Cloud Computing Standards Roadmap

Document: **NIST CCSRWG – 060**

Ninth Working Draft

March 24, 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.

TABLE OF CONTENTS

Table of Figures

Table of Tables

NIST Cloud Computing Standards Roadmap

Executive Summary

TBD

1 Introduction

1.1 Background

U.S. laws and associated policy require Federal agencies to use international, voluntary consensus standards in their procurement and regulatory activities, except where inconsistent with law or otherwise impractical.[i]

The National Institute of Standards and Technology (NIST) has been designated by Federal Chief Information Officer Vivek Kundra to accelerate the federal government's secure adoption of cloud computing by leading efforts to develop standards and guidelines in close consultation and collaboration with standards bodies, the private sector, and other stakeholders.

The NIST Cloud Computing Program was formally launched in November 2010 and was created to support the federal government effort to incorporate cloud computing as a replacement for, or enhancement to, traditional information system and application models where appropriate. The NIST Cloud Computing Program operates in coordination with other federal-wide cloud computing implementation efforts (CIO Council/ISIMC, etc.) and is integrated with federal CIO Vivek Kundra's 25-point IT Implementation Plan for the federal government. NIST has created the following Working Groups in order to provide a technically oriented strategy and standards-based guidance for the federal cloud computing implementation effort:

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

1.2 NIST Cloud Computing Vision

NIST's long term goal is to provide thought leadership and guidance around the cloud computing paradigm to catalyze its use within industry and government. NIST aims to shorten the adoption cycle, which will enable near-term cost savings and increased ability to quickly create and deploy safe and secure enterprise solutions. NIST aims to foster cloud computing practices that support interoperability, portability, and security requirements that are appropriate and achievable for important usage scenarios.

The NIST area of focus is technology, and specifically, interoperability, portability and security requirements and standards and guidance. The intent is to use the standards strategy to prioritize NIST tactical projects which support US government agencies in the secure and effective adoption of the cloud computing model to support their missions. The expectation is that the set of priorities will be useful more broadly by industry, Standards Development Organizations, cloud adopters, and policy makers.

1.3 NIST Cloud Computing Standards Roadmap Working Group

Standards Developing Organizations (SDOs) and others have and are developing supporting cloud computing documents to include standards, conceptual models, reference architectures and standards roadmaps to facilitate communication, data exchange, and security for cloud computing and its application. Still other standards are emerging to focus on technologies that support cloud computing, such as virtualization. The NIST Cloud Computing Standards Roadmap Working Group will leverage this existing, publicly available work, plus the work of the other NIST Working Groups, to develop a NIST Cloud Computing Standards Roadmap that can be incorporated into the NIST USG Cloud Computing Roadmap.

1.4 How This Report Was Produced

The NIST Cloud Computing Standards Roadmap Working Group has surveyed the existing standards landscape for security, portability, and interoperability standards/models/studies/use cases, etc. relevant to cloud computing. Using this available information, standards, standards gaps or overlaps, and standardization priorities have been identified.

Future editions of this report may consider additional areas, such as maintainability, usability, reliability and resiliency.

2. The NIST Definition of Cloud Computing

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics**,** three service models, and four deployment models.

Essential Characteristics:

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

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

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location

independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

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

Measured Service. Cloud systems automatically control and optimize resource use by leveraging a metering capability[1] at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Service Models:

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

Cloud Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

[1] Typically through a pay-per-use business model.

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

Deployment Models:

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

Community cloud. The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise.

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

Hybrid cloud. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

3. Cloud Computing Reference Architecture

The information in this clause is based upon the outcomes of the NIST Cloud Computing Reference Architecture Working Group.

3.1 Overview

The NIST cloud computing reference architecture is an extension of the NIST definition of cloud computing. It represents the three service models (*Software as a Service (SaaS)/Platform*

as a service (PaaS)/Infrastructure as a Service (IaaS)), four deployment models (*private cloud/ community cloud/public cloud/hybrid cloud*), and five essential characteristics (*on-demand self-service/broad network access/resource pooling/rapid elasticity/measured service*). It is not tied to specific vendor products, services or reference implementation, and does not define prescriptive solution that inhibits innovation.

The NIST cloud computing reference architecture presented in this clause is a generic high level architecture for discussing the requirement, structure and operation of cloud computing. It defines a set of actors, activities and functions that can be used in the process of developing cloud computing architectures, and relates to a companion cloud computing taxonomy. It contains a set of views (diagrams) and descriptions that are the basis for discussing the characteristics, uses and standards for cloud computing.

The NIST cloud computing reference architecture consists of five major actors (see Table 1). Each actor plays a role and performs a set of activities and functions.

Actor	Definition
Cloud Consumer	Person or organization that maintains a business relationship with, and uses service from, *Cloud Providers*.
Cloud Provider	Person, organization or entity responsible for making a service available to *Cloud Consumers.*
Cloud Auditor	A party that can conduct independent assessment of cloud services, information system operations, performance and security of the cloud implementation.
Cloud Broker	An entity manages the use, performance and delivery of cloud services, and negotiates relationships between *Cloud Providers* and *Cloud Consumers.*
Cloud Carrier	The intermediary that provides connectivity and transport of cloud services from *Cloud Providers* to *Cloud Consumers*.

Table 1 - Actors in Cloud Computing

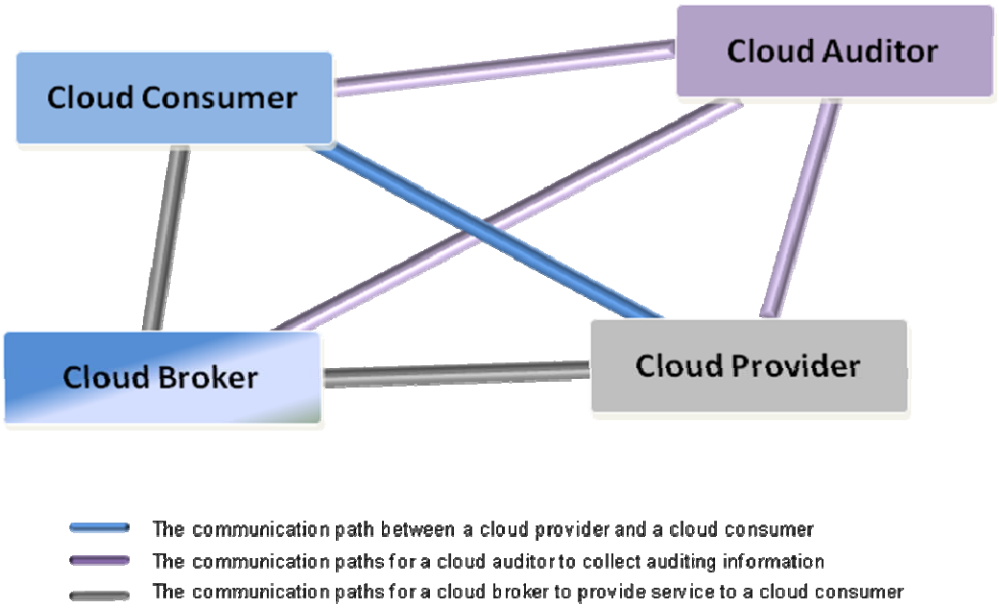

Figure 1 - Interactions between the Actors in Cloud Computing

3.2 Cloud Consumer

A cloud consumer represents a person or organization that maintains a business relationship with, and uses the service from a cloud provider, whether it is software, platform, or infrastructure as a service. A cloud consumer browses the service catalog from a cloud provider, requests the appropriate service, sets up SLAs and contracts with the cloud provider, and uses and manages the service. The cloud consumer may be billed for the service provisioned, and need to arrange payments accordingly.

Depending on the services requested, the activities and usage scenarios can be different among cloud consumers, as shown in Table 2. Some example usage scenarios are listed in Figure 2.

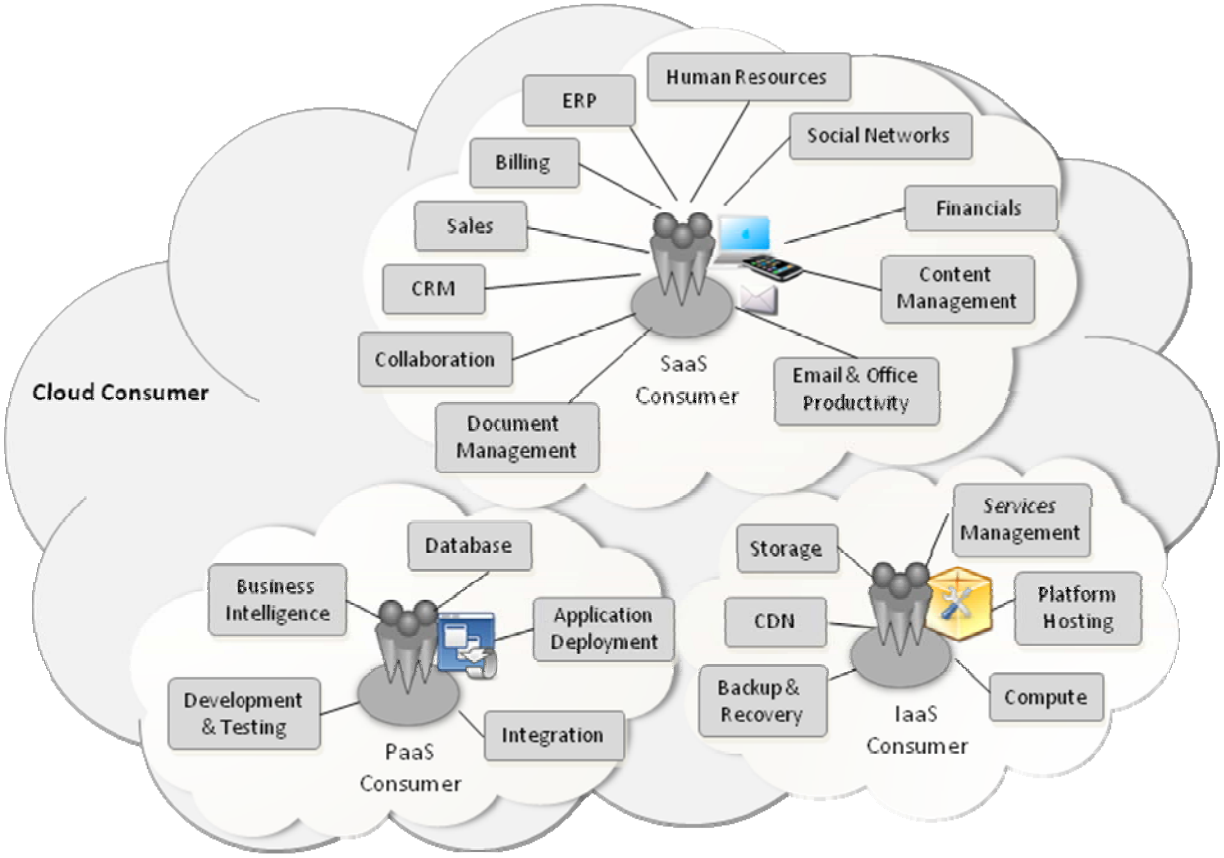

Figure 2 - Example of Services Available to a Cloud Consumer

Consumer Type	Major Activities	Example Users
SaaS	Uses application/service for business process operations	Business users
PaaS	Develops, tests, deploys and manages applications hosted in a cloud environment	Application developers, testers and administrators
IaaS	Creates/installs, manages and monitors services for IT infrastructure operations	System developers, administrators, IT managers
Provider Type	**Major Activities**	
SaaS	Installs, manages, maintains and supports the software application on a cloud infrastructure	
PaaS	Provisions and manages cloud infrastructure and middleware for the platform consumers; provides development, deployment and administration tools to platform consumers.	
IaaS	Provisions and manages the physical processing, storage, networking and the hosting environment and cloud infrastructure for IaaS consumers.	

