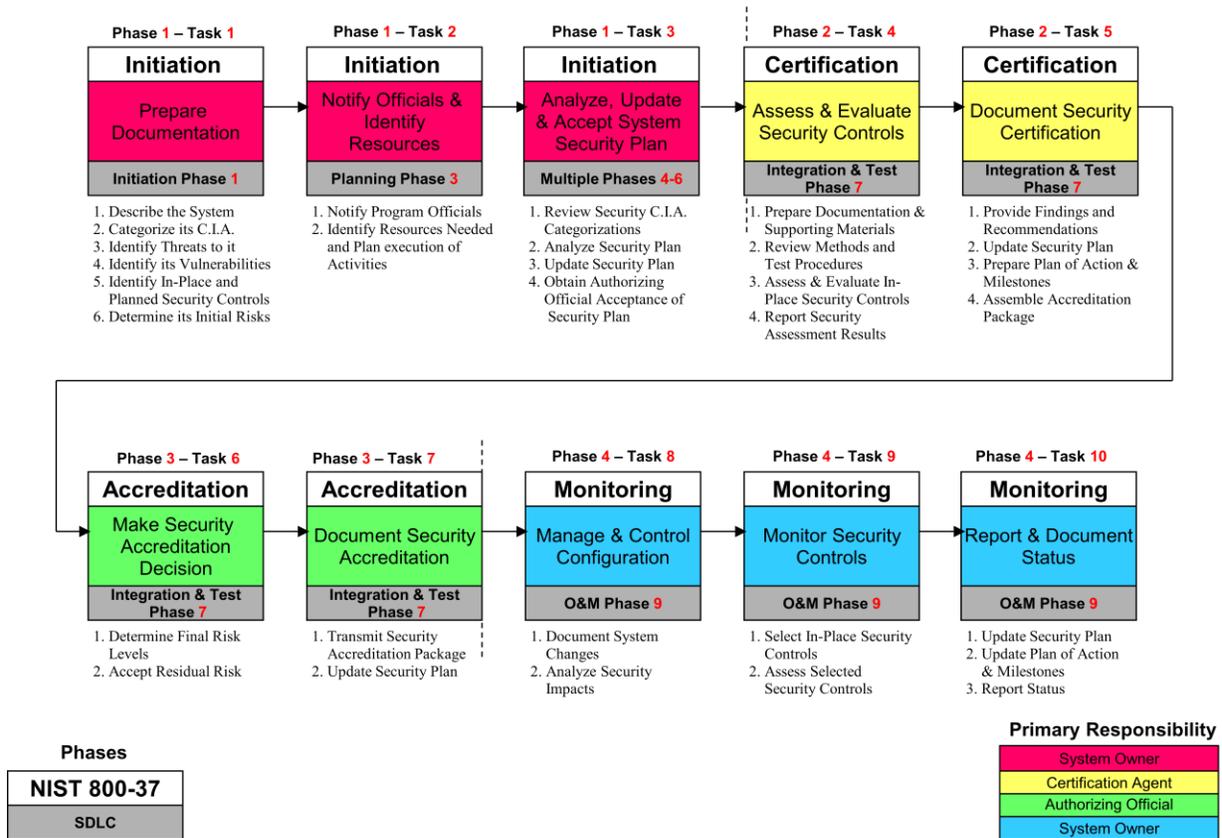
**Technical Considerations:** The risk management process changes the traditional focus of C&A as a static, procedural activity to a more dynamic approach that provides the capability to more effectively manage information system-related security risks in highly diverse environments of complex and sophisticated cyber threats, ever-increasing system vulnerabilities, and rapidly changing missions.<sup>4</sup>

**Practical Example:** The FedRAMP A&A process provides a means for assessing and authorizing cloud offerings to federal subscribers. Systems that are not migrated to a cloud environment will be subject to undergo certification and accreditation based on NIST 800-37 Risk Management Framework Guidance. A hybrid approach must be considered for those migration scenarios that include interfaces with cloud based services and applications where boundaries are established and artifacts produced by either processes are appropriately referenced.

---

<sup>4</sup> <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>

## NIST 800-37 Risk Management & Certification and Accreditation Tasks



**Importance:** Executive departments and agencies, in effectively managing 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.

**Solution Maturity:** Pending approval of OMB memo formalizing FedRAMP Program.

**Mitigation 1:** The FedRAMP program will provide a cost-effective, risk-based approach for the adoption and use of 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. <http://www.fedramp.gov>

**Sufficiency Comment:** The FedRAMP process is based on the Risk Management Framework and leverages 800-53 controls to assess systems security posture. 800-53 controls are proven adequate to secure low and moderate systems.

**References:**

<http://www.fedramp.gov>

[http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final\\_updated-errata\\_05-01-2010.pdf](http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf)

#### 4.4 Need Clear E-discovery Guidelines

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

**Technical Considerations:** Current E-discovery tools cannot be effectively implemented in a cloud environment. Present options to meet ediscovery requirements include migrating potentially large amounts of data to a non-cloud or searchable platform where tools can be effective. Such solutions are potentially costly, may be too time consuming and generally impractical depending upon the amount of data to be searched.

**Practical Example:** <sup>5</sup>Electronic discovery (E-discovery) refers to the process of identifying and producing electronically stored information (ESI) in response to civil and criminal litigation, including metadata and electronic backup materials that may be relevant to that litigation. Once such litigation is reasonably anticipated (e.g., receipt of a letter threatening a lawsuit), a party has a legal obligation to suspend destruction of such ESI, by issuing a “litigation hold” to all individuals and entities maintaining ESI on the party’s behalf. Failure to take proper and adequate steps to preserve such ESI can result in serious legal sanctions against a party. These risks may be significantly greater when using a cloud computing service, since the service provider may be unable or unwilling for technical, cost, legal or other reasons to halt routine destruction of responsive ESI, which may be maintained or commingled with data of other clients of that service for records management and disposal purposes. Likewise, the nature of cloud storage (e.g., widely dispersed servers or databases located domestically or even overseas) may complicate the ability to identify, preserve, and retrieve responsive ESI in a timely fashion, further jeopardizing the agency’s ability to meet its legal e-discovery obligations.

**Subscriber:**

---

<sup>5</sup> <http://www.cio.gov/Documents/Privacy-Recommendations-Cloud-Computing-8-19-2010.docx>

The cloud customer is responsible for preserving evidence and issuance of litigation hold notices to cloud providers who have any pertinent ESI. If required evidence is lost or damaged, the customer may be fined and/or sanctioned by the court despite any fault or failure on the part of the cloud provider, thus it is incumbent upon the cloud customer to verify that robust processes are in place to ensure preservation and facilitate ESI collection. Additionally, failure to understand where pertinent ESI is located could result in exposure of data beyond the scope of the electronic discovery request, or data belonging to customers who are not parties to the specific discovery request – possibly violating their privacy.

**Provider:**

The cloud provider is responsible for providing information to lawful authority when presented with a lawful demand for such information. This may be a one-time request for stored information or it may be a request for dynamic access to data akin to a wiretap. These requests often include a specific deadline for cooperation or surrendering of the information and the provider may face penalties if they are unable or unwilling to comply.

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

**Solution Maturity:** Involving legal/E-discovery SMEs in cloud service negotiation and mapping of business processes to identify points where Electronically Stored Information (ESI) is generated, processed, and/or stored on cloud based systems are solutions which can be implemented immediately and offer reasonable mitigation of risk. Automation of ESI collection and preservation management is potentially a future mitigation, but limitations of current tools and evolving nature of the cloud environments can result in the required interaction between the two providing a less than optimal solution at present.

**Mitigation 1:** When procuring a cloud service, customers must gain an understanding of how the cloud provider processes electronic discovery and litigation holds. The customer should acquire knowledge of key issues – such as the length of time the provider takes to enforce a litigation hold (i.e. prevent the modification and/or destruction of pertinent evidence) or respond to an electronic discovery request and what steps are required to invoke these processes, types of logs and metadata retained including lifecycles 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.

**Sufficiency Comment:** While this solution provides a reasonable level of risk mitigation, it is dependent upon the existence and availability of in-house knowledge, or the willingness to procure such, and to invest the time and resources needed. There may also be unanticipated requests for specific metadata or unstructured data not routinely collected or stored by either the cloud customer or provider that limit the availability of same and these must be handled on a case by case basis.

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

**Sufficiency Comment:** While this solution provides a reasonable level of risk mitigation, it is dependent upon the ability and willingness of the customer to map key business processes and the discovery-pertinent ESI generated, processed or stored as a result. There is a cost associated with mapping processes and storing ESI, and these should be balanced against costs of electronic discovery and litigation hold involving cloud providers, costs of hiring E-discovery specialists to handle ESI projects, and the possibility of fines and/or sanctions as the result of an inability to preserve or produce requested evidence.

**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.

**Sufficiency Comment:** Existing hypervisor platform technologies do not incorporate the necessary features to support these tasks. Providers will need to incorporate in-house or 3rd party solutions.

