Document: **NIST CCSRWG – 047**

NIST Cloud Computing Standards Roadmap

Pre-Seventh Working Draft

March 7, 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.

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 model 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.

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.

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

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

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 Conceptual Model

3.1 Overview

The cloud computing conceptual model consists of five major actors in two tiers.

- ***First-tier actors:*** The core actors in all usage scenarios
  - ***Cloud Consumer***: Person or organization that maintains a business relationship with, and uses service from, *Cloud Providers*.
  - ***Cloud Provider*****:** Person, organization or higher-level system responsible for making a service available to *Cloud Consumers*.
- ***Second-tier actors:*** The supporting actors that assist in implementing cloud services. Actors in this tier are not required or not that critical for every cloud application.

- ***Cloud Broker:*** An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between *Cloud Providers* and *Cloud Consumers*.
- ***Cloud Auditor:*** A third-party that conducts an independent audit of the operations and determines the security of the cloud implementation.

- ***Intermediary actor:***
  - ***Cloud Carrier:*** The intermediary that provides connectivity and transport of cloud services.

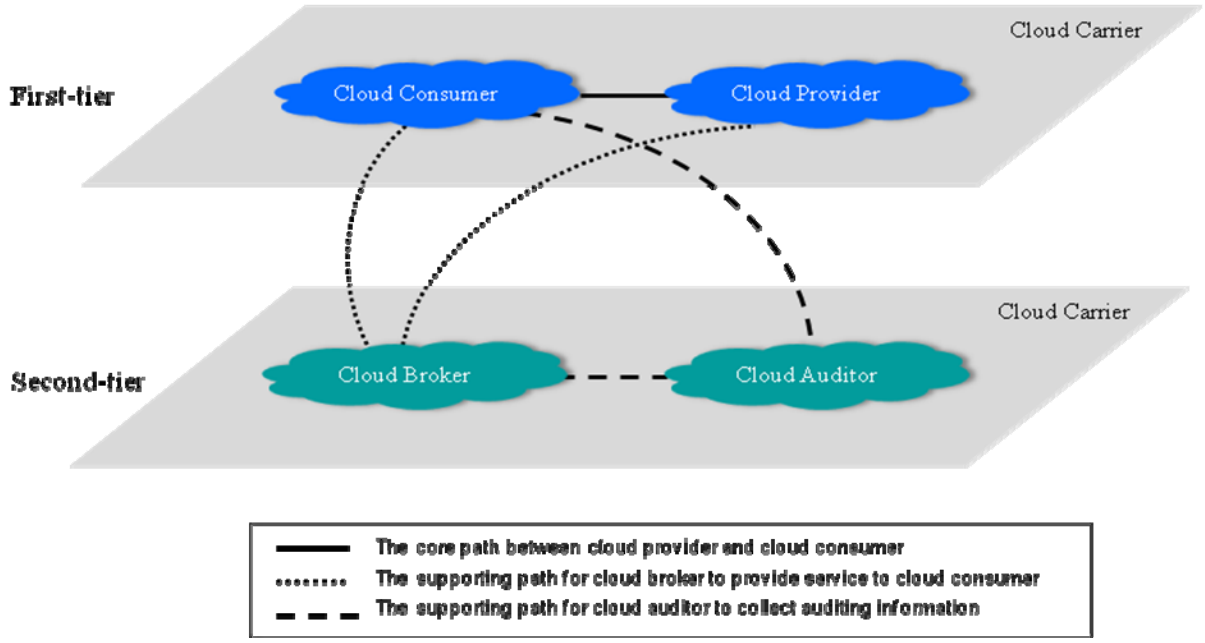

Figure 1 Five Major Actors

2.2 Cloud Consumer

Cloud consumers are categorized into three groups, based on the different usage scenarios and applications they use.