Table 2 - Cloud Consumer and Cloud Provider

3.3 Cloud Provider

Cloud providers perform services to support the business processes of cloud consumers at agreed service levels and costs. The providers perform different tasks for different service types (see Table 3). The activities of cloud providers can be discussed in greater detail from the perspectives of *Service Deployment*, *Service Orchestration*, *Cloud Service Management*, *Security* and *Privacy*.

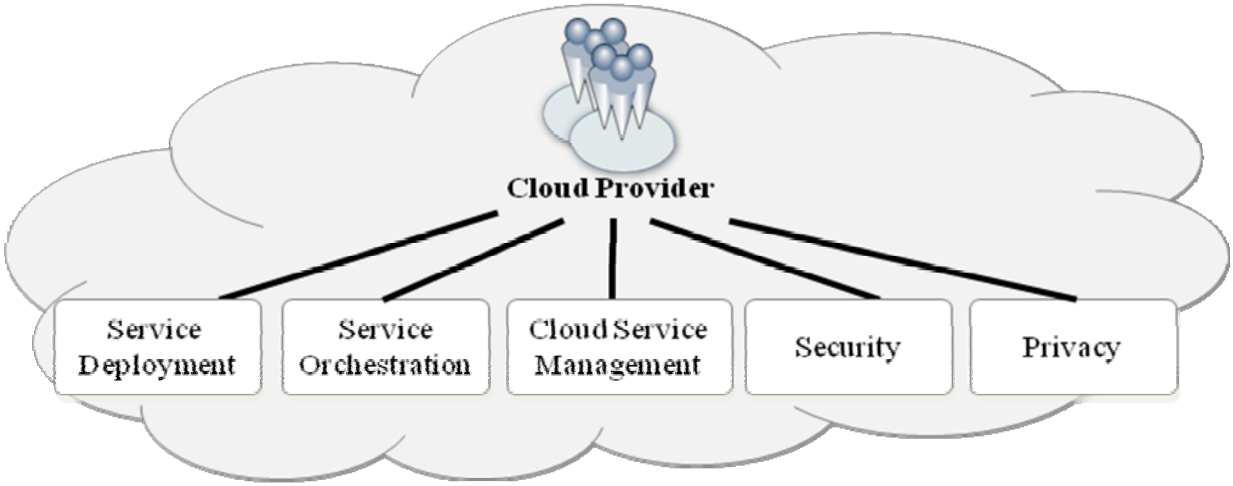

Figure 3 - Cloud Provider

- ***Service Deployment:*** A cloud system can be operated in one of the four deployment models (public cloud/private cloud/community cloud/hybrid cloud).
- ***Service Orchestration*** refers to the arrangement, coordination and management of cloud infrastructure to provide different cloud services to meet IT and business requirements. See Figure 4.
- ***Cloud Service Management*** includes all the service-related functions that are necessary for the management and operations of those services required by or proposed to cloud consumers. See Figure 5.
- ***Security*** protects information and information systems on the cloud from unauthorized access, use, disclosure, disruption, modification, or destruction.
- ***Privacy*** protects the assured, proper, and consistent collection, processing, communication, use and disposition of personal information (PI) and personally identifiable information (PII) on the cloud.

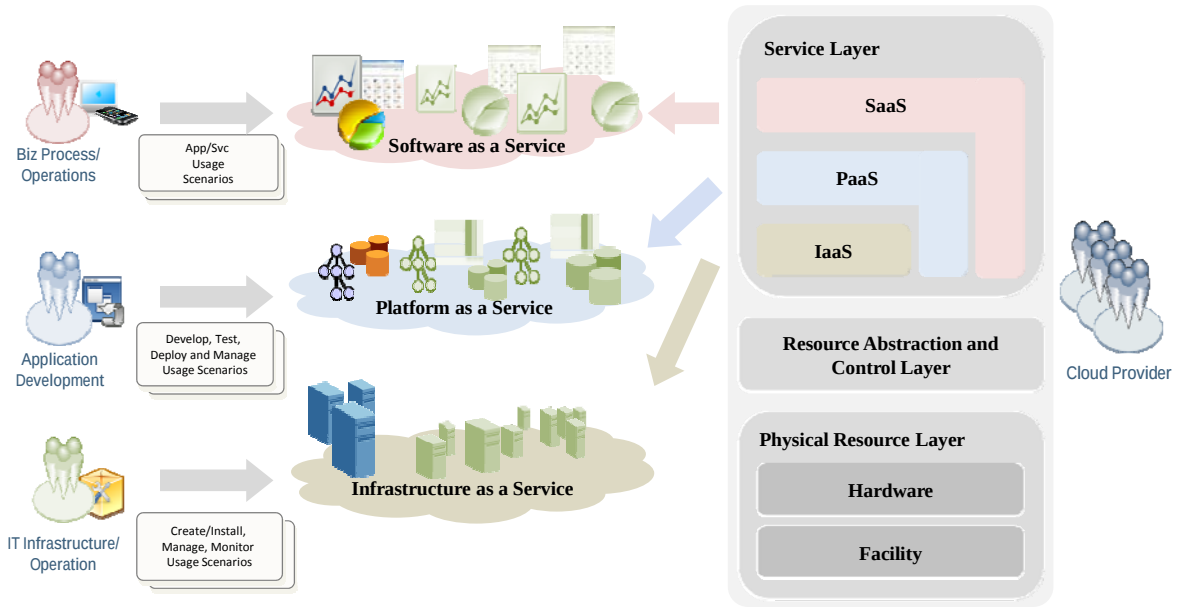

Figure 4 - Cloud Provider: Service Orchestration

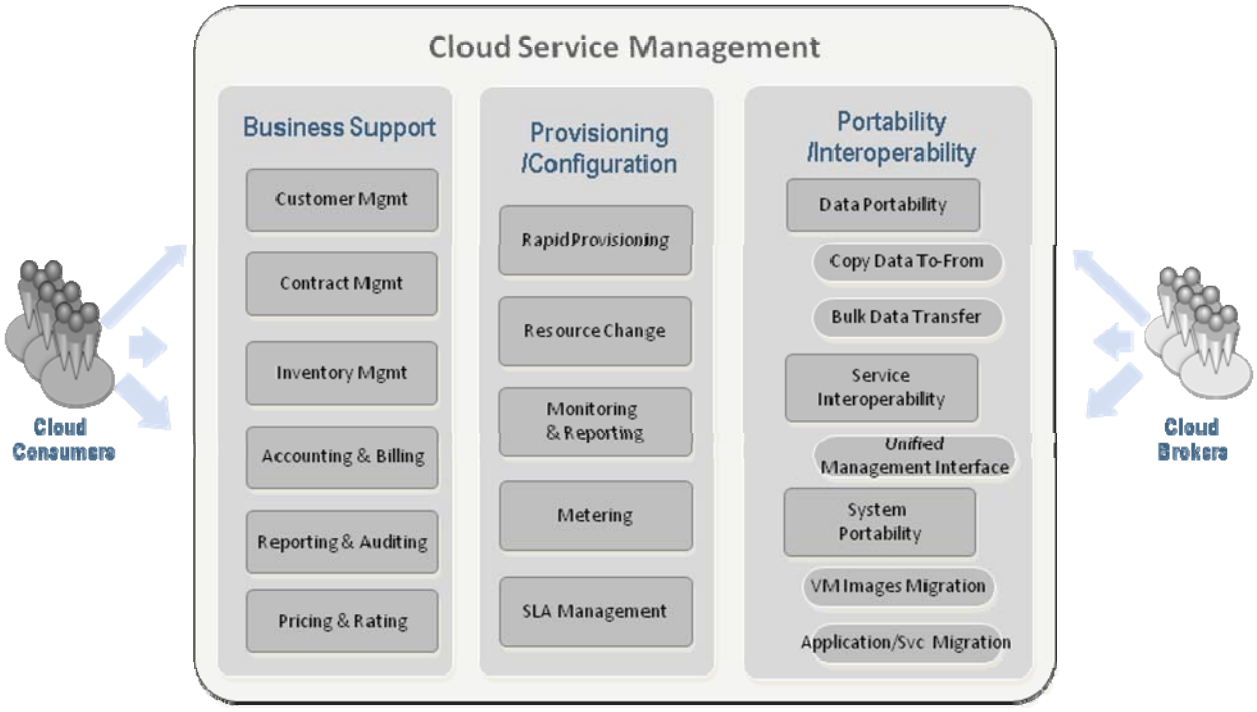

Figure 5 - Cloud Provider: Cloud Service Management

3.4 Cloud Carrier

A cloud carrier acts as an intermediary that provides connectivity and transport of cloud services between cloud consumers and cloud providers. Cloud Carriers provide access to consumers through network, telecommunication and other access devices. For example, cloud consumers can obtain cloud services through network access devices such as computers, laptops, mobile phones, mobile internet devices (MIDs), etc. The distribution of cloud services is normally provided by network and telecomm carriers, or a transport agent, where a *transport agent* refers to a business organization that provides physical transport of storage media such as high-capacity hard drives. Note that a cloud provider shall set up SLAs with a cloud carrier to provide services in consistent level of SLAs. In general, the cloud carrier may be required to provide dedicated and encrypted connections.

3.5 Cloud Broker

A cloud broker is an entity that manages the use, performance and delivery of cloud services, and negotiates relationships between *Cloud Providers* and *Cloud Consumers*.

A cloud broker may provide the following services:

- ***Service Intermediation***: An intermediation broker provides a service that directly enhances a given service delivered to one or more service consumers, essentially adding value on top of a given service to enhance some specific capability.
- ***Service Aggregation***: An aggregation brokerage service combines multiple services into one or more new services. It will ensure that data is modeled across all component services and integrated as well as ensuring the movement and security of data between the service consumer and multiple providers.
- ***Service Arbitrage***: Cloud service arbitrage is similar to cloud service aggregation. The difference between them is that the services being aggregated aren't fixed. Indeed the goal of arbitrage is to provide flexibility and opportunistic choices for the service aggregator, e.g., providing multiple e-mail services through one service provider or

providing a credit-scoring service that checks multiple scoring agencies and selects the best score.

3.6 Cloud Auditor

A cloud auditor is a party that can conduct independent assessment of cloud services, information system operations, performance and security of the cloud implementation. A cloud auditor can evaluate the services provided by a cloud provider in terms of security controls, privacy impact, performance, etc. For security auditing, a cloud auditor can make an assessment of the security controls in the information system to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system. Note that auditing is especially important for federal agencies and "agencies should include a contractual clause enabling third parties to assess security controls of cloud providers" (*by Vivek Kundra, Federal Cloud Computing Strategy, February 2011*.)