**References:**

Federal Rules of Civil Procedure (2010).

## 4.5 Cloud Privacy Guidelines

**Description:** This requirement addresses the need to build confidence that cloud solutions provide privacy data and PII protection.

**Technical Considerations:** Regardless of the cloud offering selected, all federal managers have the responsibility to ensure the protection of PII and privacy data. The NIST SP 800-122 Guide to Protecting the Confidentiality of Personally Identifiable Information (PII) was created to explain the importance of protecting the confidentiality of PII in the context of information security and explain its relationship to privacy using the Fair Information Practices, which are the principles underlying most privacy laws and privacy best practices. PII should be protected from inappropriate access, use, and disclosure. The NIST SP 800-122 document provides practical, context-based guidance for identifying PII and determining what level of protection is appropriate for each instance of PII. The document also suggests safeguards that may offer appropriate levels of protection for PII and provides recommendations for developing response plans for incidents involving PII. Organizations are encouraged to tailor the recommendations to meet their specific requirements.<sup>6</sup>

**Practical Example:** Federal managers must work to ensure that SLA and or contracts require CSPs to strictly adhere to existing guidance regarding privacy and PII information. Ultimately it is the system owner and approving authority that are responsible for the proper handling of such data.

**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](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.

---

<sup>6</sup> <http://csrc.nist.gov/publications/nistpubs/800-122/sp800-122.pdf>

**End-users (Data Owners):** The American public and other individual stakeholders - whose data is currently managed by federal agencies yet will potentially be put into the hands of third-parties in cloud-deployed systems - have a very strong and visceral concern about the protection of their personally-identifiable information. They have benefited, thus far, from precedent-setting tenets and principles, such as those defined in 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](http://www.irs.gov/irm/part11/irm_11-003-039.html)) to ensure 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>) Although related to non-agencies companies, these FIPS have come to be established globally as basic expectations of privacy rights by individuals.

The nature of cloud computing, however, leverages the economies of scale, geographies and resource balancing across an ecosystem of partners in the value chain that encompasses cloud service providers, brokers and other relevant roles. This broad and potentially non-transparent cloud provision network runs the risk of undoing the trust and confidence that individuals have come to have in the federal agencies with respect to the handling of their personal data. Therefore, the greatest challenge for the agencies is to be able to maintain and sustain that confidence level, whilst being able to ensure a chain-of-trust across the architecture, policies and legal structures that are established with their cloud providers.

**Cloud Customers (Federal Agencies):** 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. Individual agencies expounded upon these PIAs to issue additional guidance for the cross-agency privacy governance activities. (Reference to these individual PIAs and procedures for various agencies are provided below.)

The basic structure of the PIA is based upon knowing the data flows of personally-identifiable information, the secure protection and ultimate deletion/retention of the data across the lifecycle – and the ability to ensure, through pre-assessments, audits and monitoring activities, that any 3rd parties involved in the processing of the data were handling the data in a manner that consistent with and complied the governing controls defined by the agencies.

**Solution Maturity:** Efforts to address privacy issues in a cloud environment through SLAs are evolving.

**Mitigation 1:** Maintain confidence of the individuals about whom the federal agencies are managing personal data. Cloud Customers (Federal Agencies) should consistently assess the scope of the Personally-Identifiable Information that they manage within their systems through the 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 strict requirements of HIPAA and HITECH requirements) should be classified in such a manner that proper evaluations as to appropriate migration to a public, hybrid public/private or private cloud can be performed by each agency. Education and outreach efforts should be done to ensure individuals that the privacy controls that currently exist to protect their personal information will be retained despite migration to a cloud environment. (Contingent upon the ability to reach Mitigation 2)

**Sufficiency Comment:** PIA questions and PIA processes have been well-established. Cloud-specific implementations of the cloud-deployed data flow will impact the ability to assess all sections of the requisite PIAs, however, without the involvement and assurance of the Cloud Provider. The PIAs should enable the Cloud Customers to at least assess which data can be migrated to the cloud in a low-risk manner.

**Mitigation 2:** Ensure that the Cloud Providers protect the personal information to the requisite levels of protections a) that have already been established for federal agencies' existing systems and b) that are being finalized to define cloud-specific controls for federal agencies in the near-term. (These are references to NIST 800-53, NIST 800-53 Appendix J and NIST 800-146 – in varying degrees of draft form but with target completion dates by EOY). Additionally, Service Level Agreements and other legal instruments need to be established between the Cloud Customer and the Cloud Provider, although ultimately the Cloud Customer is still responsible for the protection of the data.

**Sufficiency Comment:** Cloud Customers (Federal Agencies) must limit their choice of Cloud Providers to those that will allow them the degree of assurance – and assurance testing – that is required to ensure sustained protection of the personal information in the cloud. – Controls are being defined (to be completed at EOY 2011). Alignment of Cloud Providers to those controls – and to providing “proof” of implementation of those controls is still in question.

**Mitigation 3:** Ensure that cross-jurisdictional Privacy issues are addressed and incorporated in any agencies' cloud deployment if they are going to be collecting, managing, retaining or otherwise processing data that comes under the scope of global Data Protection regulations. These global requirements may fall outside the technology scope of the NIST Reference Architecture and may not even be consistently automated across IT processes (for instance, manual collection of data that ultimately ends up in a database hosted in the cloud). Although alignment activities are taking place between the US and the EU, for instance, to address disparities in privacy-relevant legislation, there remain areas of non-agreement. This would require a proactive migration strategy that addressed physical location of the data in the cloud, awareness to and agreements about cross-border transfer of any relevant data and the ability to

ensure End User (Data Owner) involvement in the modification or requested deletion of their data (where appropriate). These are established legal rights of citizens outside of the US; yet awareness to these intricacies is not well-known across many of the agencies and system owners today. Therefore, a strategy to address these unique situations (for which there have been real-world situations in which data access was withheld/disrupted for the federal agencies due to these legal differences) will need to be addressed. Non-political mitigations (or, at least, stop-gap measures in the interim) would include geo-tagging of data to ensure that relevant data does not cross international borders, stricter or modified metadata and data classifications etc.).

**Sufficiency Comment:** Beyond Safe Harbor and Model Contracts established between the EU and the US, no other tenets exist to address the transfer of data between the two geographic jurisdictions. Cloud deployments may have to be continent-specific initially until these policies/procedures are established. The scope of federal agencies that will be impacted by international data and geopolitical considerations will need to be understood.

**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](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](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](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](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](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](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](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](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](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](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://fex.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>

## 4.6 Clarity on Security Control Roles and Responsibilities

**Description:** Who (among cloud actors such as customer and provider) should be responsible for the implementation of required security controls listed in section 1.1? It seems that the actor most able to observe and configure specific a portion of a cloud implementation would be in the best place to implement a relevant control.

**Technical Considerations:** Careful planning of security activities should be considered early in the design of a migration or cloud deployment effort. It is critical to ensure that adequate ability to meet the operational responsibilities necessary to maintain a systems security posture in planned and executed regardless of whether a system component is deployed to a cloud or self hosted.

**Practical Example:** When migrating applications and services to a cloud environment a combination of deployment and service models may be used to satisfy requirements. Such variation requires careful examination of the roles and responsibilities required to secure and monitor cloud and/or internally hosted systems.

**Importance:** The data owner (typically cloud customer) is responsible for compliance with laws and regulations including the proper security controls around their data, regardless of its location or the involvement of other parties.

### **Customer:**

When customer data is off-premise and under the control of a third party, such as the cloud provider, cloud broker, or cloud carrier, the ability of the data owner to implement security controls is often limited. In cloud computing environments, the implementation of controls is going to be largely dependent upon the type of service (IAAS, PAAS, SAAS), type of cloud (Private, Public, Community, Hybrid) type of control (i.e. physical versus logical) and the specifications of responsibility delineated in the cloud contract. Customers must choose services accordingly and understand the risks and limitations of third party control – i.e. customers can outsource the functionality of a role/responsibility, but will still hold the legal liability for failure to secure data.

### **Provider/Broker/Carrier:**

Cloud providers/brokers/carriers have increasing responsibilities for implementing and maintaining security depending on the cloud types and service types offered. For example, physical security controls in the cloud environment will have to be implemented by a party which has access to the physical property where the assets reside – a role likely performed by the cloud provider, broker, or carrier. Providers/brokers/carriers should provide statements of the security controls environment implemented and maintained by them for each of the cloud architectures and services offered.

**Solution Maturity:** Cloud Security Alliance has published a Cloud Controls Matrix which provides a good reference point and denotes applicability to cloud service type (IAAS, PAAS, SAAS) and scope (provider, tenant). The solution effectiveness is dependent upon the cloud customer's willingness to understand the provider posture on security controls for the services offered and the ability of potential providers to mitigate specific risks and meet customer requirements. The solution may also be dependent upon customer willingness to choose a cloud service or architecture (plus availability of same) that permits more direct control to address specific security requirements.