Type	Major Activities	Example Users
SaaS Consumers	Use application/service for biz process operations	Biz users
PaaS Consumers	Develop, test, deploy and manage services for application development.	Software developers, system developers, CIOs, IT managers
IaaS Consumers	Create/install, manage and monitor services for IT infrastructure operations.	System developers, IT managers

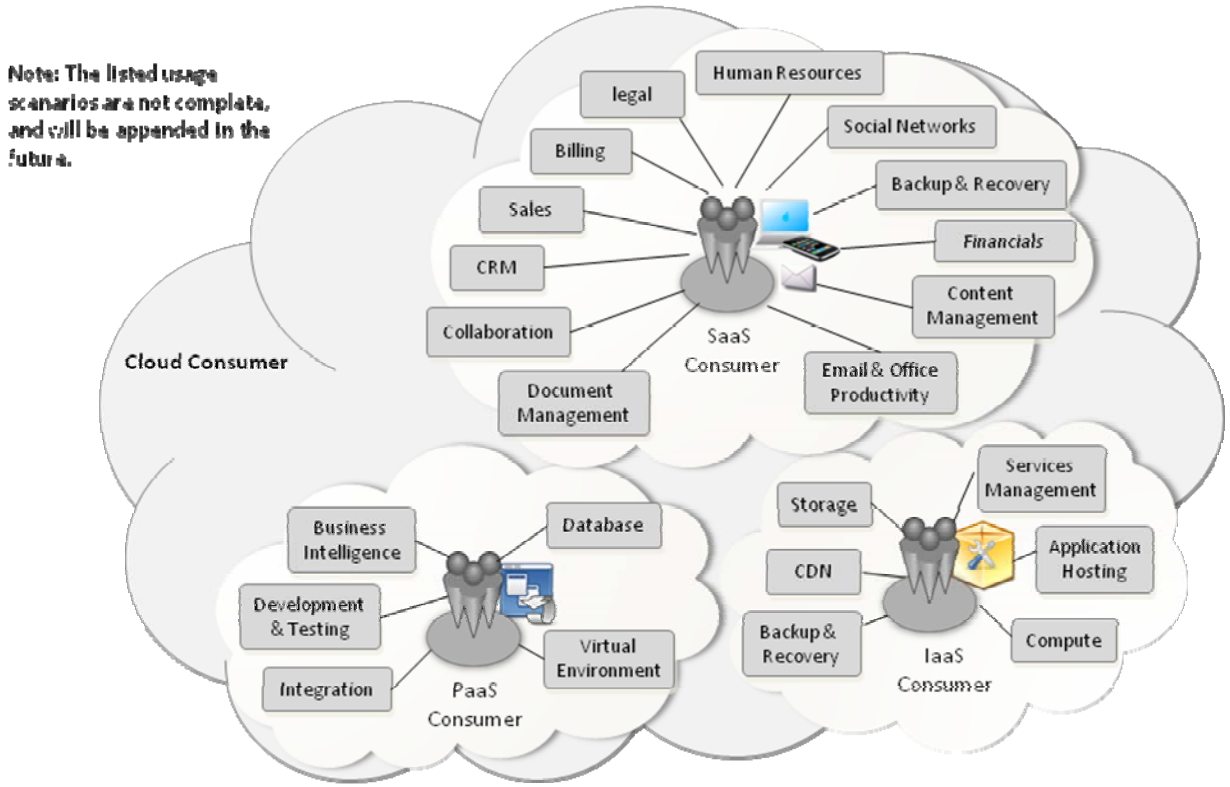

Figure 2 Cloud Consumer

3.2 Cloud Provider

Cloud service providers perform services to support the business processes of cloud service consumers at agreed service levels and costs. The providers perform different tasks for different service types.