4. Cloud Computing Use Cases

Cloud computing use cases describe the consumer requirements in using cloud computing service offerings. Analyzing business and technical cloud computing use cases and the applicable standards provides an intuitive, utility-centric perspective in surveying existing standardization efforts and identifying gaps. This section leverages the business and technical use case outputs from other NIST Cloud Computing Program working groups and presents an analysis on how existing cloud-related standards fit the needs of USG cloud consumers and where the gaps for standardizations are.

4.1. Business Use Cases

The Business Use Case Working Group has produced a template for documenting specific use cases. It includes a Concept of Operations where Current System and Desired Cloud

Implementation are described. The template also requires information about "how the current system integrates with other systems, what are security requirements, do network considerations vary among users (local versus remote, for example), etc" to aid in migration. A set of business use case describing candidate USG agency cloud deployment examples are being drafted. The real stories captured in these business use cases not only help us understanding the background and business drivers behind the adoptions of cloud computing in USG agencies, they also help surface general USG agency consumer concerns and realistic issues encountered in security, interoperability and portability. These business use cases can help us summarize the key technical requirements that need to be addressed using cloud-related standards in these areas.

The "Cloud First" Business Use Case called out by the Federal CIO is a more general expansion of this analysis to multiple interacting Current Systems and Cloud Implementations. This expansion is to support evolving business processes as Cloud deployments are implemented. It requires interoperability and portability across multiple Cloud deployments and enterprise systems.

4.2. Technical Use Cases

The SAJACC Working Group has produced a set of preliminary use cases developed for the SAJACC project for the first pass through the SAJACC process. Through a series of open workshops, and through public comment and feedback, NIST will continue to refine these use cases and add new use cases as appropriate. These use cases are technical in nature, capturing the more generic and cross-cutting technical requirements of cloud consumers. They are descriptions of how groups of users and their resources may interact with one or more cloud computing systems to achieve specific goals, such as "how to copy data objects into a cloud".

There is a natural mapping from the high level business use cases to the SAJACC technical use cases, where the business operational stories of specific agency consumers will imply specific technical requirements expressed in SAJACC technical use cases. For example, the business

use case of an agency consumer's move of its virtualized computing infrastructure to an IaaS cloud vendor implies the technical requirement of "*VM control: manage virtual machine instance state*" to be met. The rest of this section drives through the high level business use cases to the general technical requirements expressed and analyzes where cloud standards help address these requirements.

4.3. Deployment Scenario Perspective

The "Cloud First" Business Use Case requires more complex interactions between USG agency cloud consumer and cloud providers. There are three main groups of interaction scenarios:

Single Cloud

Scenario 1. Deployment on a Single Cloud
Scenario 2. Manage resources on a Single Cloud
Scenario 3. Interface Enterprise Systems to a Single Cloud
Scenario 4. Enterprise Systems migrated or replaced on a Single Cloud

Multiple Clouds – (serially, one at a time)

Scenario 5. Migration between Clouds
Scenario 6. Interface across Multiple Clouds
Scenario 7. Work with a Selected Cloud

Multiple Clouds – (simultaneously, more than one at a time)

Scenario 8. Operate across Multiple Clouds

The figure below illustrates the different generic scenarios

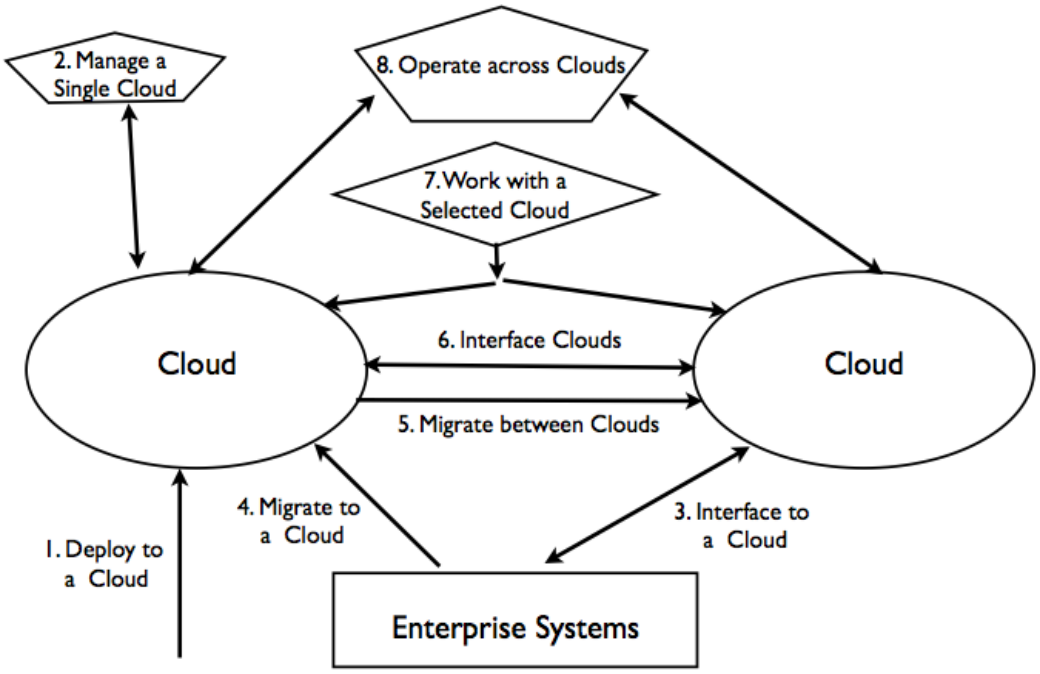

Figure 6 - High Level Generic Scenarios

These technical use cases must also be analyzed in the context of their deployment models and the resultant way cloud actors must interact. These considerations identify two fundamental dimensions to the spectrum of cloud computing use cases:

- Centralized vs. Distributed, and
- Within vs. Crossing Trust Boundaries

These deployment cases will drive the requirements for cloud standards. They can be identified through the following matrix:

	a.) Within Trust Boundary	b.) Crossing Trust Boundary
1.) Centralized i.e., one administrative cloud domain	Deployment Case 1A	Deployment Case 1B
2.) Distributed, i.e., crossing administrative cloud domains	Deployment Case 2A	Deployment Case 2B

Table 3 - Deployment Cases for High Level Scenarios

Deployment Case 1. In the Centralized Deployment cases, there is one Cloud Provider under consideration at a time. Each Cloud Provider may service multiple Cloud Consumers. Each Cloud Consumer has a simple client-provider interaction with the Provider.

Deployment Case 1A. This deployment case is typically a *private cloud* within a single administrative domain and trust boundary wherein policy and governance can be enforced by non-technical means. Use Cases within this Deployment Case may require standards to support the following basic technical requirements:

- Simple, consumer-provider authentication
- VM management
- Storage management
- Service level agreements (SLAs) and performance/energy monitoring
- Service discovery
- Workflow management
- Auditing
- Virtual Organizations in support of Community Cloud Use Cases

Deployment Case 1B. This deployment case is typically (commercial) *public cloud* within a single administrative domain but is outside of any trust boundary that a client could use to enforce policy and governance. Clients must rely on the Cloud Provider to enforce policy and governance through technical means that are "baked into" the infrastructure. Use Cases within this Deployment Case may require standards to support the following additional technical requirements:

- SLAs in support of governance requirements, e.g., national or regional regulatory compliance
- Stronger authentication mechanisms, e.g., PKI Certificates, etc.
- Certification of VM isolation through hardware and hypervisor support
- Certification of storage isolation through hardware support
- Data encryption

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.

Deployment Case 2A. This deployment case is typically a federated cloud of two or more administrative cloud domains, but where the Cloud Providers can agree "out of band" how to mutually enforce policy and governance -- essentially establishing a common trust boundary. Use Cases within this Deployment Case may require standards to support the following basic technical requirements:

- P2P Service discovery
- P2P SLA and performance monitoring

- P2P Workflow management
- P2P Auditing
- P2P Security Mechanisms for Authentication, Authorization
- P2P Virtual Organization Management

Deployment Case 2b. This deployment case is typically a *hybrid cloud* where apps cross a private-public trust boundary, or even span *multiple public clouds*, where both administrative domains and trust boundaries are crossed. Consumers must rely on the Cloud Provider to enforce policy and governance through technical means that are "baked into" the infrastructure. Apps and services may be distributed and need to operate in a P2P manner. Use Cases within this Deployment Case will require all the standards of the other Deployment Cases, in addition to the following more extensive technical requirements:

- P2P SLAs in support of governance requirements

The Use Cases presented in this section will be analyzed with regards to their possible *deployment scenarios* to determine their requirements for standards. This analysis will be subsequently used to evaluate the likelihood of each of these Deployment Cases. Clearly the expected deployment of these Use Cases across the different Deployment Cases will not be uniform. This non-uniformity will assist in producing a *prioritized roadmap* for cloud standards. Likewise, in reviewing existing standards, these Use Cases – in conjunction with their possible Deployment Cases – will be used to identify and prioritize *gaps* in available standards.

Based on this analysis, we note that Scenarios 1 through 4 could, in fact, be deployed on either a private cloud or a public cloud. Hence, the different standards noted in Deployment Cases 1A and 1B will be required. Scenarios 5, 6, and 7 all involve the notion of the serial use of multiple clouds. Presumably these different clouds, used serially, could be either private or public. Hence, Deployment Cases 1A and 1B would also apply, but there are additional requirements to achieve portability, e.g., API commonality. Finally, Scenario 8 could involve a

federated/community cloud or a hybrid cloud. Hence, Deployment Cases 2A and 2B would apply here.

To summarize the detailed technical use cases for this analysis, the following areas of technical requirements are common across all scenarios:

1. Creating, accessing, updating, deleting data objects in Clouds
2. Moving VMs and virtual appliances between Clouds
3. Selecting the best IaaS vendor for private externally hosted Cloud
4. Tools for monitoring and managing multiple Clouds
5. Moving data between Clouds
6. Single sign on access to multiple Clouds
7. Orchestrated processes across Clouds
8. Discovering Cloud resources
9. Evaluating SLAs and penalties
10. Auditing Clouds

5 Cloud Computing Standards

Standards are already available in support of many of the functions and requirements for cloud computing described in Sections 3 and 4. While many of these standards were developed in support pre-cloud computing technologies, such as those designed for Web Services and the Internet, they also support the functions and requirements of cloud computing. Other standards are now being developed in specific support of cloud computing functions and requirements, such as virtualization.

To assess the state of standardization in support of cloud computing, the NIST Cloud Computing Standards Roadmap Working Group has compiled an Inventory of Standards Relevant to Cloud Computing http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/StandardsInventory.

5.1 Information and Communication Technologies (ICT) Standards Life Cycle

The figure below is a high level conceptualization of how ICT standards are developed and standards-based ICT products, processes and services are deployed. This figure is not meant to imply that these processes occur sequentially. Many of the processes illustrated can and should be done somewhat concurrently. Some of these processes (e.g., reference implementations / product / process / service / test tools development; testing; deployment) occur outside of the SDO process. These processes provide input and feedback to improve the standards, profiles, test tools, etc.