**Mitigation 1:** Provider-subscriber guidelines

Document roles and responsibilities definitions or guidelines for cloud provider and consumer/subscriber— Clarity about how responsibility for protection of information remains with a system owner but the terms of a contract between a system owner and a cloud provider can place an obligation on the provider to protect information. To satisfy due diligence and due care requirements for securing their data, cloud customers must ensure the contract with the cloud provider/broker/carrier specifies the responsible party (or parties) for implementing and maintaining security controls and provides rights of action for failure to implement or maintain same.

**Sufficiency Comment:** The solution provides a reasonable degree of risk mitigation but is dependent upon customer willingness to define security roles/responsibilities and negotiate same with the cloud provider(s).

**Mitigation 2:** Cloud type/service selection

In cases where a larger degree of direct control over security roles/responsibilities and the ability to implement security controls is needed, cloud customers may consider utilization of a type of service (i.e. PAAS, or IAAS instead of SAAS) and/or a cloud type (on-premise versus off-premise, private versus public cloud) which will allow that requirement to be fulfilled.

**Sufficiency Comment:** The solution can provide reasonable degree of risk mitigation but there may be increased cost in choosing services and/or architectures that allow more direct control and these costs should be weighed against security requirements and acceptable risk levels.

**References:**

CSA Cloud Controls Matrix.

## 4.7 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.

**Technical Considerations:** Mitigations for this requirement should be addressed via contractual mechanisms and policy considerations.

**Practical Example:** Not being able to fully trust cloud operators is one of the primary reasons for the reluctance to adopt cloud services and solutions. Despite this, cloud operators are still hesitant to divulge all their security practices in fear of giving away intellectual property. This is problematic for cloud service consumers that have invested significant amount of resources over the past several years (with Sarbanes-Oxley, FISMA, etc.) gaining transparency into their own IT infrastructure.

**Importance:** For cloud service consumers, it is critical to know the security practices of their cloud operators in order to maintain and improve on the security of their data and their IT environment. This means that cloud consumers need to know what their cloud operators are doing and if their cloud operators are effectively performing those functions. In addition, cloud consumers must be able to randomly and independently verify their cloud operators' practices.

**Solution Maturity:** The practices and technology currently exist to address this impediment. Security control frameworks exist (such as CoBIT, ISO 27000, NIST SP 800-53, etc.) in order to architect and assess the security posture of an organization. NIST-validated SCAP-compliant tools exist to identify and verify vulnerability and configuration data, and virtualization management modules also exist to provide additional insight into the virtual infrastructure. The shift that may need to occur is tailoring the frameworks and the technology for the cloud and a cultural shift to provide the needed transparency.

### **Mitigation 1:** Operator human resources practices

Through standard and effective SLAs, consumers can specify background screening requirements for operator staff hiring, and can require regular training so that operator employees (including contractors and third party users) understand their responsibilities and use best practices. Consumers and operators can also agree to apply separation of duty and to monitor unauthorized activities by malicious insiders.

**Sufficiency Comment:** There is still a need to develop and adopt consistent SLAs of high quality and completeness for cloud services.

### **Mitigation 2:** Operator self-certification and third party verification

To gain consumers' trust, cloud operators seek self-certification of compliance with legal and regulatory requirements (for example, SAS 70 or ISO 27002 compliant). There is also a significant increase in third party independent audit of operators' information security management from policies all the way to specific controls.

**Sufficiency Comment:** A consumer still needs to verify if its specific requirements for security are met by the operator and how operator's measures stack up against other operators.

**Mitigation 3:** Operator transparency

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

**Sufficiency Comment:** Refer to section 4.1, Lack of Visibility for Customers, for more details.

**Mitigation 4:** Reviews from other consumers and industry groups can serve as a useful resource for research on the services and offered and the general reputation of cloud providers. However, one should remain mindful that much of the information that may be available online may be subjective or unreliable so sources of information should be carefully weighed when considering this type of information.

**Sufficiency Comment:** Potentially reliable but is still a rapidly expanding online resource.

**References:**

TBD.

## 4.8 Business continuity and disaster recovery

**Description:** Even 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 adds more complexity to IT infrastructure and obfuscates responsibility between cloud provider and customer. There is an elevated concern of business continuity and disaster recovery in this new paradigm.

**Technical Considerations:** While migrating a service or application to a cloud environment may provide some inherent level of redundancy, system owners must ensure that cloud providers provide documented DR and COOP plans.

**Practical Example:** As pointed out in NIST SP 800-144 (section 4.8 Availability), outages due to high level disruptions do occur in cloud computing services despite redundant architectures designed for high availability. In addition, the value concentration in clouds makes them the preferred targets of malicious attacks. Consequently, it is important to plan and implement business continuity and disaster recovery in consideration of the characteristics of cloud computing:

The intricate business relationship and managerial responsibility (ownership and governance) in cloud computing complicate business continuity planning and implementation and effective disaster recovery.

The large scale, complexity and dynamics of cloud infrastructure make it even more difficult for maintaining business continuity and performing disaster recovery.

Comingled data and services result in comingled responsibilities and contingency plans.

Clustering, quick provisioning and on-demand self-service provide conduits for disruptions to propagate through the cloud environment.

However, economies of scale in cloud computing have the potential to provide a full range of backup sites and to carry out effective testing and exercises in order to validate security controls of contingency planning.

**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.

**Solution Maturity:** Documented security policies and procedures, clustering technologies, alternate sites and backup have all been used in traditional IT contingency planning and implementation. How effective they are in cloud computing (with order of magnitude increase in scale, complexity and dynamics) remains to be seen. Clear roles and responsibilities and close collaboration between cloud provider and customer are necessary for any effective business continuity and disaster recovery. It is also necessary to coordinate incident activities with contingency planning activities between cloud provider and customer.

**Mitigation 1:** Consistent Policies and Procedures

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 customer' contingency plan for individual application and/or system into cloud provider's overall contingency plan.

Establish management succession and escalation procedures between cloud provider and customer.

Reduce the complexity of the recovery effort.

**Sufficiency Comment:** To be effective, this mitigation needs unprecedented close collaboration between cloud provider and customer and among customers. Since recovery priorities are based on mission requirements and data sensitivity, applications should be partitioned according to sensitivity as suggested in mitigations of multi-tenancy.

**Mitigation 2:** Clustering and redundancy

Shared storage clusters.

Hardware level clustering.

VM clusters.

Software clustering (application servers and database management systems).

**Sufficiency Comment:** The key to clustering security is isolation (logical separation), and its sufficiency is addressed in the mitigations of multi-tenancy.

**Mitigation 3:** Alternate sites and backup

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.
- Supplement cloud provider's backup schemes with customer's non-cloud sites.

**Sufficiency Comment:** A hot backup is a preferred solution to down time for high available systems, but may be expensive compared to alternatives.

**Mitigation 4:** Effective testing and exercises

Test the contingency plan periodically to verify its effectiveness (including personnel training) and if it's updated to reflect recent changes.

Perform contingency plan testing and exercises against high level disruptions to discover deep-rooted risks.

Provide production-like testing environments for contingency plan testing and exercises.

**Sufficiency Comment:** Agencies are constrained to simulate level 3 and 4 disruptions in their core production environments. And there are resource limitations to provide production-like testing environments. As a result, contingency plan testing and exercises are usually tabletop scenario exercises or at best consist of some failovers in peripheral equipments.

**References:**

NIST SP 800-34 Rev 1

NIST SP 800-53 Rev 3

NIST SP 800-144

Federal Cloud Security Guidelines (2011)

## 4.9 Technical Continuous Monitoring Capabilities

**Description:** There exists a lack of technical continuous monitoring capabilities necessary to support monitoring of cloud environments. This need can be especially challenging with multi-data center clouds using 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 even further aggregated to support higher level operational pictures and management decisions. The data needs to also reflect both the security posture of the cloud and the security posture of customer's use of the cloud.

**Practical Example:** Currently questions remain regarding how specific information regarding the security posture of the environment in which a particular subscribers data may reside. Current monitoring solutions were not designed for highly mobile environments or multi-tenant environments with potentially largely disparate monitoring and reporting requirements.

**Importance:** For cloud providers, it is critical that they be able to gain situational awareness of their cloud environment and to provide evidence to their customers that the cloud infrastructure is secure. It also may be important to provide customers feedback on the security of their use of the cloud.

**Solution Maturity:** Much of the foundation for eliminating this impediment exists in the area of security automation standards. This is especially true for asset, configuration, and vulnerability management. However, the higher level model for providing situational awareness is still in the draft stages.