Type	Major Activities
SaaS Provider	Install, manage and maintain the software
PaaS Provider	Manage the cloud infrastructure for the platform
IaaS Provider	Maintain the storage, database, message queue or other middleware, or the hosting environment for virtual machines

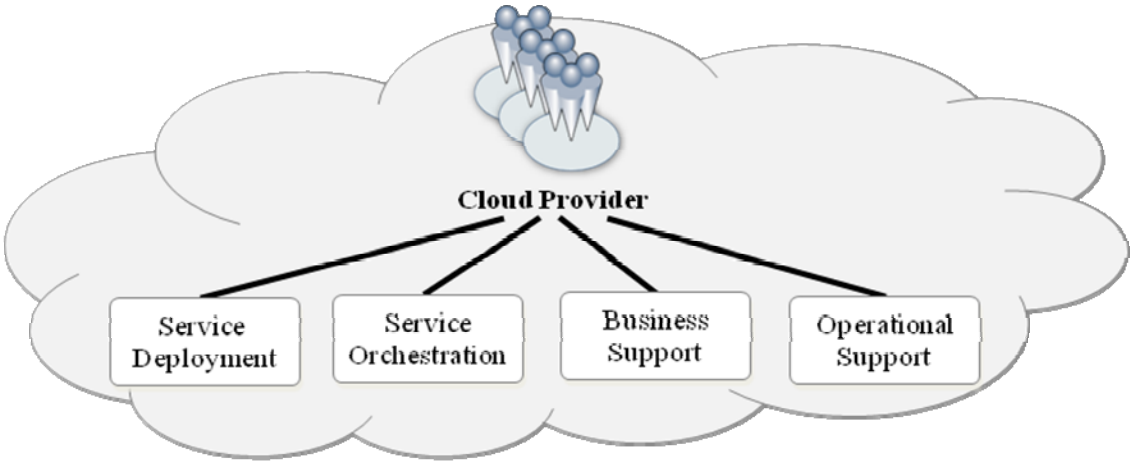

Figure 3 Cloud Provider

3.3 Cloud Carrier

The Cloud Carrier acts as an interconnection intermediary between cloud consumers and cloud providers. While an active participant of the cloud community the carrier does not create new information or alter information between the cloud consumer and cloud provider. Costs associated intermediary services are transacted outside of the cloud consumer/provider relationships. Cloud Carriers:

- Provide access to consumers through network and telecommunication access devices
  - Examples of network access devices include computers, laptops, mobile phones, mobile internet devices (MIDs), etc.

- Distribution normally provided by Network and Telecomm Carriers, or a transport agent
  - ***Transport agent:*** A business organization that provides physical transport of storage media such as high-capacity hard drives.
- For enterprise application deployed in the cloud, a cloud carrier with consistent SLA shall be required. In general, the cloud carrier may be required to provide dedicated and encrypted connections.

3.4 Cloud Broker

Cloud Broker is an entity that acts as an intermediary in establishing relationships between cloud tenants (consumers and providers) which manage the use, performance, and delivery of cloud services. Cloud Broker may create, alter content and transform information between the cloud consumer and cloud provider, including time shifting information delivery. Costs associated with Cloud Broker Services may be transacted outside of the cloud consumer/provider relationships.

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.

3.5 Cloud Auditor

A cloud auditor is a third-party that can conduct independent audits of the operations and determine the security of the cloud implementation.

A cloud auditor conducts 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.

- ***Security Assessment:*** Assess the management, operational, and technical controls of the cloud system with a frequency depending on risk, but no less than annually.
- ***Security Certification:*** A security certification is conducted for accrediting the cloud system. The security certification is a key factor in all security accreditation (i.e., authorization) decisions and is integrated into and spans the system development life cycle.
- ***Security Accreditation:*** The organization authorizes (i.e., accredits) the cloud system for processing before operations and updates the authorization or when there is a significant change to the system.