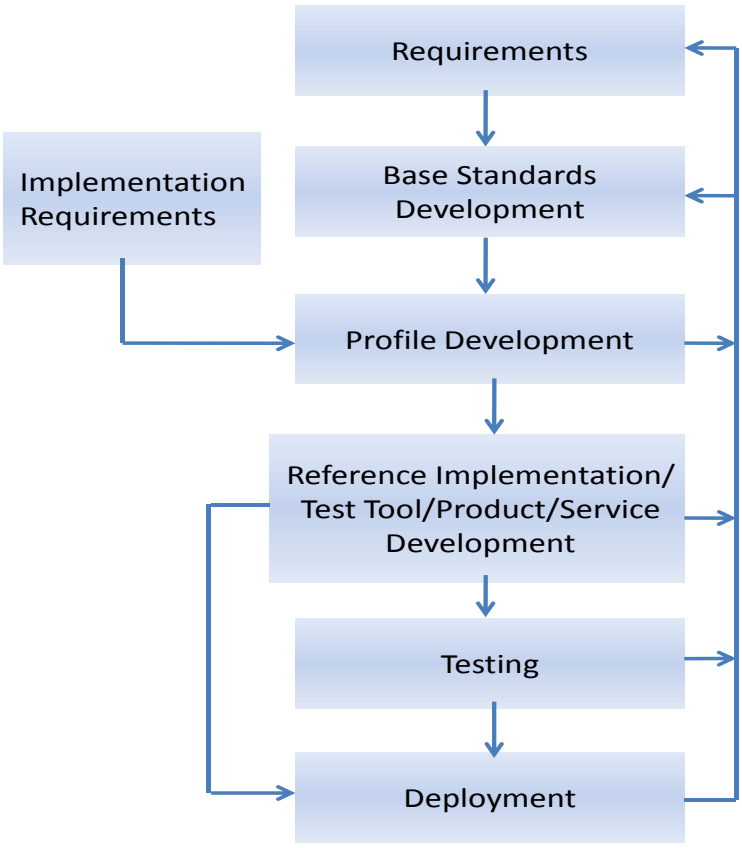

Figure 7 - ICT Standards Life Cycle

5.2 Categorizing the Status of Standards

Innovation in ICT means that ICT standards are constantly being developed, approved, and maintained. Revisions to previous editions of standards may or may not be backward compatible. Table 4 below is intended to provide an indication of the maturity level of a standard. Some SDOs require two or more implementations before final approval of a standard. Such implementations may or may not be commercial products or services. In other cases, an SDO may be developing a standard while conforming commercial products or services are already being sold.

Standard Maturity Level	Description
No Standard	SDOs have not initiated any standard development projects.
Under Development	SDOs have initiated standard development projects. Open source projects initiated.
Approved Standard	SDO approved standard is available to public. Some SDOs require multiple implementations before final designation as a "standard".
Reference Implementation	Reference implementation available
Testing	Test tools are available. Testing and test reports are available.
Products/Services	Standards-based products/services are available.
Market Acceptance	Widespread use by many groups. De facto or de jure market acceptance of standards-based products/services.
Sunset	Newer standards (revisions or replacements) are under development.

Table 4 - Standards Maturity Model

F.2 Cloud Computing Standards for Interoperability

As it would be expected there are a broad range of capabilities and functions available in the various cloud provider interfaces currently available. This may indicate that we are still in the early days of cloud computing and consolidation has not yet occurred. While standardization of cloud interfaces are maturing, commonalities among provider interfaces can help us understand the key interoperability requirements and features.

The interfaces that are presented to cloud users can be broken down into two major categories, with interoperability determined separately for each category. As show in the diagrams below, each type of cloud offering presents an interface of each category.

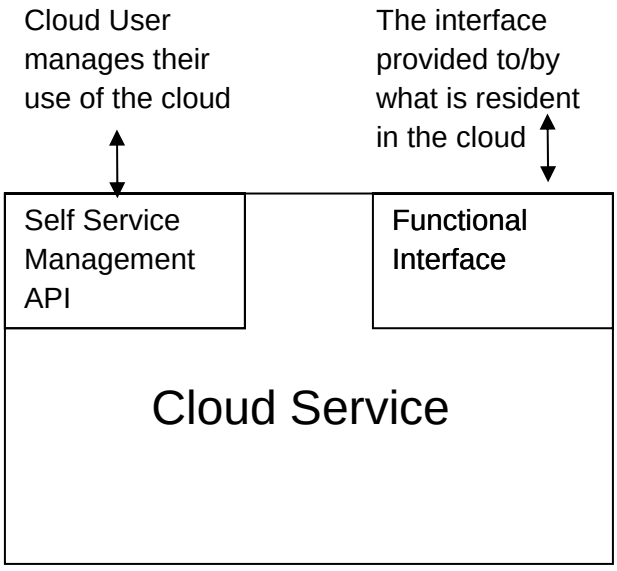

Figure 8 - Cloud Service presents an interface of each category

The interface that is presented to (or by) the contents of the cloud encompasses the primary *function* of the cloud service. This is distinct from the interface that is used to *manage* the use of the cloud service. For an Infrastructure as a Service cloud offering, as shown in the diagram below, the **Functional Interface** is a virtualized CPU, Memory and I/O space typically used by an operating system (and the stack of software running in that OS instance).

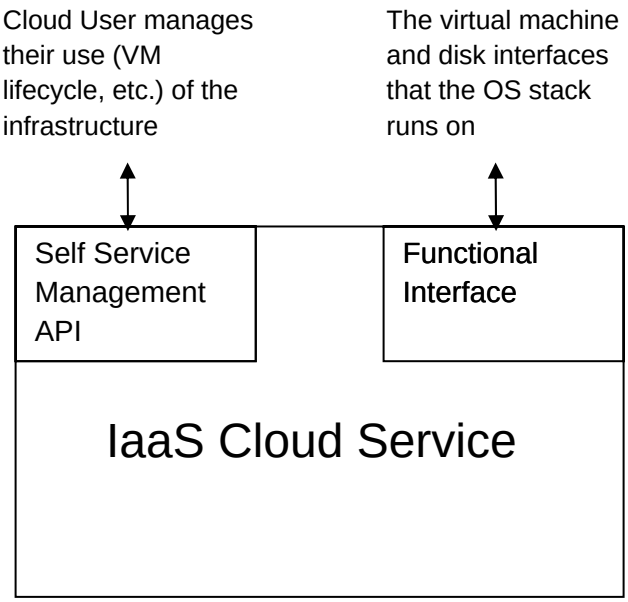

Figure 9 - An Infrastructure as a Service (IaaS) Interface

The cloud user utilizes the **Management Interface** to control their use of the cloud service by starting, stopping and manipulating virtual machine images and associated resources. It should be clear from this that the Functional Interface for an IaaS cloud is very much tied to the architecture of the CPU that is being virtualized. This is not a cloud specific interface and no effort is being put into a de jure standard for this interface since de facto CPU architectures are the norm.

The self-service IaaS management interface however is a candidate for interoperability standardization and there are several efforts in this space. The OCCI interface from the Open Grid Forum is an example of a standard IaaS resource management interface. The CDMI standard is an example of both storage management interface as well as a storage functional interface. There is a rapid proliferation of various proprietary interfaces as well as all competing to become a de facto means of interoperability.

For PaaS, as shown below, again we see the differentiation needed between these two categories of interfaces.

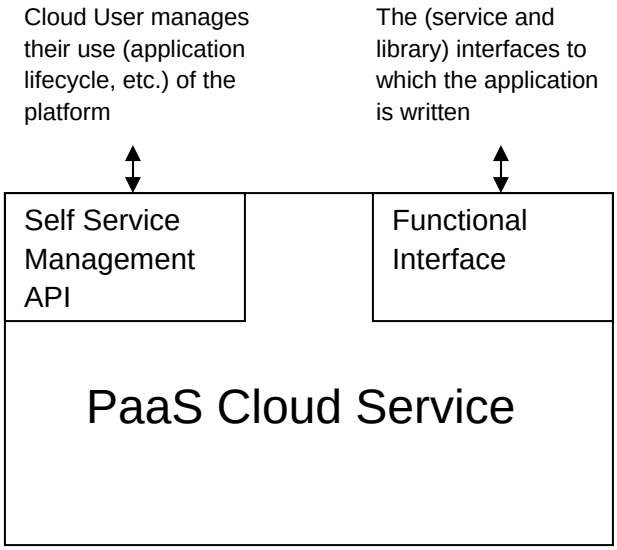

Figure 10 - Platform as a Service (PaaS) Interface

The Functional Interface of a PaaS offering is a runtime environment with a set of libraries and components to which the application is written. This could be offered in different languages and may or may not take advantage of existing application platforms standards such as J2EE or .Net. The Management Interface of a PaaS offering, however, may be very similar to the Management Interface of an IaaS offering. Instead of the lifecycle of virtual machines and their resources, the PaaS self-service interface is concerned with the lifecycle of applications and the platform resources they depend on. In addition, instead of being metered and billed on the basis of virtual hardware resources, the interface typically exposes metrics for platform service and runtime container usage. Interoperability of Paas self-service management interfaces can be achieved separately from the interoperability of the PaaS functional interfaces, although there seem to be very few efforts concentrating on PaaS management interfaces today.

For Software as a Service offering, as shown below, the Functional Interface is the same as the application interface of the software itself. In the case where a SaaS application is consumed through a web browser, there may be many standards that are used to achieve interoperability between what is essentially a web server and the user's browser, such as IP (v4, v6), TCP, HTTP. SSL/TLS, HTML and JavaScript/JSON, none of these web standards are cloud specific, and these same standards are being used in the many web browser-based Management Interfaces.

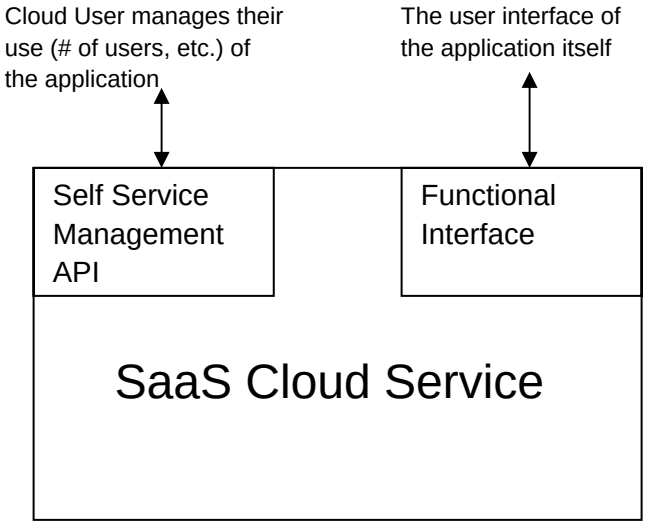

Figure 11 – Software as a Service (SaaS) Interface

The self-service Management Interface of a SaaS offering is typically concerned, not with lifecycle, but with the administration and customization of application functionality for each user of the offering. Through this interface, for example, additional users can be added (along with their credentials and permissions), additional features can be ordered for each user (usually in packaged sets), and an accounting of each user's consumption of the offering is available. Interoperability of a SaaS Management Interface may need to wait for SaaS offerings to be similar enough in their feature sets before a standard can be pursued.

Most of these interfaces will be tested and analyzed by NIST to validate its capabilities against the list of cloud computing use cases. At the same time work is continuing in the SDOs to further the interests of cloud computing interoperability – including the maintenance of standards to reflect implementation experience, development of new standards for agreed upon functions and/or protocols, and the profiling of existing standards.

F.3 Cloud Computing Standards for Portability

As described earlier in the document, the portability requirements for cloud computing are focused on two areas: system portability and data portability.