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

CyberScope is a secure web-based application requiring two factor authentication that collects both automated and manual data from federal agencies, used to assess and report on the agencies IT security posture. CyberScope receives both live data feeds and data entry by agency staff. It 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.<sup>7</sup>

The types of information that OMB requires to be reported through CyberScope are broader in scope than the status of individual assets, which are the focus of the CAESARS reference architecture. Nevertheless, the CAESARS reference architecture can directly support the achievement of some of the OMB 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 OMB reporting objectives is full situational awareness of all agency IT assets.<sup>8</sup>

**Sufficiency Comment:** Once adopted and implemented the CEASARS framework will allow agencies to implement CM more rapidly by leveraging CM compliant tools eliminating the need for custom integration efforts.

## References:

---

<sup>7</sup> <https://www.cippguide.org/2010/11/02/cyberscope/>

<sup>8</sup> <http://www.dhs.gov/xlibrary/assets/fns-caesars.pdf>

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](http://csrc.nist.gov/publications/drafts/nistir-7756/Draft-nistir-7756_feb2011.pdf)

## 5 Focused Technical

The focused technical issues are amenable to automated mitigation mechanisms.

### 5.1 Visibility for Customers

**Description:** Today, cloud customers have very limited visibility into security measures, incident alert and audit information generated by providers relating to cloud resources and customers' data and applications. Most of this information is used internally, and providers are reluctant to share it in fear of divulging proprietary information and introducing an avenue of attack. Serious questions remain:

- How can cloud customers observe their workloads to be aware of their security, privacy, compliance, health and general status?
- And how can customers instruct the cloud as to what kind of information they are interested in?

Technically, what was once a rack of physical machines interconnected via a network switch may be consolidated in the cloud into a single virtual machine hosting many applications with a virtual switch enabling communication among virtual machines. Movement of virtual machines within the same physical machine or to different ones in different data centers may occur dynamically, making monitoring and management of virtual machines much more complex. The prevalent monitoring and management tools are mostly piecemeal and specific to each provider. To achieve the needed level of transparency in the cloud, there are needs and undertakings for unified tools that deeply monitor network, computer, storage resources and their virtual counterparts and broadly manage provisioning, configuration, change control and compliance.

**Importance:** Cloud customers are ultimately liable for security and privacy. Specifically, security/privacy/compliance regulations state that the liability cannot be outsourced, and providers do not take on full responsibility for security in their SLAs. Also customers are mostly liable for health and general status. Providers usually compensate for the subscription cost of an outage, but not for the actual damage or loss of business. Customers, seeking confidence in the cloud, will favor providers who can build high visibility and provide greater transparency (refer to NIST SP 800-144, Guidelines on

Security and Privacy in Public Cloud Computing and NIST SP 800-146, Draft Cloud Computing Synopsis and Recommendations).

**Solution Maturity:** Monitoring and management tools are currently being unified to be more effective in the cloud environment. There are also efforts going on for multi-cloud mashup (e.g., Web 2.0 based). Standards are maturing for alert exchange, and standards are emerging that let customers instruct the cloud as to what kind of information they are interested in (e.g., CloudAudit).

**Mitigation 1:** Agreement and cooperation between providers and customers to implement customer-specific requirements and to provide transparency

As pointed out in 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 means should be specified for customers to instruct the cloud as to what to monitor and to be alerted about.

**Sufficiency Comment:** Still in debate is the appropriate level of abstraction and summarization of monitoring/audit results that protects providers' sensitive information and satisfies customers' needs for transparency.

**Mitigation 2:** Effective Monitoring

- 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 in the cloud architecture using cloud APIs.
- Centralized monitoring and analysis of configuration files and log files and alert capability.
- SCAP is an alert format standard mandated by US government that can help providers push alerts to customers in a standard format.

**Sufficiency Comment:** Once fully standardized and faithfully implemented, this may be sufficient. But sufficiency depends very much on what kinds of monitoring providers will support.

**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 an open, extensible and secure APIs.

**Sufficiency Comment:** It is still evolving at this stage (unclear whether there are actual implementations yet).

**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).
- Tools that push monitoring to customers and allow customers to configure what's interesting to them (e.g., Amazon CloudWatch).

**Sufficiency Comment:** Once multiple clouds are involved, there is a strong need for flexible ways to further unify monitoring and management tools.

**Active WG Members:** Ulrich Lang, others TBD.

**References:**

- [www.cloudaudit.org](http://www.cloudaudit.org)
- <http://scap.nist.gov>
- <http://aws.amazon.com/cloudwatch/>
- <http://www.emc.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

## 5.2 Control for Customers

**Description:** Today, customers have very limited control over security policies enforced by cloud providers on their behalf. There is also very little automation available to help customers to implement technical controls (policies) across their cloud applications. Questions are raised:

How can cloud customers maintain effective control over their workloads even though the protection mechanisms and the locations of workloads may not be known to them?

How can customers instruct cloud as to what kind of security policies they want to be enforced at various control layers such as data object, VMs/Applications, virtual network and geographic location?

**Importance:** Moving IT to cloud 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. Further, customers and providers need to be able to define and enforce security policies at various control layers in consideration of:

Much of the security policies are specific to customers, not providers. In other words, the policies depend on the specifics of customer business, business processes, regulations, standards, etc.

Providers cannot independently know and enforce security requirements because these requirements are customer-specific.

There are no/few tools available for customers to specify security policies in their terms and to feed them into cloud for control enforcement.

Available tools are piecemeal and provider-specific, which renders a multi-cloud mashup unmanageable.

**Solution Maturity:** Any effective controls depend on effective identity and access management (section 4.5). In addition, fine-grained access control over data objects has long been used in traditional IT and is mature enough to be applied in clouds. Controls at VMs/applications and virtual network layers are provider-specific and less verifiable. Standards such as XACML start to be used to instruct access policies across cloud applications, for example, for Web services proliferated by SOA, Web 2.0 and cloud. Model driven security possesses potential to integrate dynamic security requirements into cloud through low level abstraction. Gaining visibility (section 4.1) and continuous monitoring (section 3.9) also help to establish customers' confidence of security controls in cloud.

**Mitigation 1:** Selection and Use of Appropriate Cloud Models

Public, hybrid, community and private with increasingly greater customer control over tenants.

SaaS, PaaS and IaaS with increasingly greater customer control over infrastructure.

Externally hosted and internally hosted with increasingly greater customer control over location.

External provider operated, outsourced and internally operated with increasingly greater customer control over personnel.

**Sufficiency Comment:** Once the appropriate cloud models are selected based on control needs, customers and providers need to cooperate further to specify and enforce protection mechanisms at various control layers.

**Mitigation 2:** Control of Data Objects

Establish and maintain data object ownership.

Use 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.

Keep change history.

Manage data lifecycle (section 4.3).

**Sufficiency Comment:** Access control over data objects is widely used and mature. Customers need to verify that providers protect data at rest, in transit, and especially when it is processed.

**Mitigation 3:** Control of VMs and Applications

Perform and verify VM hardening based on Federal and/or generally accepted standards.

Use automated tools to assess and report baseline security configurations and patch updates of VMs (especially long dormant and roll back ones).

Sanitize and protect virtual machine images.

Secure APIs (based on externalized, unified and fine-grained authorization management, for example) to allocate, start, stop and de-allocate VMs/applications.

**Sufficiency Comment:** This mitigation is mostly provider-specific. A more joint-up enforcement stack is required to be fully effective.

**Mitigation 4:** Control of Virtual Network

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

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

**Sufficiency Comment:** Visibility, maturity and faithfulness of security protection in virtual network are provider-dependent and still evolving.

**Mitigation 5:** Control of Geographic Location

Identify and select data center locations.

Enforce and verify security and compliance constraints for trans-border data flow in self service, data replication, workload management, and cloud bursting.

**Sufficiency Comment:** Various types of security and privacy regulations and laws at national, state and local levels make compliance a complicated issue. Cloud providers are at very early stage in implementing geographic restriction and in achieving right balance between on-demand characteristics and compliance constraints.

**References:**

[www.modeldrivensecurity.org](http://www.modeldrivensecurity.org)

[www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=xacml](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml)

### 5.3 Data Protection Concerns

**Description:** The loss of confidentiality, integrity, or availability for customer data can impose a wide variety of impacts. Cloud customers need to understand the extent of the data protection that a cloud offers (even if limited) so that they can make rational risk-based decisions about when to store data in a cloud. FIPS 199 provides a categorization scheme (low-impact, moderate-impact, high-impact) for data and systems and describes impacts in terms of confidentiality, integrity, and availability. The suitability of a cloud to store or process customer data varies both on the impact level of the data and on the assurances that the cloud can offer that the data is protected. The technical ability to protect data varies widely depending on how the data is accessed. A number of access scenarios are possible, including:

**In transit to or from a provider:** Data that a customer wishes to upload into a cloud must be protected in transit; similarly, data that a customer wishes to download from a cloud must be protected in transit.

**Passively stored with no shared access:** Data that should be accessed only by the originating customer needs to be protected against access attempts by all other entities, while preserving the availability for the originating customer.

**Passive stored with selective shared access:** Data that that should be accessed only by entities that have been authorized by the originating customer for specific access modes (e.g., read, write, delete) needs to be protected against access attempts by unauthorized entities or accesses in unauthorized modes, while preserving availability for authorized customers.