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

Figure 4 illustrates the different generic scenarios

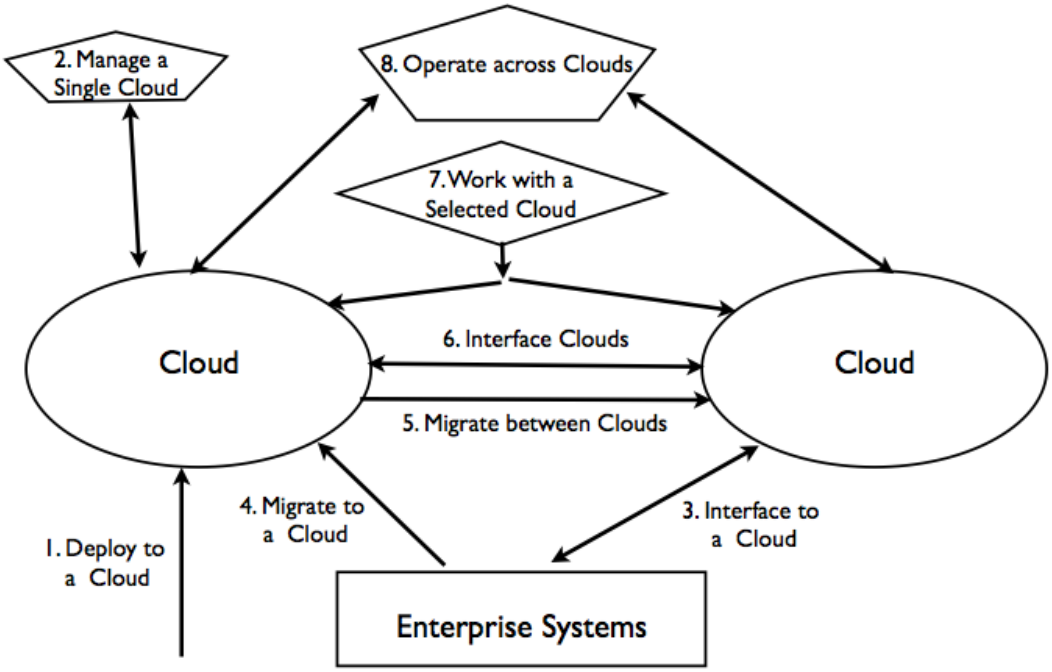

Figure 4 High Level Generic Scenarios

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

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

4.4 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

The technical requirement areas can be mapped to standardization status of those areas as *standards available*, *standards being developed*, and *new standards needed.* Along the following map, priorities of adopting standards or priorities in standardization in the area is also given based on the analysis of the business use cases.

4.4.1 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)

4.4.2 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)

4.4.3 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

4.4.4 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

4.4.5 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

4.4.6 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

4.4.7 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

4.4.8 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:

4.4.9 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

4.4.10 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.

5 Cloud Computing Standards

Standards work is underway in all of these areas. There are several standards that have been published from consortia. The Standards Roadmap Working Group has created a Standards Catalog that is available at:
http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/StandardsInventory
Many of the standards are for pre-Cloud technologies such as those designed for Web Services and the Internet. However, they are still relevant in addressing the technical requirements in cloud era.

5.1 ICT Standards Life Cycle

Figure 5 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 processes occur sequentially. Many of the processes illustrated can and should be done somewhat concurrently. Some of these processes (i.e., 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.