The rapid adoption of virtual infrastructure has popularized the practice of packaging, transporting and deploying pre-configured and ready-to-run systems, including all needed applications and the operating systems into virtual machines. The development of a standard, portable meta-data model for the distribution of virtual machines to and between virtualization and cloud platforms will enable the portability of such packaged workloads on any cloud computing platform. Some cloud workload formats contain a single VM only, modern enterprise applications are often constructed using a multiple tiered model, where each tier contains one or more machines. A single VM model is thus not sufficient to distribute a complete multi-tier system. In addition, complex applications require install-time customization of networks and other customer specific properties. Furthermore, a virtual machine image is packaged in a run-time format with hard disk images and configuration data suitable for a particular hypervisor. Run-time formats are optimized for execution and not for distribution. For efficient software distribution, a number of additional features become critical, including platform independence, compression, verification, signing, versioning, and software licensing management.

Over the last year much progress has been made on new standards in this area. Open Virtualization Format (OVF) from the Distributed Management Task Force (DMTF), for example, was developed to address portability concerns between various virtualization

platforms. It consists of meta-data about a virtual machine images or groups of images that can be deployed as a unit. It provides an easy way to package and deploy services as either a virtual appliance or used within an enterprise to prepackage known configurations of a virtual machine image or images. It may contain information regarding the number of CPUs, memory required to run effectively, and network configuration information. It also can contain digital signatures to ensure the integrity of the machine images being deployed along with licensing information in the form of a machine readable EULA (End User License Agreement) so that it can be understood before the image(s) is deployed.

A future direction of workloads data and metadata standardization is to help improve the automation of intercloud workload deployment. Concepts such as standardized SLAs (Service Level Agreements), sophisticated inter virtual machine network configuration and switching information and software license information regarding all of the various components that make up the workload are possibilities.

Another aspect of portability in the cloud environment is that of data (and metadata) portability between clouds, for example, between storage cloud services and between compatible application services in SaaS layer. For cloud storage services, as much of the actual data movement needs to be done in bulk moves of massive numbers of objects, retaining the data organization (into containers for example) and retaining the associated metadata are main portability requirements. This can be achieved by standardizing a canonical format for data and it is associated metadata to be packaged up and moved in bulk, including the format of the data and metadata package on physical media used in bulk transfers between the clouds in addition to on-the-wire format. CDMI, for example, does standardize a data packaging format that retains the organization and metadata associated with the data, and provides operations for packaging (serialization) and unpackaging (de-serialization) cloud data, which can be used to transfer data over the wire or via physical media.

F.4 Cloud Computing Standards for Security

The three cybersecurity objectives, ensuring the confidentiality, integrity, and availability of information and information systems, are particularly relevant as these are the high priority concerns and perceived risks related to cloud computing. Cloud Computing implementations are subject to local physical threats as well as remote, external threats. Consistent with other Application Areas, the threat sources include accidents, natural disasters and external loss of service, hostile governments, criminal organizations, terrorist groups, intentional and unintentional introduction of vulnerabilities through internal and external authorized and unauthorized human and system access, including but not limited to employees and intruders. The characteristics of Cloud Computing, significantly, multi-tenancy, and the implications of the three Service Models and four Deployment models heighten the need to consider data and systems protection in the context of logical as well as physical boundaries.

Possible types of attacks against Cloud Computing services include the following:

- Compromises to the confidentiality and integrity of data in transit to and from a cloud provider;
- Attacks which take advantage of the homogeneity and power of cloud computing environments to rapidly scale and increase the magnitude of the attack;
- Unauthorized access (through improper authentication or authorization, or vulnerabilities introduced during maintenance) to software, data, and resources in use by a cloud service consumer by another consumer;
- Increased levels of network-based attacks which exploit software not designed for an Internet threat model and vulnerabilities in resources which were formerly accessed through private networks ;
- Limited ability to encrypt data at rest in a multi-tenancy environment;
- Portability constraints resulting from non-standard application programming interfaces (APIs) which make it difficult for a cloud consumer to change to a new cloud service provider when availability requirements are not met;

- Attacks which exploit the physical abstraction of cloud resources and exploit a lack of transparency in audit procedures or records;
- Attacks that take advantage of virtual machines that have not recently been patched because they have been turned off; and
- Attacks which exploit inconsistencies in global privacy policies and regulations.

Security Objectives

Major security objectives for a Cloud Computing implementation include the following:

- Protect customer data from unauthorized disclosure or modification. This includes supporting identity management such that the customer has the capability to enforce identity and access control policies on users accessing cloud services. This includes the ability of a customer to make access to its data selectively available to other users.
- Protect from supply chain threats. This includes ensuring the trustworthiness and reliability of the service provider as well as the trustworthiness of the hardware and software used.
- Restrict unauthorized access to Cloud Computing infrastructure resources. This includes implementing security domains that have logical separation between computing resources (e.g. logical separation of customer workloads running on the same physical server by virtual machine [VM] monitors [hypervisors] in a multitenant environment) and using secure-by-default configurations.
- Design web applications deployed in a cloud for an internet threat model and embedding security into the software development process.
- Protect internet browsers from attacks to mitigate end-user security vulnerabilities. This includes taking measures to protect internet-connected personal computing devices by applying security software, personal firewalls, and patch maintenance.
- Display access control and intrusion detection technologies at cloud provider, and independent assessment to verify they are in place. This includes (but does not rely on) traditional perimeter security measures in combination with the domain security model.

Traditional perimeter security includes restricting physical access to network and devices, protecting individual components from exploitation through security patch deployment, default most secure configurations, disabling all unused ports and services, role based access control, monitoring audit trails, minimizing the use of privilege, antivirus software; and encrypting communications.

- Define trust boundaries between service provider(s) and consumers to ensure that the responsibility for providing security is clear.
- Support portability such that the customer can take action to change cloud service providers when needed to satisfy availability, confidentiality and integrity requirements. This includes the ability to close an account on a particular date and time, and to copy data from one service provider to another.

6. Cloud Computing Standards Gaps, Overlaps

The following diagram is the Combined Conceptual Reference Diagram from the NIST Reference Architecture Working Group. It has slightly changed since the 5th draft. Each sub-box could be used to drill down and map relevant standards in a table format. See below for a possible approach for security, interoperability, and portability standards.

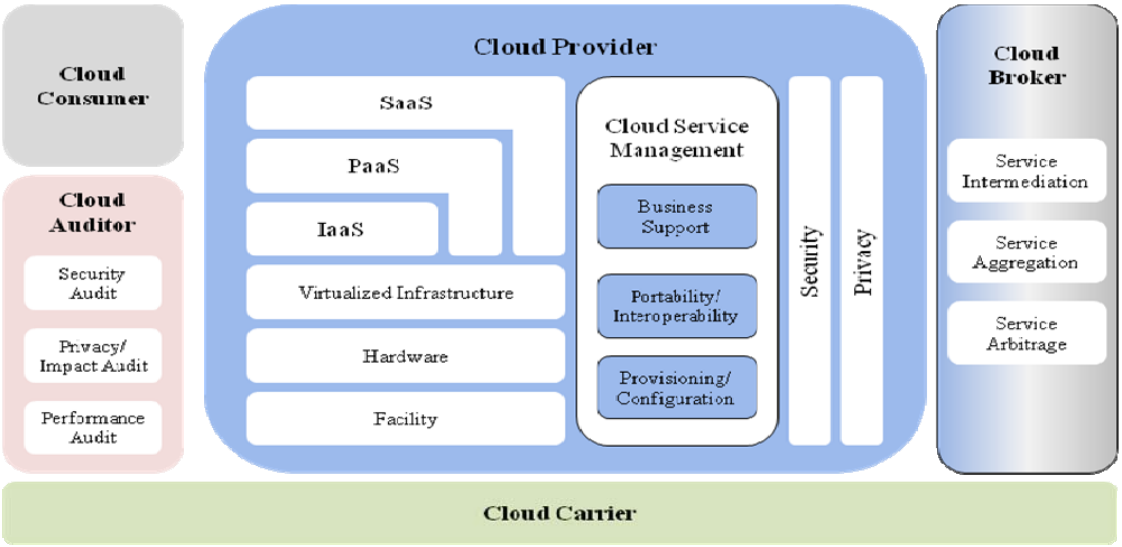

Figure 12 - The Combined Conceptual Reference Diagram

6.1 Security Standards Mapping

The table below maps standards to the security categories in the NIST Cloud Computing Taxonomy and gives their status (ref: Table 4, Standards Maturity Model). Some of the listed standards apply to more than one category and are therefore listed more than once.

Categorization	Available Standards and SDO	Status
Authentication & Authorization	RFC 5246: Secure Sockets Layer (SSL)/ Transport Layer Security (TLS); IETF	Approved Standard Market Acceptance
	RFC 3820: X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile; IETF	Approved Standard Products/Services?
	RFC 5849: OAuth (Open Authorization Protocol); IETF	Approved Standard Market Acceptance
	OpenID Authentication; OpenID	Approved Standard Market Acceptance
	eXtensible Access Control Markup Language (XACML); OASIS	Approved Standard Market Acceptance
	Security Assertion Markup Language (SAML); OASIS	Approved Standard Market Acceptance
Confidentiality	RFC 5246: Secure Sockets Layer (SSL)/ Transport Layer Security (TLS); IETF	Approved Standard Market Acceptance
	XML Encryption Syntax and Processing; W3C	Approved Standard Market Acceptance
	Key Management Interoperability Protocol (KMIP); OASIS	Approved Standard Market Acceptance?

Categorization	Available Standards and SDO	Status
Integrity	XML signature (XMLDSig); W3C	Approved Standard Market Acceptance
Identity Management	Service Provisioning Markup Language (SPML); OASIS	Approved Standard
	FIPS 201-1: Personal Identity Verification (PIV) of Federal Employees and Contractors, NIST	Approved Standard Marketplace Acceptance Draft FIPS 201-2 is out for comments
Security Monitoring & Incident Response	NIST SP 800-126: Security Content Automation Protocol (SCAP)	Approved Standard Testing Market Acceptance
	X.1520: Common vulnerabilities and exposures for determination; ITU-T	Approved Standard Pending Approval, April 2011
	X.1521; Common Vulnerability Scoring System; ITU-T	Approved Standard Pending Approval, April 2011
	PCI Data Security Standard; PCI	Approved Standard Market Acceptance
Security Policy Mgmt	eXtensible Access Control Markup Language (XACML); OASIS	Approved Standard Market Acceptance
Availability		

Table 5 - Security: Categorization

6.2. Interoperability Standards Mapping

The table below maps standards to the security categories in the NIST Cloud Computing Taxonomy and gives their status (ref: Table 4, Standards Maturity Model).

Categorization	Available Standards and SDO	Status
Service Interoperability	Web Services Interoperability Basic Security Profile 1.1 and 1.0; WS-I	Approved Standard Market Acceptance?
	Open Cloud Computing Interface (OCCI); Open Grid Forum	Approved Standard
	Cloud Data Management Interface (CDMI); Storage Networking Industry Association (SNIA)	Approved Standard

Table 6 - Interoperability: Categorization

6.3. Portability Standards Mapping

The table below maps standards to the security categories in the NIST Cloud Computing Taxonomy and gives their status (ref: Table 4, Standards Maturity Model). Some of the listed standards apply to more than one category and are therefore listed more than once.