**Passively stored public access:** Data that should be accessible anonymously in some authorized modes (e.g., read) but that should not be accessed in other modes except by authorized customers.

**Actively processed:** Data that is accessed by a computation running in a cloud (e.g., a VM, PaaS, or SaaS application) but that otherwise may not be shared or may be shared only selectively.

**Account termination:** Data that should be maintained for a fixed period of time.

**Deletion:** The authorized erasure of customer data.

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

**Solution Maturity:** Immature and evolving.

**Mitigation 1** Data management covers the concepts of ensuring the integrity and availability of information in transit, being processed, and in storage within the cloud. Another unique aspect of cloud usage is data segregation and isolation, as data from other sources could possibly be comingled between organizations. Therefore, encryption of the data is imperative to ensure confidentiality of data.

**Sufficiency Comment:** Methods and techniques for accomplishing encryption are evolving. Cloud computing offers unique challenges to encrypting data as it exists in various states.

Data at rest:

- Prevent tampering of data, copying, altering, or deleting from the original
- Applying hashes or certificates to ensure authenticity
- Method(s) to allow searching and updating encryption algorithms

Data processing:

- Treatment of information processed within the cloud
- Processes in place to prevent data leakage

Data in transit:

- Remote VPN connections/Public ISP access
- Mobile wireless devices
- Use of intranet, cross agency or cross department
- Using direct encryption of data, with hashing/signatures<sup>9</sup>

**Mitigation 2:** Employ a comprehensive Information Lifecycle Management Program to ensure the protection and proper handling of data throughout then various phases of data management. The Cloud Security Alliance has developed a useful model of information lifecycle management, with phases of Create, Store, Use, Share, Archive, and Destroy<sup>10</sup>. This is illustrated in *Figure 1: Information Lifecycle Management Phases*. The security requirements in this lifecycle are dictated by the types of data.



**Figure 1: Information Lifecycle Management Phases**

---

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

<sup>10</sup> <http://www.cloudsecurityalliance.org/csaguide.pdf>

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, but other models in the private sector will be useful for formulating additional pertinent controls.

**Sufficiency Comment:** The lifecycle management approach is proven to be effective in ensuring the proper handling of data. The cloud environment introduces additional challenges to this model such as delete and destroy.

## 5.4 Risk of Account Compromise

**Description:** Benefits of cloud computing include its easy accessibility. A customer can use cloud computing services anywhere he/she has Internet access. However, the Internet is full of threats such as phishing, pharming and spyware, whose purpose is to steal usernames and passwords (credentials). Facing this Internet security threat environment, customers adopting cloud computing are concerned about how user accounts are protected from hijack to avoid misuse.

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

There is additional attack surface exposure due to increased complexity and dynamics in the infrastructure.

There are new APIs/interfaces that are less battle-tested.

A hijacked account may be used to steal information, manipulate data, and defraud others under the customer's identity.

A hijacked account may be used to attack other tenants as an insider in the multi-tenancy environment.

**Solution Maturity:** Mitigations such as strong authentication, encrypted credentials, and secure APIs/interfaces have been used to protect user accounts from hijack. But, as pointed out in the Symantec Security Threat Report, the easiest vulnerability for attackers to exploit is our trust of friends and colleagues. Users tend to click the links and attachments in an email they received from a trusted source (perceived or real). Thus, these mitigations should complement effective user training and proactive monitoring and auditing of unauthorized activities.

**Mitigation 1:** Strong authentication

Enforce strong password usage and change passwords periodically.

Use multi-factor authentication.

Prompt users for their passwords during sessions especially when there are suspicious events.

Allow logins coming only from a white listed address range.

Use biometrics.

**Sufficiency Comment:** There is already a password proliferation, resulting in security compromising behavior and increased burden on help desk. Possible considerations in cloud computing include integration with customers' existing identity management processes and single sign-on (SSO).

**Mitigation 2:** Encrypted Credentials

Provide a dedicated VPN.

Use HTTPS and LDAPS.

Enable secure cookies.

Use strong cryptographic PKI keys.

**Sufficiency Comment:** The effectiveness of encrypted credentials depends largely on secure key management.

**Mitigation 3:** Secure APIs/interfaces

Provide common security models for cloud APIs/interfaces (e.g., WS\*, WS-I, SAML for web services).

Protect application security using secure APIs/interfaces (e.g., input validation/escaping/encoding against injection exploits such as SQL injection and cross site scripting).

**Sufficiency Comment:** Cloud APIs/interfaces are still evolving (all the way up to the level of cloud federation).

## 5.5 Identity and Access Management (IAM) and Authorization

**Description:** Unauthorized access to sensitive information in public, private and hybrid clouds is always a major security concern for both customers and providers. Even though identity and access management (IAM) has long been used to manage users and their access to resources in the traditional IT environment, the question is what kinds of

IAMs in terms of identity proofing, strength of credentials and access control mechanisms should be implemented for effective federal cloud-based authentication and authorization.

**Technical Considerations:** Due to the broad network access characteristic of cloud computing, remote authentication of individual user is a common practice and it presents a technical challenge to establish the needed level of confidence in user identity commensurate with information sensitivity. In the very dynamic cloud environment, the control of who can access what, where and when should employ the context-aware enforcement of authorization policies and should be standards-based to be scalable and robust in multi-clouds. Over the years, federal agencies have made a sizable technical investment in identity and access management (IAM) systems in the traditional IT environment. To take full advantage of this investment, agencies would like to see these IAM systems to be integrated with the cloud services, especially during the initial adoption stage.

**Importance:** High. The identity and access management (IAM) should be effective and scalable, especially when multiple clouds are involved. The effectiveness and scalability will only be achieved through the seamless extension of controls from agencies to the cloud. Key to the extension of controls is to establish trust relationships between cloud customers and cloud providers, and, possibly, identity, credential and attribute providers.

**Solution Maturity:** Immature. Although the foundations for solutions exist, practical federal-wide high-assurance identity and access management (IAM) systems are not publicly available yet, especially when suitable for sensitive data in multi-clouds. Currently, agencies are only testing the water with small projects and less sensitive data.

**Mitigation 1:** Using provider's IAM for cloud-based services while keeping agency's IAM for internal systems.

This is the most straightforward implementation, but there is an initial creation effort of user identities and account provisioning in the cloud.

**Sufficiency Comment:** The effectiveness depends on how well the provider's IAM matches the agency's security requirements. There may be duplication of IAMs and difficulties in replicating and synchronizing authentication and authorization policies from agency to the cloud. If provider's IAM is proprietary and domain-centric, it may encounter difficulties to scale up in multi-clouds.

**Mitigation 2:** Integrating agency's IAM with cloud-based services

Cloud providers 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.

**Sufficiency Comment:** The effectiveness depends on how secure agency's IAM is integrated into the cloud, especially at endpoints. There may be a need and associated costs to integrate incompatible IAMs. It may be difficult to scale up, especially in multi-clouds.

**Mitigation 3:** Claim-based Federated Identity Management

This is a single sign-on (SSO) solution that an external identity system is depended upon to give cloud services all the 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 then make authentication and authorization decisions based on the supplied claims. At present, there are many types of issuing authorities, from domain controllers that issue Kerberos tickets, to certificate authorities (CAs) that issue X.509 certificates.

The solution may also use unifying standards such as SAML to exchange authentication and authorization decisions between security domains (for example, identity providers and service providers).

**Sufficiency Comment:** The effectiveness depends on how trustworthy and secure the issuing authorities and key managements are. Currently, Kerberos systems are mainly used within a security domain and SAML is mainly binding with web services.

**Mitigation 4:** Digital Identity

The solution uses the 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 take their digital IDs to cloud services and to be validated on the spot similarly to the way driver's licenses are used in the real world. This solution exhibits the scalability and flexibility needed in multi-clouds.

**Sufficiency Comment:** The effectiveness depends on how trustworthy the digital identity issuers are and how securely the digital identities are managed. There may be complications when the digital identity is phished, compromised, or simply lost.

**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, the solution employs 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 agency to specify and deploy access control policies to match its mixture of assets and business functions, and then to plug in

additional policies as business and infrastructure evolve. The unity is meant to express access control policies in a single language and format in multi-clouds.

**Sufficiency Comment:** XACML needs to be more widely applied beyond web services.

**Active WG Members:** Bill Butler, others TBD.

**References:**

- NIST SP 800-63 Rev. 1
- 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

## 5.6 Multi-tenancy Risks and Concerns

**Description:** Cloud computing provides the potential of cost saving through resource sharing. Different tenants use services on the same cloud simultaneously. As a result, there are warranted security concerns:

A tenant may have access to other tenants' virtual machines, network traffic, actual/residual data, etc.

A tenant may impact the normal operation of other tenants, steal their data, steal their identities, etc.