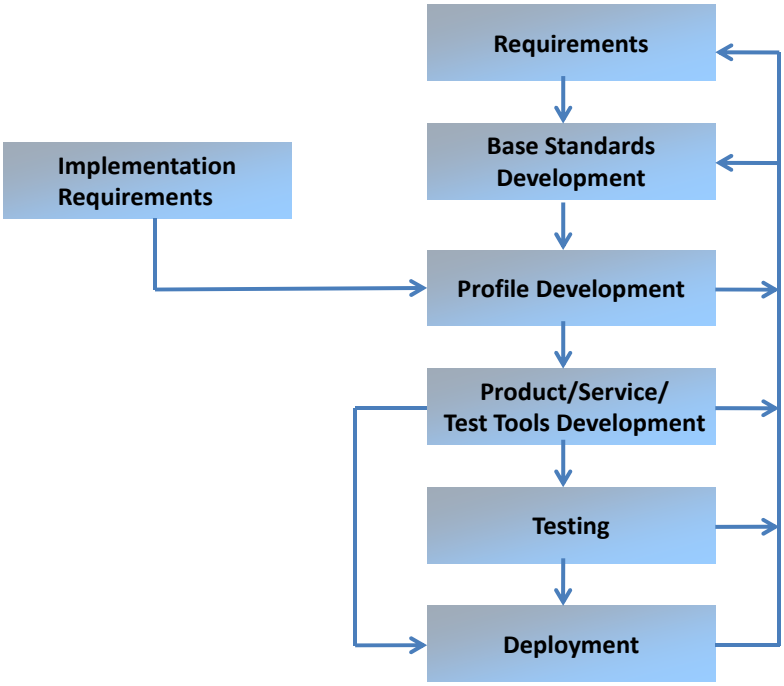

Figure 5 ICT Standards Life Cycle

5.2 Categorizing the Status of Standards

Standard Maturity Level	Description
1. No Standards	SDOs have not initiated any standard development projects.
2. Under Development	SDOs have initiated standard development projects. Open source projects initiated.
3. Approved	SDO approved standard is available to public.

Standard Maturity Level	Description
4. Reference Implementation	Reference implementation available
5. Testing	Test tools are available. Testing and test reports are available.
6. Product/Services	Standards-based products/services are available.
7. Market Acceptance	Widespread use by many groups. De facto or de jure market acceptance of standards-based products/services.
8. Sunset	Newer standards (revisions or replacements) are under development.

Table 1 Standards Maturity Model

5.3 USG Approach to Selecting Standards

5.3.1 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?
If the standard is proprietary, are there nevertheless many products readily available from a variety of vendors?
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 as 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?

5.3.2 DOD

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 2 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 3 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*".