Categorization	Available Standards and SDO	Status
Data Portability	Cloud Data Management Interface (CDMI); Storage Networking Industry Association (SNIA)	Approved Standard
System Portability	Open Virtualization Format (OVF); Distributed Management Task Force (DMTF)	Approved Standard - OVF v1.1.0 - INCITS 469-2010 Market Acceptance

Table 7 - Portability: Categorization

6.4 Roadmap Analysis

There are several facets of cloud service interfaces that are candidates for standardization including:

* Management APIs
* Data Exchange Formats
* Federated Identity
* Resource Descriptions
* Data Storage APIs

Based on the above candidate areas and the analysis of the business use cases, the following standards opportunities are examined:

6.4.1 Use Case: Creating, accessing, updating, deleting data objects in Clouds

Benefits: Cross-Cloud applications

Standardizations Needed: Standard interfaces to metadata and data objects

Possible Standards: CDMI

Priority: Near term

Availability: Now for CDMI 1.0 (Level 5)

6.4.2 Use Case: Moving VMs and virtual appliances between Clouds

Benefits: Migration. Hybrid Clouds. Disaster Recovery. Cloudbursting

Standardizations Needed: Common VM description format

Possible Standards: OVF from DMTF

Priority: Near term because OVF is available and an official standard

Availability: Now for OVF (Level 6)

6.4.3 Use Case: Selecting the best IaaS vendor for private externally hosted Cloud

Benefits: Provide cost-effective reliable deployments

Standardizations Needed: Resource and performance requirements description languages

Possible Standards: TBD

Priority: Medium term

Availability: TBD

6.4.4 Use Case: Portable tools for monitoring and managing Clouds

Benefits: Simplifies operations as opposed to individual tools for each Cloud

Standardizations Needed: Standard management interfaces to IaaS resources

Possible Standards: DMTF Cloud Management WG, OGF OCCI

Priority: Medium term

Availability: TBD

6.4.5 Use Case: Moving data between Clouds

Benefits: Migration between Clouds. Cross-cloud applications

Standardizations Needed: Standard metadata/data formats for movement between Clouds.

Vendor mappings between Cloud data and standard formats

Standardized query languages (e.g. for NoSQL for IaaS)

Possible Standards: TBD

Priority: Near term to avoid lock-in

Availability: TBD

6.4.6 Use Case: Single sign-on access to multiple Clouds

Benefits: Simplified access. Cross-cloud applications

Standardizations Needed: Federated identity and authorization

Possible Standards: OpenID, OAuth, OASIS, CSA outputs

Priority: Medium term

Availability: TBD

6.4.7 Use Case: Orchestrated processes across Clouds and Enterprise Systems

Benefits: Enhanced applications

Standardizations Needed: Standards for APIs and data movement

Possible Standards: Existing SOA standards and new Intercloud standards from IEEE

Priority: Long term because new standards must be developed and tested

Availability: TBD

6.4.8 Use Case: Discovering Cloud resources

Benefits: Selection of appropriate Clouds for applications

Standardizations Needed: Description languages for available resources. Catalog interfaces

Possible Standards: DMTF, TM Forum

Priority: Medium term

Availability:

6.4.9 Use Case: Evaluating SLAs and penalties

Benefits: Selection of appropriate Cloud resources

Standardizations Needed: SLA description language

Possible Standards: TBD

Priority: Long term because it is a hard problems

Availability: TBD

6.4.10 Use Case: Auditing Clouds

Benefits: Ensure regulatory compliance. Verify information assurance.

Standardizations Needed: Auditing standards and verification check lists

Possible Standards: CSA Cloud Audit

Priority: Near term because it is needed to avoid risky deployments

Availability: TBD

Ongoing Roadmap analysis should track the development of the standards and update the Standards Inventory as necessary.

7. USG Cloud Computing Standards Priorities

8. Conclusions and Recommendations

8.1. Recommendations for Accelerating the Development and Use of Cloud Computing Standards

8.1.1. Strategic Recommendations

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

- Agencies should actively participate in standards development projects that are of high priority to their agency missions. Cloud Standards naturally evolve in maturity, and agencies are advised to help accelerate this maturity by asking for vendors to show compliance with cloud standards in their roadmaps. When multiple vendors then offer an implementation of a standard, consider requiring that standard in RFPs.
- Agencies should support the concurrent development of conformity and interoperability assessment schemes to accelerate the development and use of technically sound standards and standards-based products, processes and services.
- Agencies should specify cloud computing standards in their procurements and grant guidance.
- The USG interagency group should be chartered to recommend specific cloud computing standards and best practices for USG-wide use that will support USG

requirements for interoperability, portability, and security in cloud computing applications.

8.1.2. Tactical (i.e., Near Term) Recommendations

- A listing of standards relevant to cloud computing should be maintained at Standards.gov.

Bibliography

Distributed Management Task Force (DMTF)

- Interoperable Clouds White Paper

DSP-IS0101 Cloud Interoperability White Paper V1.0.0
This white paper describes a snapshot of the work being done in the DMTF Open Cloud Standards Incubator, including use cases and reference architecture as they relate to the interfaces between a cloud service provider and a cloud service consumer.

- Architecture for Managing Clouds White Paper

DSP-IS0102 Architecture for Managing Clouds White Paper V1.0.0
This white paper is one of two Phase 2 deliverables from the DMTF Cloud Incubator and describes the reference architecture as it relates to the interfaces between a cloud service provider and a cloud service consumer. The goal of the Incubator is to define a set of architectural semantics that unify the interoperable management of enterprise and cloud computing.

- Use Cases and Interactions for Managing Clouds White Paper

DSP-IS0103 Use Cases and Interactions for Managing Clouds White Paper V1.0.0
This document is one of two documents that together describe how standardized interfaces and data formats can be used to manage clouds. This document focuses on use cases, interactions, and data formats.

Global Inter-Cloud Technology Forum (GICTF)

Use Cases and Functional Requirements for Inter-Cloud Computing

Published on August 2010

http://www.gictf.jp/doc/GICTF_Whitepaper_20100809.pdf

This whitepaper describes three areas of advantages of Inter-Cloud computing, which are, assured or prioritized performance, availability, and convenience of combined services. Several use cases of Inter Cloud Computing are provided with details according to these three areas, such as assured performance against transient overload, disaster recovery and service continuity for availability, and federated service provisions, followed by sequential procedures, functional requirements for each use case.
Essential functional entities and interfaces are identified to meet these described requirements.

NIST Special Publication 800-125 Guide to Security for Full Virtualization Technologies,

NIST Special Publication 800-144 DRAFT Guidelines on Security and Privacy Issues in Public Cloud Computing,

Annex A

Definitions

Information and Communications Technologies (ICT) -- encompasses all technologies for the capture, storage, retrieval, processing, display, representation, organization, management, security, transfer, and interchange of data and information.
[SOURCE: This report]

Interoperability The capability to communicate, execute programs, or transfer data among various functional units under specified conditions. [SOURCE: American National Standard Dictionary of Information Technology (ANSDIT)]

Maintainability A measure of the ease with which maintenance of a functional unit can be performed using prescribed procedures and resources. Synonymous with serviceability.
[SOURCE: American National Standard Dictionary of Information Technology (ANSDIT)]

Network Resilience – A computing infrastructure that provides continuous business operation (i.e., highly resistant to disruption and able to operate in a degraded mode if damaged), rapid recovery if failure does occur, and the ability to scale to meet rapid or unpredictable demands.
[SOURCE: The Committee on National Security Systems Instruction No 4009"National Information Assurance Glossary." CNSSI-4009]

Portability The capability of a program to be executed on various types of data processing systems with little or no modification and without converting the program to a different language.
[SOURCE: American National Standard Dictionary of Information Technology (ANSDIT)]

Portability: 1. The ability to transfer data from one system to another without being required to recreate or reenter data descriptions or to modify significantly the application being

transported. 2. The ability of software or of a system to run on more than one type or size of computer under more than one operating system.
[SOURCE: Federal Standard 1037C, Glossary of Telecommunication Terms, 1996]

Privacy,Information privacy is the assured, proper, and consistent collection, processing, communication, use and disposition of personal information (PI) and personally identifiable information (PII) throughout its life cycle.
[SOURCE: OASIS]

Reference implementation is an implementation of a standard to be used as a definitive interpretation for the requirements in that standard. Reference implementations can serve many purposes. They can be used to verify that the standard is implementable; validate conformance test tools; and support interoperability testing among other implementations. A reference implementation may or may not have the quality of a commercial product or service that implements the standard.
[SOURCE: This report]"

Reliability A measure of the ability of a functional unit to perform a required function under given conditions for a given time interval.
[SOURCE: American National Standard Dictionary of Information Technology (ANSDIT)]

Resilience is the ability to reduce the magnitude and/or duration of disruptive events to critical infrastructure. The effectiveness of a resilient infrastructure or enterprise depends upon its ability to anticipate, absorb, adapt to, and/or rapidly recover from a potentially disruptive event.
[SOURCE: CRITICAL INFRASTRUCTURE RESILIENCE FINAL REPORT AND RECOMMENDATIONS, NATIONAL INFRASTRUCTURE ADVISORY COUNCIL, SEPTEMBER 8, 2009]

Resilience is the adaptive capability of an organization in a complex and changing environment.
[SOURCE: ASIS International, ASIS SPC.1-2009, American National Standard, Organizational Resilience: Security, Preparedness, and Continuity Management System – Requirements with Guidance for Use.]

Security refers to information security. Information security means protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide:

- **Integrity,** which means guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity;
- **Confidentiality,** which means preserving authorized restrictions on access and disclosure, including means for protecting personal privacy and proprietary information; and
- **Availability,** which means ensuring timely and reliable access to and use of information.

[SOURCE: Title III of the E-Government Act, entitled the Federal Information Security Management Act of 2002 (FISMA)]

Standard document, established by consensus and approved by a recognized body that provides for common and repeated use, rules, guidelines or characteristics for activities or their results, aimed at the achievement of the optimum degree of order in a given context. Note: Standards should be based on the consolidated results of science, technology and experience, and aimed at the promotion of optimum community benefits. [SOURCE: ISO/IEC Guide 2:2004, Standardization and related activities - General Vocabulary, definition 3.2]

Standard document may provide the requirements for: a product, process or service; a management or engineering process; or a testing methodology. An example of a product standard is the multipart ISO/IEC 24727, *Integrated circuit card programming interfaces*. An example of a management process standard is the ISO/IEC 27000, *Information security management systems*, family of standards. An example of an engineering process standard is

ISO/IEC 15288, System life cycle processes. An example of a testing methodology standard is the multipart ISO/IEC 19795, *Biometric Performance Testing and Reporting*.