**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.

**Solution Maturity:** Physical separation is a mature enough practice even in traditional IT environments. Despite key management limitations, data encryption has been accepted in eCommerce and Federal IT systems. Application partitioning facilitates putting critical components in more secure environments, but its assurance of security needs to be further verified. Logical separation in cloud computing remains a general concern and its maturity will be vendor dependent in the near future. Based on their maturity levels, it is suggested to use

physical separation or a combination of data encryption, application partitioning and logical separation (defense in depth) to address the risk of multi-tenancy. All mitigations should complement the identity management and access control best practices.

**Mitigation 1: Data encryption**

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 key separately from data with higher privileges and preferably accessible only through procedures/programs.

Change key periodically and data unencrypted and re-encrypted with the new key.

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.

**Sufficiency Comment:** By itself, encryption is not sufficient to mitigate the risks from multi-tenancy. Encrypted data is not as vulnerable to disclosure as plaintext data but is still vulnerable to loss and possibly corruption. Key management must be performed correctly and at scale or the cryptography does not provide value. Performance may be affected.

**Mitigation 2: Application Partitioning**

Separate access control functionality from business processing functionality.

Separate logic processing functionality from data access functionality.

Separate user functionality from system management functionality.

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).

**Sufficiency Comment:** By itself, localization is not sufficient to mitigate the risks from multi-tenancy but it can localize the reach of security risks and hence reduce risks.

**Mitigation 3: Logical separation**

Support holistic logical separation of the resources at all the layers: computing (virtualization), networking (vSwitches and vLANs) and storage (logical separation of files with access controls).

Secure the virtualization server (hypervisor isolation settings to limit accesses).

Secure the virtual network by working hand-in-hand with the physical network security, especially against man in the middle attacks (MAC spoofing and ARP poisoning).

Harden the Virtual Machine (VM) so that the virtualization layer is not exposed to attack.

**Sufficiency Comment:** If logical separation is faithfully implemented, it addresses much of the multi-tenancy impediment. The difficulty is in achieving assurance that an implementation is correct.

**Mitigation 4:** Physical separation

Special virtual environments with physical separation of the full-stack cloud infrastructure can be provisioned to customers with special security requirements.

This kind of special virtual environments can be provisioned in a cookie cutter way to respond to increasing demands.

**Sufficiency Comment:** Consider private cloud for even higher demand for separation/isolation.

**References:** [01] through [08] Inserted in section 5.

## 5.7 Cloud Based Denial of Service

**Description:** Because cloud customers depend on functional networks to access their Resources, and because networks are often not under the control of customers, there is a risk that the cloud may not be reachable. Note: high latency on the cloud carrier network may have the same net effect as DoS.

In spite of clauses in SLAs implying high availability and minimal downtimes for subscribers, service or utility outages are inevitable due to man-made causes (e.g., malicious attacks or inadvertent administrator errors) or natural causes (e.g., floods, tornados, etc.) (NIST 800-146, p8-3). Cloud providers depend on functional networks to provide resource access to cloud consumers, and because cloud carrier networks are often not under the control of cloud providers, there is a risk that the cloud may not be reachable or infrastructure and/or data not available or severely degraded. It is possible that the cloud based data is available but the network is not. It is possible that the cloud carrier network to be available but the cloud provider service and data are not available. Both scenarios equate to cloud service denial from the cloud consumer perspective.

**Importance:** Denial of Service (DoS) attacks have existed since the early days of network computing, but cloud computing has significantly increased the attack surface. Internally accessed applications become remotely accessible as cloud services, thus the exposure to DoS. Due to multi-tenancy, DoS can be launched by insiders through shared resources, for example, via side channel attacks. Malicious users can initiate distributed DoS using the vast resources of

cloud to the level of severity never seen before. It is important to address DoS issues in the cloud environment.

**Solution Maturity:** One prevalent solution to withstand DoS is to provide diversity and redundancy in networking and data processing and to use multiple cloud service providers (refer to section 3.8, Business Continuity and Disaster Recovery). Another traditional solution still beneficial in the cloud environment is to use modern network devices that recognize DoS signatures and prevent DoS traffics. However, the separation and isolation mechanisms in cloud to protect against inside DoS attacks are still evolving and vendor-dependent (refer to section 4.6, Risk from Multi-tenancy). Finally, cloud computing depends on strong identity and access management (IAM) to keep malicious users at bay (refer to section 4.5, Identity and Access Management (IAM) and Authorization not Deployed).

**Mitigation 1:** Cloud Consumer adopts hybrid approach with Cloud Broker to contract with two or more Cloud Providers (Cloud Broker service aggregation) (NIST, p23).

Cloud Provider adopts hybrid approach by contracting with two or more cloud carriers for access network to cloud consumer(s) (NIST RA, p24)

NIST Cloud Computing Standards Roadmap – 092 5 Jul, 2011 scenario eight discusses cloud consumer access across multiple clouds, simultaneously, as a mitigation strategy

Deployment Case 2: In the distributed deployment cases, a single cloud consumer has an application that may be distributed across two or more cloud providers and administrative domains simultaneously. While the cloud consumer may have simple consumer-provider interactions with their application and the providers, more complicated Peer-to-Peer (“P2P”) interactions may be required -- between both the consumer and provider and also between the providers themselves.

**Sufficiency Comment:** See usage scenario one of the NIST RA (2011). 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 experiences an outage as well. Cloud provider connected to two cloud carriers, which connect to cloud consumer providing redundant paths. An outage experienced by one cloud carrier will not result in total loss of cloud provider access to cloud consumers. This solution will require the development of interoperability standards and contractual relationships between cloud providers. Also more research is required into the effects of these mitigation strategies have on cloud based risk to the cloud consumer.

Action: Create usage scenario four (NIST RA (2011) to depict 4 node topology

Cloud consumer, (2) cloud carriers and cloud provider

Cloud brokers, through service aggregation or arbitrage provide diverse network for cloud

**Mitigation 2:** Cloud Consumer contract Cloud Carrier (or Cloud Broker) for diverse network access from customer site(s) to the by Cloud Carrier network. Cloud Consumer site(s) access diversity can take the form of ingress/egress, route, switch, serving wire center and interconnection points (Cloud Broker service Intermediation) (NIST, p23).

Cloud provider must provide a highly redundant, high availability service, resource abstract and control physical resource layers (NIST , RA, p15). These three layers are all required to provide a cloud service to cloud consumers within the cloud provider boundary. Cloud provider must be required to have COOP/DR plans in place by cloud customer contract / SLA.

**Sufficiency Comment:** See usage scenario one of the NIST RA (2011). Network diversity between the Cloud Consumer site(s) and Cloud Carrier points of presence is a sound strategy to address the risk of denial of service (DoS).

The physical building and the three layers (service, resource and physical) must be redundant, resilient, and fault resistant. NIST 800-146 (p 8- suggests the following mitigation strategies:

SLA Evaluation

Portability of Workloads

Interoperability between Cloud Providers

Disaster Recovery

**Mitigation 3:** Cloud Consumer contract Cloud Carrier, or Cloud Broker, for redundant customer premises equipment (CPE) with failover (FO) capability to provide high availability network access to compliment diverse network access to Cloud Provider network. Cloud provider must provide redundant data instances (protect data availability) through various strategies and control measures.

**Sufficiency Comment:** See usage scenario one of the NIST RA (2011). The Cloud Carrier, through its transport agent for example, shall provide required equipment as part of the cloud based service contract with appropriate SLA's. The equipment SLA's include provisions such as (equipment monitoring, service, upgrades, repair, replacement and technology refresh).

CSA Cloud Control Matrix: Control Area: Operations Management - Equipment Maintenance Policies and procedures shall be established for equipment maintenance ensuring continuity and availability of operations.

MA-2 Controlled Maintenance

MA-3 Maintenance Tools

MA-4 Non-Local Maintenance

MA-5 Maintenance Personnel

MA-6 Timely Maintenance

\*\* NIST SP 800-53 Controls

**References:**

CSA Cloud Control Matrix

Federal Risk and Authorization Management Program (FedRAMP)

NIST Cloud Computing Reference Architecture (30 March 2011), Version 1

NIST SP 800-53

**5.8 Incident Response**

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

**Importance:** This requirement highlights the need for updated 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 lifecycle consist of Preparation, Detection and Analysis, Containment, Eradication, and Recovery, and Post-Incident Activity<sup>11</sup>.