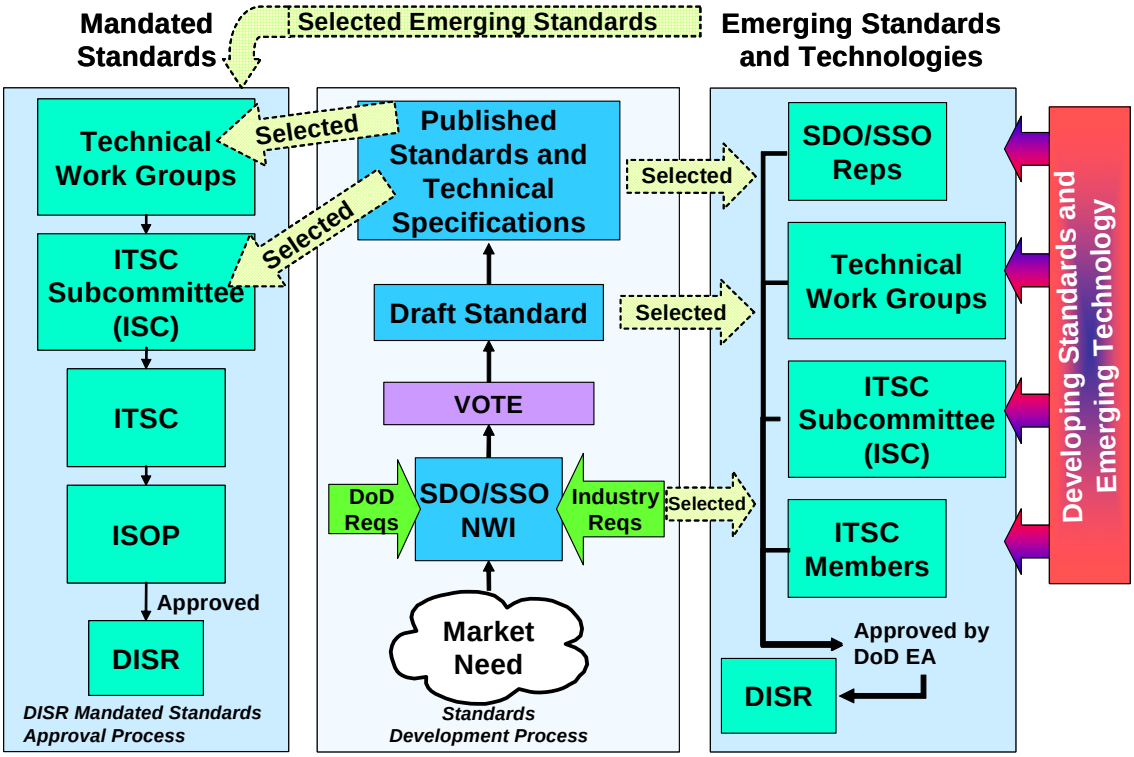

Figure 6 DOD DISR Standards Selection Process

5.4 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.

One of the key activities now underway is the SAJACC (Standards Acceleration to Jumpstart the Adoption of Cloud Computing) program. At the very end of 2010 new work groups were launched along with a Wiki to enable the industry and interested participants to start to develop a plan for evaluating current interface standards for cloud computing. As part of that work a Cloud Interface Catalogue of interfaces has been established. This provides a list of some of the cloud computing interfaces and their capabilities and links to more information about these interfaces. This provides a good place for consumers as well as cloud service providers to review what is currently available.

However, a good thing to notice is that there are many similarities between these provider interfaces. Another thing to notice is that while many of the features that have been implemented are similar, some of the interfaces are stronger in one or more areas.

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.

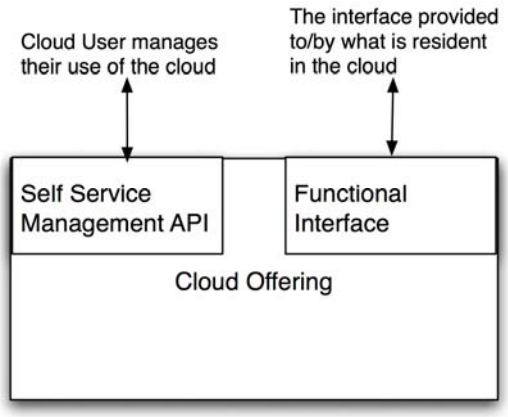

Figure 7

The interface that is presented to (or by) the contents of the cloud encompasses the primary *function* of the cloud offering. This is distinct from the interface that is used to *manage* the use of the cloud offering. 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).

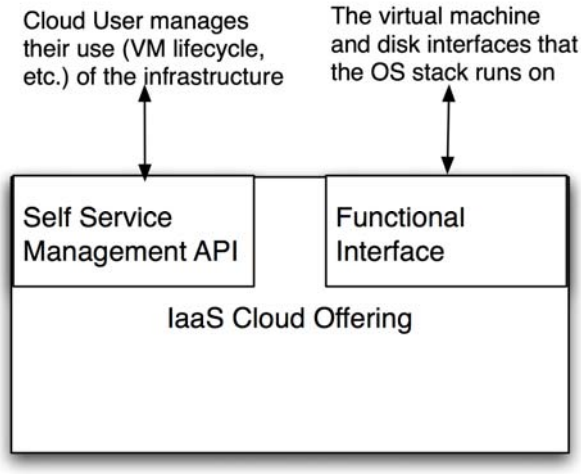

Figure 8

The cloud user utilizes the **Management Interface** to control their use of the cloud offering 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 API however is a candidate for de jure standardization and there are several efforts in this space. The OCCI interface from the Open Grid Forum is an example of a standard IaaS Management Interface. There is a rapid proliferation of various proprietary interfaces as well as all competing to become a de facto means of interoperability.

For cloud offerings in the Platform as a Service market, as shown below, again we see the differentiation needed between these two categories of interfaces.