Standards Developing Organization (SDO) is any organization that develops and approves standards using various methods to establish consensus among its participants. Such organizations may be: accredited, such as ANSI-accredited IEEE; or international treaty based, such as the ITU-T; or international private sector based, such as ISO/IEC; or an international consortium, such as OASIS or IETF; or a government agency. SOURCE: [This report]

Usability The extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use. [SOURCE: ISO 9241-11:1998 Ergonomic requirements for office work with visual display terminals (VDTs) -- Part 11: Guidance on usability and ISO/IEC 25062:2006 Software engineering -- Software product Quality Requirements and Evaluation (SQuaRE) -- Common Industry Format (CIF) for usability test reports]

Annex B

Acronyms

CDMI	Cloud Data Management Interface
CDN	Content Delivery Network
CIO	Chief Information Officer
CPU	Central Processing Unit
CRM	Customer Relationship Management
DISR	Defense IT Standards Registry
DoD	Department of Defense
I/O	Input/Output
IaaS	Cloud Infrastructure as a Service
ICT	Information and Communications Technologies
ISIMC	Information Security and Identity Management Committee
IT	Information Technology
MID	mobile internet devices
OCCI	Open Cloud Computing Interface
OS	Operating System
P2P	Peer-to-Peer
PaaS	Cloud Platform as a Service
PDA	Personal Digital Assistant
PI	Personal Information
PII	Personal Identifiable Information

PKI	Public key infrastructure
SaaS	Cloud Software as a Service
SAJACC	Standards Acceleration to Jumpstart Adoption of Cloud Computing
SDOs	Standards Developing Organizations
SLA	Service Level Agreement
USG	United States Government

Annex C

Standards Developing Organizations

Global Information and Communications Technologies (ICT) standards are developed in many venues. Such standards are created through collaborative efforts that have a global reach, are voluntary and widely adopted by the marketplace across national borders. These standards are developed not only by national-member based international standards bodies, but also by consortia groups and other organizations.

In July 2009, a Wiki site for Cloud Computing Standards coordination was established: Cloud-standards.org is. The goal of the site is to document the activities of the various SDOs working on Cloud Computing Standards.

The following is a list of SDOs that have standards projects and standards relevant to Cloud Computing.

Alliance for Telecommunications Industry Solutions (ATIS)

ATIS develops standards and solutions addressing a wide range of industry issues in a manner that allocates and coordinates industry resources and produces the greatest return for communications companies.
ATIS creates solutions that support the rollout of new products and services into the information, entertainment and communications marketplace. Its activities provide the basis for the industry's delivery of:

- Existing and next generation IP-based infrastructures;
- Reliable converged multimedia services, including IPTV;
- Enhanced Operations Support Systems and Business Support Systems; and
- Greater levels of service quality and performance.

ATIS is accredited by the American National Standards Institute (ANSI).

CloudAudit

The goal of CloudAudit is to provide a common interface and namespace that allows cloud computing providers to automate the Audit, Assertion, Assessment, and Assurance (A6) of their infrastructure (IaaS), platform (PaaS), and application (SaaS) environments and allow authorized consumers of their services to do likewise via an open, extensible and secure interface and methodology.

CloudAudit is a volunteer cross-industry effort from the best minds and talent in Cloud, networking, security, audit, assurance and architecture backgrounds.

The CloudAudit/A6 Working group was officially launched in January 2010 and has the participation of many of the largest cloud computing providers, integrators and consultants.

Cloud Management Working Group (CMWG)

The CMWG will develop a set of prescriptive specifications that deliver architectural semantics as well as implementation details to achieve interoperable management of clouds between service requestors/developers and providers. This WG will propose a resource model that at minimum captures the key artifacts identified in the Use Cases and Interactions for Managing Clouds document produced by the Open Cloud Incubator.

Using the recommendations developed by DMTF's Open Cloud Standards Incubator, the cloud management workgroup (CMWG) is focused on standardizing interactions between cloud environments by developing specifications that deliver architectural semantics and implementation details to achieve interoperable cloud management between service providers and their consumers and developers.

Distributed Management Task Force (DMTF)

Open Virtualization Format (OVF)

DSP0243 Open Virtualization Format (OVF) V1.1.0

OVF has been designated as ANSI INCITS 469 2010

This specification describes an open, secure, portable, efficient and extensible format for the packaging and distribution of software to be run in virtual machines.

Open Cloud Standards Incubator

DMTF's Open Cloud Standards Incubator focused on standardizing interactions between cloud environments by developing cloud management use cases, architectures and interactions. This work was completed in July 2010. The work has now transitioned to the Cloud Management Working Group.

Institute of Electrical and Electronic Engineers (IEEE)

With approximately 350,000 members, the Institute of Electrical and Electronic Engineers (IEEE) is the world's largest technical professional society. The IEEE Standards Association (IEEE-SA) coordinates the efforts of experts throughout the IEEE in the development of standards such as key standards in the areas of computers, power and healthcare, and has 20,000 plus participants worldwide, including individuals in corporations, organizations, universities, and government agencies. An example IEEE of cyber security standards is the wireless local area network (WLAN) computer communication security standards (e.g., IEEE 802.11 series).

The Internet Engineering Task Force (IETF)

The Internet Engineering Task Force (IETF) issues the standards and protocols used to protect the Internet and enable global electronic commerce. The IETF develops cyber security standards for the Internet. Current activities include Public Key Infrastructure Using X.509 (PKIX), Internet Protocol Security (IPsec), Transport Layer Security (TLS), Secure Electronic Mail (S/MIME V3), DNS Security Extensions (DNSSEC), and Keying and Authentication for Routing Protocols (karp). Another IETF standard is the Incident Object Description Format

(IODEF), which provides a framework for sharing information commonly exchanged by Computer Security Incident Response Teams (CSIRTs) about computer security incidents. IODEF is an underpinning for the National Information Exchange Model (NIEM), which enables jurisdictions to effectively share critical information on cyber incident management, security configuration management, security vulnerability management, etc.

International Society of Automation (ISA)

The International Society of Automation (ISA) develops consensus standards for automation and industrial control systems. Since 1949, over 150 standards have been developed by over 4,000 industry experts around the world. The ISA Standards Committee, ISA99, Industrial Automation and Control System Security, is developing a multipart standard for security for industrial automation and control systems. A sister committee is ISA100, Wireless Systems for Automation.

International Organization for Standardization/International Electrotechnical Commission Joint Technical Committee 1 Information Technology (ISO/IEC JTC 1)

ISO/IEC JTC 1, Information Technology, develops international ICT standards for global markets. ISO and IEC are private sector international standards developing organizations. In 1987, ISO and IEC established a joint Technical Committee by combining existing ICT standards groups within ISO and IEC under a new joint Technical Committee, JTC 1. JTC 1 members are National Standards Bodies of different countries. Presently, there are 66 members. Approximately 2100 technical experts from around the world work within JTC 1. There are presently 18 JTC 1 Subcommittees (SCs) in which most of JTC 1 standards projects are being developed.

JTC 1 SC 27 (IT Security Techniques) is the one JTC 1 SC that is completely focused on cyber security standardization. SC 27 approved the establishment of a six-month study period (SP) that ends in April 2011. The purpose of the SP is to investigate the security requirements for cloud computing and what would be a feasible program of standards work to meet these requirements. The study period involves SC27 WG 1 (Information Security Management), WG

4 – Security Control and Services, and WG 5 – Identity Management, Privacy Technology and Biometrics. It is likely that SC 27 will proceed with some form of Cloud work by October-2011). Many other JTC 1 SCs are directly involved in specific standards critical to cyber security, including SC 6 (public key infrastructure [PKI] certificates), SC 7 (software and systems engineering), SC 17 (identification cards and related devices), SC 22 (programming languages, software environments and system software interfaces), and SC 37 (biometrics). In October 2009, JTC 1 established a new SC 38 for standardization in the areas of web services, Service Oriented Architecture (SOA), and cloud computing. SC38 initiated a Cloud Computing Study that will end in September 2011. The primary purpose of the Study is to analyse Cloud computing standardization activities and to recommend new SC38 cloud computing standardization projects.

International Organization for Standardization Technical Committee 68 (ISO TC 68)
ISO TC 68, Financial Services, develops international standards in the field of banking, securities and other financial services. ISO TC 68 Subcommittee 2 (SC 2) develops international standards on security management and techniques applicable to general banking operations such as public key management and encryption algorithms.

International Organization for Standardization Technical Committee 223 ISO TC 223
ISO TC 223, Societal Security, develops standards in the area of societal security, aimed at increasing crisis management and business continuity capabilities (i.e. through improved technical, human, organizational, and functional interoperability) as well as shared situational awareness, amongst all interested parties.

ITU Telecommunication Standardization Sector ITU-T

The ITU-T develops international standards for the ICT infrastructure including voice, data, and video. ITU-T established a Focus Group on Cloud Computing (FG Cloud) - http://www.itu.int/en/ITU-T/focusgroups/cloud/Pages/default.aspx . The charter of the FG Cloud is to investigate standards needed to support services/applications of cloud computing that make use of telecommunication networks; specifically:

* identify potential impacts on standards development and priorities for standards needed to promote and facilitate telecommunication/ICT support for cloud computing
* investigate the need for future study items for fixed and mobile networks in the scope of ITU-T
* analyze which components would benefit most from interoperability and standardization
* familiarize ITU-T and standardization communities with emerging attributes and challenges of telecommunication/ICT support for cloud computing
* analyze the rate of change for cloud computing attributes, functions and features for the purpose of assessing the appropriate timing of standardization of telecommunication/ICT in support of cloud computing

The Focus Group is collaborating with the worldwide cloud computing communities (e.g., research institutes, forums, academia) including other SDOs and consortia. The ITU-T Study Groups involved in standards relevant to cloud computing include: SG-13 (Next Generation Networks); and SG-17 (Network Security).

Kantara Initiative

Kantara Initiative was established on April 20, 2009, by leaders of several foundations and associations working on various aspects of digital identity, aka “the Venn of Identity”. It is intended to be a robust and well-funded focal point for collaboration to address the issues we each share across the identity community: Interoperability and Compliance Testing; Identity Assurance; Policy and Legal Issues; Privacy; Ownership and Liability; UX and Usability;

Cross-Community Coordination and Collaboration; Education and Outreach; Market Research; Use Cases and Requirements; Harmonization; and Tool Development.

Organization for the Advancement of Structured Information Standards OASIS

Founded in 1993, OASIS is a not-for-profit consortium. OASIS develops open standards for the global information society. The consortium produces Web services standards along with standards for security, e-business, and standardization efforts in the public sector and for application-specific markets. OASIS has more than 5,000 participants representing over 600 organizations and individual members in 100 countries. OASIS has a number of projects related to Cloud Computing including: ID Cloud, SSTC, WSSX, E- gov, and iD Trust Community of Practice. OASIS security, access and identity policy standards relevant to cloud computing include: SAML, XACML, SPML, WS-Security Policy, WS-Trust

The Open Cloud Consortium (OCC)

OCC is a member driven organization that develops reference implementations, benchmarks and standards for cloud computing. The OCC operates clouds testbeds, such as the Open Cloud Testbed and the OCC Virtual Network Testbed. The OCC also manages cloud computing infrastructure to support scientific research, such as the Open Science Data Cloud.

Open Cloud Computing Interface (OCCI) Working Group

The purpose of this group is the creation of a practical solution to interface with Cloud infrastructures exposed as a service (IaaS). We will focus on a solution which covers the provisioning, monitoring and definition of Cloud Infrastructure services. The group should create this API in an agile way as we can have advantages over other groups if we deliver fast. Overlapping work and efforts will be contributed and synchronized with other groups.

- Open Cloud Computing Interface Specification
- Open Cloud Computing Interface Terms and Diagrams

OGF and SNIA have collaborated on a Cloud Storage for Cloud Computing whitepaper.

Open Grid Forum (OGF)

Open Grid Forum (OGF) is a leading standards development organization operating in the areas of grid, cloud and related forms of advanced distributed computing. The OGF community pursues these topics through an open process for development, creation and promotion of relevant specifications and use cases.