<sup>11</sup> NIST SP 800-61, Computer Security Incident Handling Guide

**Figure 2: NIST Incident Response Life Cycle**

A cloud computing-oriented architecture introduces a number of complexities when conducting the incident response life cycle, though the steps of the life cycle remain the same. When federal departments and agencies consider embracing cloud computing technologies, they should update their incident response plan to include cloud computing considerations, as incidents will occur once agencies begin to process operational information in a cloud environment. This is not a trivial task, as some Computer Security Incident Response Team (CSIRT) organizations have already attempted to address the issue of incident response in a distributed, dynamically scalable cloud environment, and discovered that nuances exist in the incident response process. For example, the compromise of a credential in a cloud computing environment is not as simple as in a system environment, where the account credentials may be disabled. The distributed nature of the cloud environment and the data abstraction requires additional knowledge of the underlying resources, possible avenues for credential compromise, and strong logging to determine abuse.<sup>12</sup>

Some of the key challenges for federal departments and agencies when developing a cloud-oriented incident response plan include data separation, reporting and compliance, data safeguarding and minimization, forensic acquisition including live incident response, and forensic investigation. All these key challenges can fit within the NIST incident response life cycle model, but may require additional tools, techniques, and procedures. For example, data separation in a public cloud may be particularly difficult, as data may reside across multiple shared resources. If data separation is a requirement, consider a private cloud or select a public cloud with strong assurances to maintain data separation. Weak data separation may also make acquisition problematic, as acquisition techniques will need to include data minimization processes to remove data that is out of scope, though this also runs the risk of accidentally excluding relevant evidence data.

The right to capture data is also problematic when considering government private and public clouds. Consider that under international law, information system data is governed under the laws and regulations of the country where the data resides. If government data resides within a foreign country's servers as part of a public cloud, acquisition and overall control of the data may have additional challenges. Consider that if data should have strong acquisition controls and access restrictions that it remains located within areas under United States control.

These concerns may be weakly addressed using a Service Level Agreement with public clouds, but for strongest control, system owners should consider a government private cloud dedicated to government customers.

**Preparation:**


---

<sup>12</sup> Setting up a Grid-CERT: experiences of an academic CSIRT. (2007) Moller, Klaus.

An incident response plan should take the potential for security breaches into account. Several NIST publications provide essential guidance for developing security-related aspects of an incident response plan. *Minimum Security Requirements for Federal Information and Information Systems* (FIPS 200) and *Standards for Security Categorization of Federal Information and Information Systems* (FIPS 199), provide a framework for categorizing information and information systems, and provide baseline security requirements and security controls for incident handling and reporting. The procedures organizations should use to implement FISMA requirements are found in two primary documents: NIST Special Publication 800-61, *Computer Security Incident Handling Guide*;<sup>13</sup> and the Concept of Operations for the United States Computer Emergency Readiness Team (US-CERT), the federal security incident-handling center located within the Department of Homeland Security.<sup>14</sup>

Perhaps a larger difficulty will be the differentiation between public and private clouds, where ownership of the equipment and physical access rights for forensic imaging may not be as clear. These items should be clearly identified in any Service Level Agreements for cloud computing, and recognize that the same level of access for private clouds may not be afforded to public cloud customers. Service Level Agreements should clearly articulate government-cloud-specific requirements relating to incident response, such as the right to capture and the responsibilities of the General Support System (GSS), in this case the public cloud provider, to support incident response requirements.

SLA requirements and preparation should include the identification of proper incident response analysts monitoring your cloud environment, and pre-establishment of relationships with security teams monitoring your shared cloud infrastructure. It is important to note that while a centralized incident response capability may exist to coordinate the response activity, such as a service provided by the cloud provider, GSA, or US-CERT, the responsibility for data security always remains with the Authorizing Official (AO). This may be easier with a private cloud infrastructure, where government incident response teams may work together. The use of public cloud infrastructure will require collaboration with public incident response teams that may not be dedicated for government support.

In addition, the data abstraction of a cloud computing environment may be problematic for traditional log collection and analysis, which should be considered during the preparation stage. Security Incident Event Manager (SIEM) tools may require additional capacity to handle the increased data flow from a consolidated cloud environment, such as additional storage capacity, bandwidth capacity, and analysts to review and process SEIM alerts. Abstracted APIs and

---

<sup>13</sup> See *Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology* (<http://csrc.nist.gov/publications/nistpubs/800-61/sp800-61.pdf>).

<sup>14</sup> The responsibilities of US-CERT are outlined in 44 U.S.C. § 3546. Its complete set of operating procedures may be found on the US-CERT website at [www.us-cert.gov/federal/reportingRequirements.html](http://www.us-cert.gov/federal/reportingRequirements.html).

shared infrastructure may also mean log files are comingled, or worse may not be logged to your agencies requirements if at all. A thorough review of the logging environment at multiple levels of the cloud architecture should be conducted to assess the thoroughness of logging being conducted.

### **Detection and Analysis:**

Federal departments and agencies have a wide array of incident reporting and compliance requirements that may be strenuous for public cloud providers to fully meet. One challenging problem in the government's use of cloud computing is that of data spills, specifically those involving PII and/or classified data.

OMB has stated in its implementing guidance that “[s]afeguarding [PII] in the possession of the government and preventing its breach are essential to ensure the government retains the trust of the American public.”<sup>15</sup> Maintaining the public's trust greatly depends on an organization's procedures for detecting, reporting, and responding to privacy incidents involving the suspected or confirmed breach of PII and any data collected by the agency.

OMB Memorandum 07-16 requires organizations to develop and maintain an incident response policy and notification plan. Even with the implementation and monitoring of privacy and security controls, however, it is impossible to prevent all risks associated with government operations. It is inevitable that Federal organizations will experience incidents, and agencies should take this into consideration when evaluating public or private clouds that may safeguard PII. Being prepared to respond to and mitigate these risks before substantial damage is done is critical to the success of incident response.

In addition, Cloud Service Providers that process sensitive government information will eventually experience a classified spill, usually through human error where organizational employees with access to classified material inadvertently store, process, or transmit classified data in an unauthorized cloud environment. Few technical mechanisms exist to prevent a classified spill, rather users should be trained to prevent spillage, and mechanisms should be in place to detect, contain, and eradicate the classified data spill to an acceptable level. Classified spills are required to be reported to US-CERT within one hour.

The causes for classified spills are wide-ranging, from user error and improper handling to aggregation of sensitive data.<sup>16</sup> Responding to a classified data spill is different than responding to other types of incidents, as it requires specialized reporting, handling, and response procedures. From the moment of detection, the incident rapidly moves to containment. Specialized policy is identified in the National Policy on Classified Information Spillage.<sup>17</sup> This includes reporting to appropriate authorities and treating the affected media with the

---

<sup>15</sup> OMB M-07-16, *Safeguarding Against and Responding to the Breach of Personally Identifiable Information*.

<sup>16</sup> “FAQ on Incidents and Spill” (August 2007) CNSS-079-07.

<sup>17</sup> CNSS Policy No. 18. June 2006

appropriate level of security commensurate with the classification of the data spilled. For example, this means that incident handlers should have the appropriate clearance to the classification level of the spilled data, and response procedures should follow National Security Systems (NSS) policy, regardless of the offending agency policy.

Implications for handling classified data spills in a cloud environment are dramatic. Requirements for isolation and containment of spilled data may have implications in a shared cloud environment. If data is spilled into a public cloud, response requirements may include seizure and isolation of certain aspects of the public cloud until such time the classified spill can be properly remediated. This is a scenario public cloud providers or users should be willing to respond to, and in a purely shared environment this may be a complex task. If the data is spilled into a private cloud at a government facility, more controls may be in place to mitigate the impact of classified spill containment. As stated in CNSS Instruction No. 1001:

“Unless otherwise determined by the information owner, in cases where the spillage occurred within agency-controlled space, sanitization is not required until such time as the affected systems are removed from agency control. In such cases, immediate actions will be required to ensure that the spillage is isolated and contained, and that unauthorized access is precluded based on risk management decisions and operational considerations related to the loss of information services. Preclusion of unauthorized access may include software overwriting of affected data sectors in the interest of meeting operational needs. When the media is released from agency control, sanitization is required.”<sup>18</sup>

In general, if the spill is only one level down in classification (ie. Secret to unclassified), remediation of the spill may include secure wiping of the storage media before reuse. However, if the spill is across two levels (ie. Top Secret or above to unclassified) then physical destruction is usually required. Remediation activities require containment and safeguarding at the appropriate level within the agency-controlled space, and may impact the resources available to the shared cloud customers. Spill into a public cloud is by default a non-agency controlled space and may require immediate sanitization with an impact on shared cloud customers, which should be built-into and considered in any shared public cloud environments.

Another key aspect that addresses government specific guidelines includes the recent mandates for continuous monitoring. The cloud environment lends itself to the continuous monitoring paradigm, as the levels of data abstraction requires API hooks and interfaces to represent the current state of the cloud health and security posture. A properly architected cloud environment, such as one that provides strong audit controls, will provide for a continuous monitoring paradigm. This is the appropriate location to implement continuous monitoring solutions, those that are customized specific to the cloud environment and make use of cloud-

---

<sup>18</sup> National Instruction On Classified Information Spillage. (February 2008) CNSS Instruction No. 1001

specific capabilities, like virtualization, hypervisor monitoring, and data-centric health, monitoring, and management tools.