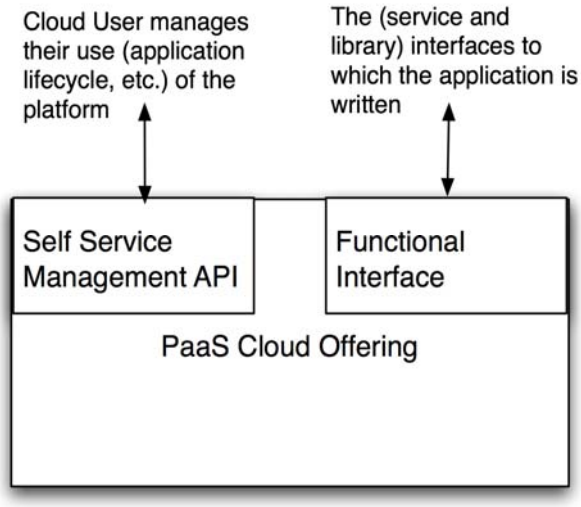

Figure 9

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 server 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 API 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 APIs can be achieved separate from the interoperability of the PaaS functional interfaces, although there seem to be very few efforts concentrating on this today.

For Software as a Service offering, as shown below, the Functional Interface is the same as the user interface of the application itself. 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. TLS, HTML and JavaScript/JSON. None of these web standards are cloud

specific, but these same standards are being used in the many Management Interfaces of this and the other types of offerings, interestingly enough.

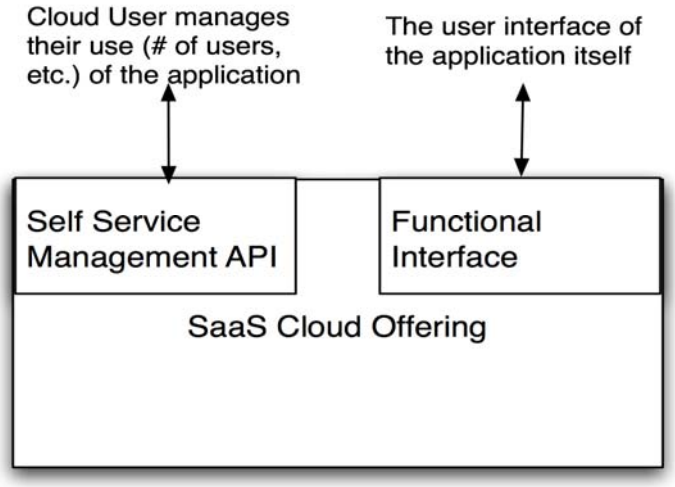

Figure 10

The self-service Management Interface of a SaaS offering is concerned, not with lifecycle, but with the application functionality for each user of the offering. Through this interface, 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.

Lastly, there are cloud storage (or Data Storage as a Service) offerings as shown below.

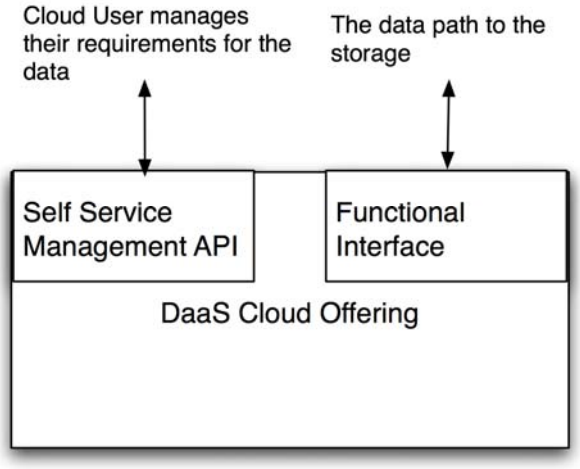

Figure 11

The Functional Interface for a DaaS offering is the means by which an application reads and writes their data to the storage cloud. For cloud computing situations, this interface may be one of several existing block (Fibre Channel, iSCSI) or file-based (NFS, CIFS/SMB) protocols for connecting to the orchestrated hypervisors running under the virtual machine guests. The Functional Interface may also be based on web standards and have a container and object orientation. This allows access to the data objects from any web browser, or from any of the other types of cloud offerings (IaaS, PaaS or SaaS). Interoperability between DaaS Functional Interfaces is key to minimizing the proliferation of interfaces to which an application must be coded to in order to read and write its data.