OGF engages partners and participants throughout the international arena to champion architectural blueprints related to cloud and grid computing and the associated specifications to enable the pervasive adoption of advanced distributed computing techniques for business and research worldwide.

Advanced computing built on OGF standards enables organizations to share computing and information resources across department and organizational boundaries in a secure, efficient manner. Organizations throughout the world use production distributed architectures built on these features to collaborate in areas as diverse as scientific research, drug discovery, financial risk analysis and product design. The capacity and flexibility of distributed computing enables organizations to solve problems that until recently were not feasible to address due to interoperability, portability, security, cost and data-integration constraints.

Clouds, grids and virtualized distributed architectures reduce costs through automation and improved IT resource utilization and improve organizational agility by enabling more efficient business processes. OGF's extensive experience has enabled distributed computing built on these architectures to become a more flexible, efficient and utility-like global computing infrastructure.

Standardization is the key to realizing the full vision and benefits of distributed computing. The standards developed by OGF enable the diverse resources of today's modern computing

environment to be discovered, accessed, allocated, monitored and managed as interconnected flexible virtual systems, even when provided by different vendors and/or operated by different organizations.

Object Management Group (OMG)

The OMG was founded in 1989 and develops standards for enterprise integration. Its membership is international and is open to any organization, both computer industry vendors and software end users. Specific cloud-related specification efforts have only just begun in OMG, focusing on modeling deployment of applications & services on clouds for portability, interoperability & reuse.

Storage Networking Industry Association (SNIA)

SNIA Cloud TWG

The SNIA has created the Cloud Storage Technical Work Group for the purpose of developing SNIA Architecture related to system implementations of Cloud Storage technology. The Cloud Storage TWG:

Acts as the primary technical entity for the SNIA to identify, develop, and coordinate systems standards for Cloud Storage.
Produces a comprehensive set of specifications and drives consistency of interface standards and messages across the various Cloud Storage related efforts.
Documents system-level requirements and shares these with other Cloud Storage standards organizations under the guidance of the SNIA Technical Council and in cooperation with the SNIA Strategic Alliances Committee

SNIA Cloud Data Management Interface (CDMI)

The CDMI specification is now a SNIA Architecture standard and will be submitted to the INCITS organization for ratification as an ANSI and ISO standard as well.

SNIA CDMI Reference Implementation

The first working draft release of the Reference Implementation of CDMI is now available for download.

SNIA Terms and Diagrams

SNIA and OGF have collaborated on a Cloud Storage for Cloud Computing whitepaper. A demo of this architecture has been implemented and shown several times. More information can be found at the Cloud Demo Google Group.

Cloud Data Management Interface (CDMI) now has a working draft reference implementation available. Download and implement: http://snia.org/cloud

The Trusted Computing Group TCG

The TCG is a not-for-profit organization formed to develop, define and promote open, vendor-neutral, industry standards for trusted computing building blocks and software interfaces across multiple platforms. TCG has approximately 100 members from across the computing industry, including component vendors, software developers, systems vendors and network and infrastructure companies.

World Wide Web Consortium W3C

Founded in 1994, the W3C is a non-incorporated international community of 334 Member organizations that develop standards in support of Web technologies. The W3C work in the area of cyber security standards includes secure transferring data from one domain to another domain or between applications with well defined document authentication. XML Encryption and XML Signature are key pieces of the XML security stack.

Annex D

Conceptual Models and Architectures

General reference models:

- Distributed Management Task Force (DMTF): Cloud Service Reference Architecture
- Cloud Computing Use Case Discussion Group: a taxonomy for cloud computing
- IBM: Cloud Reference Architecture
- Cloud Security Alliance: Cloud Reference Model
- Cisco Cloud Reference Architecture Framework
- IETF: Cloud Reference Framework

Reference models focusing on specific application requirements:

- Open Security Architecture: Secure Architecture Models
- GSA: FCCI (Federal Cloud Computing Initiative)
- Juniper Networks: Cloud-ready Data Center Reference Architecture
- SNIA standard: Cloud Data Management Interface
- Elastra: A Cloud Technology Reference Model for Enterprise Clouds

[1] Trade Agreements Act of 1979, as amended (TAA) , the National Technology Transfer and Advancement Act (NTTAA), and The Office of Management and Budget (OMB) Circular A-119 Revised: Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities

Annex E

Examples of USG Criteria for Selection of Standards

USG Approach to Selecting Standards

F.2 USG Analysis Model for Selection of Private Sector Consensus Standards to be E-Gov Standards

The NIST E-Gov Standards Resource Center at Standards.gov includes the following list of questions that USG agencies can use when evaluating private sector consensus standards for agency use:

Applicability of standard

Is it clear who should use the standard and for what applications?
How does the standard fit into the Federal Enterprise Architecture (FEA)?
What was done to investigate viable alternative standards (i.e., due diligence) before selecting this standard?

Availability of standard

- Is the standard published and publicly available?
- Is a copy of the standard free or must it be purchased?
- Are there any licensing requirements for using the standard?

Completeness of standard

- To what degree does the candidate standard define and cover the key features necessary to support the specific E-Gov functional area or service?

Implementations to standard

- Does the standard have strong support in the commercial marketplace?
- What commercial products exist for this standard?
- Are there products from different vendors in the market to implement this standard?
- Are there any existing or planned mechanisms to assess conformity of implementations to the standard?

Interoperability of standard

- How does this standard provide users the ability to access applications and services through Web services?
- What are the existing or planned mechanisms to assess the interoperability of different vendor implementations?

Legal considerations

- Are there any patent assertions made to this standard?
- Are there any IPR assertions that will hinder USG distribution of the standard?

Maturity of standard

- How technically mature is the standard?
- Is the underlying technology of the standard well-understood (e.g., a reference model is well-defined, appropriate concepts of the technology are in widespread use, the technology may have been in use for many years, a formal mathematical model is defined, etc.)?
- Is the standard based upon technology that has not been well-defined and may be relatively new?

Source of standard

- What standards body developed and now maintains this standard?
- Is this standard a de jure or de facto national or international standard?
- Is there an open process for revising or amending this standard?

Stability of standard

- How long has this standard been used?
- Is the standard stable (e.g., its technical content is mature)?
- Are major revisions or amendments in progress that will affect backward compatibility with the approved standard?
- When is the estimated completion date for the next version?

F.2 Department of Defense (DOD)

The DISR mandates the minimum set of IT standards and guidelines for the acquisition of all DoD systems that produce, use or exchange information. The Defense Information Systems Agency (DISA) is the executive agent for the DISR. The DoD IT Standards Registry (DISR) is updated three times a year.

Initial Standards Selection Criteria for Inclusion in the DISR

A number of criteria should be considered when evaluating a standard for inclusion in the DISR. Selection criteria include:

1) the source of the standard;
2) openness;
3) technology relevance;
4) maturity;
5) marketplace support;

6) "usefulness/utility"; and,

7) risk.

Criteria	Description
Source of the Standard	Recognized authority
	Cooperative stance
	Feedback
	Process
	Consensus
Openness	Ownership/IPR
	User Participation
	Vendor Participation
Technology Relevance	
Maturity	Planning Horizon
	Stability
	Revision Content & Schedule
Marketplace Support	Acceptance
	Commercial Viability
Usefulness/Utility	Well Defined Quality Attributes
	Services & Application Interoperability
Risk	Performance, maturity & stability issues

Table 8 - DOD Selection Criteria and Description Summary

Standards Source

DOD policy articulates a preference hierarchy based on the source (owner/sponsor/publisher) of the standard. Note that the 5th Priority, Military, has its own internal priority of international first and then DOD MIL-STDs.

The standards preference hierarchy is:

Priority	Standards Source Hierarchy	Example
1st	International	ISO, IEC, ITU
2nd	National	ANSI
3rd	Professional Society; Technology Consortia; Industry Association	IEEE; IETF; W3C; OASIS; GEIA
4th	Government	FIPS
5th	Military	MIL-STDS, STANAGS

Table 9 - DOD Standards Sources Preferences

The standard must be recognized as being available from a reputable and authoritative source. The responsible SDO/SSO must have an established position within the relevant technical, professional, and marketplace communities as an objective authority in its sphere of activity. This means that the standard has been created and approved/adopted/published via a formal process and configuration management of the standard has been established. Accreditation implies acceptance by a recognized authoritative SSO.

The Standards Selection Criteria also provides guidance for moving through the standards lifecycle that changes the category of a standard from "*emerging*" to "*mandated*" to "*inactive/retired*".

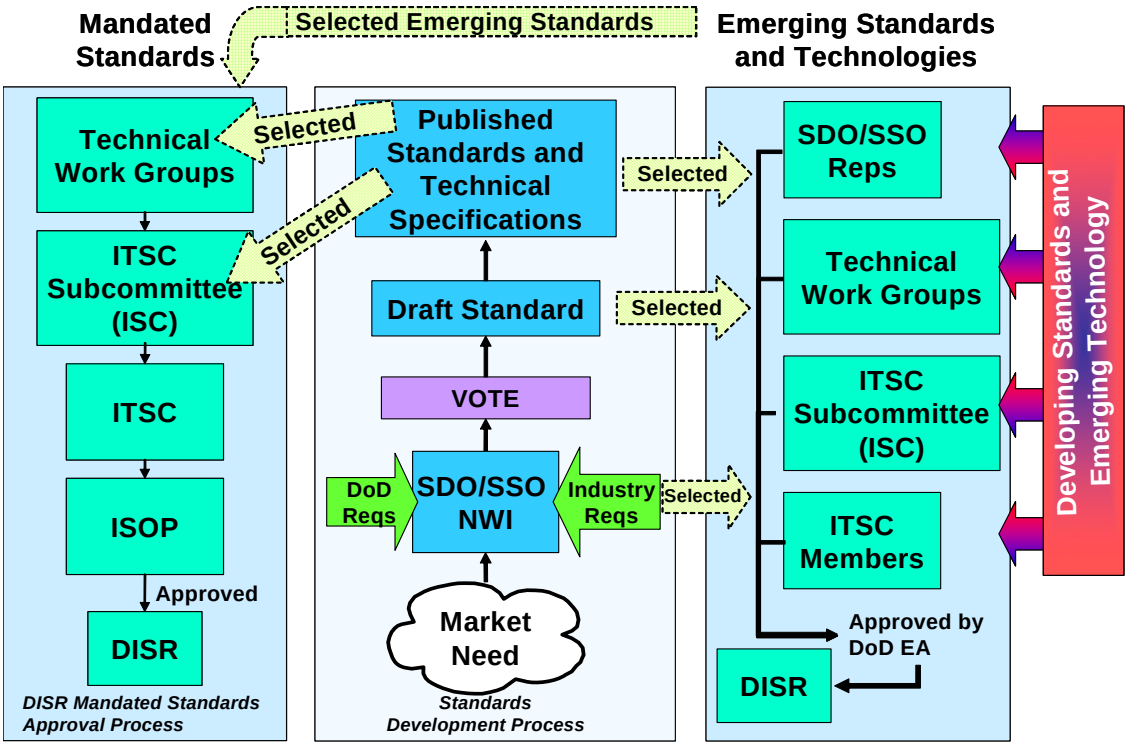

Figure 13 - DoD DISR Standards Selection Process