### **Containment, Eradication, and Recovery:**

Probably the most complex step in the NIST incident response life cycle for cloud computing is the third phase, as containment and eradication are particularly complex in this paradigm. We already addressed some of the complexities of containment and eradication within a classified spill incident, which operates under a different incident paradigm, but there are additional nuances to this stage of the incident response lifecycle for cloud computing. The physical location of data as it resides within a cloud makes it difficult to locate and contain. Data should be identified as required to be stored CONUS or OCONUS, and should be locatable and accessible to federal incident response officials if required by the agency.

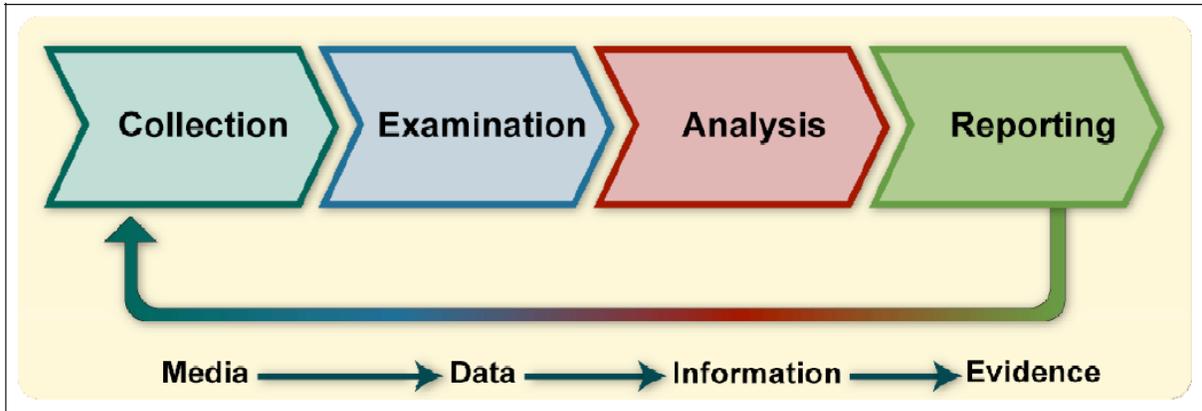
A cloud environment inherently relies on a shared infrastructure, so data separation is an inherent consideration. However, this level of separation tends to be for data management and network efficiency, but may not provide the level of separation necessary for federal customers. Separation requirements and concerns over data comingling should be clearly articulated within the cloud SLA, from both the customer and cloud provider perspective. If a customer has a spill which requires IT asset seizure and forfeiture, how can shared cloud customers ensure they do not lose data during the seizure. In addition, seized media may include non-agency information, which could possibly be sensitive, and should be properly accounted for and controlled. Consider that your agency may have media seized as a result of a co-located customers incident, will you have concerns over your data being captured in the investigation sweep. While this is a common concern for both public and private government clouds, a private government cloud will be shared amongst other federal agencies that have similar data protection requirements. A public cloud scenario may share resources with commercial industry, private citizens, other countries, and other customers that are not constrained to the same federal security requirements.

### ***Technical Considerations***

The technical considerations of responding to an incident in a cloud environment are significantly different than traditional incident response and forensic techniques, though the elemental guidelines still apply. NIST identifies the forensic process in SP 800-86, which still applies in a cloud environment<sup>19</sup>.

---

<sup>19</sup> Guide to Integrating Forensic Techniques into Incident Response. (August 2006). NIST SP 800-86



**Figure 3: The Forensic Process**

The forensic process is unique in a cloud environment because data is abstracted from the computing capability, and the computing environment is abstracted from the customer interface. However, to maintain a strong chain of custody and conduct forensically sound investigations that will hold up in court, the investigation and acquisition process for cloud computing facilities should maintain strong evidentiary controls. Unfortunately, little guidance exists on how to acquire and conduct forensics in a cloud environment, though many of the techniques identified in NIST 800-86 and other best practice documents still apply. As already stated, the issue of physical hard drive seizure may be problematic or impossible in a cloud environment, requiring investigators to utilize other means of collecting evidence.

Evidence collection in a cloud environment may rely more heavily on live response techniques to capture volatile information, such as memory collection and analysis techniques, especially where an endpoint or physical machine is a component of the intrusion. Thorough logging may become much more critical in a cloud environment, as well as network forensic techniques that help analysts rebuild a timeline of the incident and gather volatile evidence, even if the physical hard drive is unavailable. Of course, all these capabilities require that appropriate logging and capture is in place before the incident occurs. These requirements may be easier to implement in a private cloud environment where there is less concern over logging or capturing data of other cloud customers.

Once evidence is collected, the second phase of the forensic process is examination, which may be problematic as evidence will likely exist in cloud-based file systems and formats. For example, if evidence is captured through hard drive seizure of a Hadoop cloud environment, the evidence should be extracted from the Hadoop file system and unnecessary data eliminated. Current forensic tools may not support cloud file system formats, or will require shims and translators to read evidence data and capture to a forensically sound image.

The third phase is analysis, after evidence has been collected and translated into a format readily accessible for a forensic analyst to begin their work. Unfortunately, the distributed

nature of a cloud environment may result in a much larger data set of files, file systems, log files, memory dumps, images, and other data extracts for investigators to examine, relative to the few hard drives currently provided in a physical seizure and investigation. Additional data minimization may also be required as the acquisition process may capture data not owned by the agency, or in a public cloud scenario, data owned by the federal government. Industry has recognized some of this need and begun to evaluate more powerful forensic tools, but their support for cloud system investigations is still in the early stages.<sup>20</sup>

### **Post-Incident Activity:**

The fourth phase of the forensic process, and the fourth phase of the incident response process, reporting, may not have significant changes under a cloud computing model. However, as this is still an evolving field, strong reporting and table top exercises and drills as required under NIST guidance for incident response in a cloud computing environment will help organizations identify weaknesses in their training, data acquisition, logging, examination, and analysis. Organizations should use these tests to identify weaknesses in their cloud incident response plans and work to remediate those weaknesses. They may discover that certain cloud incident response activities are better handled by a central organization, such as large data log analysis and acquisition, while other activities should be handled internally, such as internal investigations, counter intelligence issues, and those of a sensitive or restricted nature.

The Privacy Act requires organizations to make public information regarding procedures for an individual to access his or her information and to correct or amend inaccurate information. Organizations should also have in place policies and procedures for managing privacy complaints or inquiries when information is used in a cloud environment. Such procedures should ensure that all complaints are recorded, tracked, and addressed.<sup>21</sup> Where feasible, organizations should establish an automated tracking process to capture and manage privacy complaints, to promote compliance with written policies and procedure, and to ensure all complaints are addressed.

Organizations should track security incidents and privacy complaints related to cloud computing for purposes of internal and external reporting. This information should be used to identify areas within an organization that may require further review or education and training.

## **6 References:**

<sup>20</sup> Ayers, Daniel. (2009) A second generation computer forensic analysis system. Digital Investigation.

<sup>21</sup> OMB M-08-21 *FY 2008 Reporting Instructions for the Federal Information Security Management Act and Agency Privacy Management* requires organizations to include in their FISMA report the number of written privacy complaints and divides them into the following categories: Process and Procedural, Redress, Operational, and Referrals (July 14, 2008). Accurate categorization and reporting of complaints can also aid in directing education and training resources to mitigate areas of greatest concern.

NIST Cloud Computing Standards Roadmap – 092 5 Jul, 2011

NIST Cloud Computing Reference Architecture v1. Dated 30 March, 2011

CSA Cloud Control Matrix

Draft Cloud Computing Synopsis and Recommendations -

<http://csrc.nist.gov/publications/drafts/800-146/Draft-NIST-SP800-146.pdf>

[01] Draft Cloud Computing Synopsis and Recommendations -

<http://csrc.nist.gov/publications/drafts/800-146/Draft-NIST-SP800-146.pdf>

[02] Proposed Security Assessment & Authorization for U.S. Government Cloud Computing -

<http://www.cio.gov/pages.cfm/page/Federal-Risk-and-Authorization-Management-Program-FedRAMP>

[03] Security Guidance for Critical Areas of Focus in Cloud Computing V2.1 -

<https://cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf>

[04] Top Threats to Cloud Computing V1.0 -

<https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>

[05] SaaS, PaaS, and IaaS: a Security Checklist for Cloud Models -

<http://www.csoonline.com/article/print/660065>

[06] Cloud – 10 Risks with Cloud IT Foundation Tier -

[https://www.owasp.org/index.php/Cloud-10\\_Risks\\_with\\_Cloud\\_IT\\_Foundation\\_Tier](https://www.owasp.org/index.php/Cloud-10_Risks_with_Cloud_IT_Foundation_Tier)

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

[08] 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](http://www.trustedcomputinggroup.org/files/resource_files/1F4DEE3D-1A4B-B294-D0AD0742BA449E07/Cloud%20Computing%20and%20Security%20Whitepaper_July29.2010.pdf).

[09] Symantec Internet Security Threat Report, Trends for 2010, Volume 16, April 2011.