The Management Interface for a DaaS offering is concerned with monitoring the usage and services supporting the data stored in the cloud. The simplest and most intuitive way to manage data in the cloud is to state what the data requirements are and associate these with groupings of the data. If the cloud offering can meet these requirements (for backup, archive, replication, etc.), then it can bill the user for the data services that were used to achieve this. Standardizing the Management Interface for DaaS offerings is essential for moving the data and it is associated requirements from cloud to cloud and achieving these requirements in all cases. The CDMI

standard is an example of both a DaaS Management Interface as well as a DaaS Functional Interface based on web standards.

A provider's Management Interface defines how the developer and consumer interact with the provider for purposes of managing the use of the cloud offering. This architecture differentiates between service endpoints that accept (and respond to) messages over a protocol based on some message exchange pattern and the data elements and operations that an interface can support (data artifacts). The interface comprises both management interfaces and data artifacts. Elements of Management Interfaces, such as Service Catalogs and Service Managers, are programming interfaces (for example, APIs). Through these interfaces, developers and consumers interact with providers to request, deploy, administer, and use services. Examples of likely Management Interface elements are:

A Service Catalog, through which service offerings are offered, requested, and managed A Security Manager, through which the security-related aspects of a cloud are managed.

A Service Manager, through which instances of deployed services are managed and modified. Data artifacts are exchanged over the Management Interfaces. In this context, a data artifact definition describes the semantic content and the specific format (for example, the XML schema definition that describes the XML payload). Examples of data artifact types include service requests, service level-agreements (SLAs) and other contracts, service templates, service offerings, and images that contain applications. For example, a customizable contract template that includes the customer request, SLA, and security requirements is needed to support the service catalog interface. SLA, security requirements, and resource specifications are used to build offerings.

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 standards development organizations 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.

The work will continue throughout 2011 to evaluate and consolidate cloud computing interfaces with goals of reducing complexity, improving interoperability and ultimately improving customer flexibility and choice.

5.5 Cloud Computing Standards for Portability

The rapid adoption of virtual infrastructure has enabled the development of a standard, portable meta-data model for the distribution of virtual machines to and between virtualization and cloud platforms. Packaging an application together with the operating system on which it is certified, into a virtual machine that can be easily transferred from an ISV, through test and development and into production as a pre-configured, pre-packaged unit with no external dependencies, is extremely attractive. Such pre-deployed, ready to run applications packaged as virtual machines (VMs) in cloud computing are called work loads. In order to make this concept practical on a large scale it is important that the industry adopts a vendor-neutral standard for the packaging of such VMs and the meta-data that are required to automatically and securely install, configure, and run the workloads on any cloud computing platform.

Some cloud workload formats contain a single VM only, modern enterprise applications model service oriented architectures (SOA) with multiple tiers, where each tier contains one or more machines. A single VM model is thus not sufficient to distribute a multi-tier service. 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 portability, platform independence, verification, signing, versioning, and licensing terms.

Over the last year much progress has been made on new standards for improved cloud interoperability and reduced vendor locking. Existing standards are being adapted to address cloud computing portability such as the Open Virtualization Format (OVF) from the Distributed Management Task Force (DMTF).

OVF, for example, was originally 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 the terms can be understood before the image(s) is deployed.

OVF is currently being explored to validate if other metadata should be added 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.

This and other standard packaging formats need to be explored to enable portability of workloads for cloud computing.

An additional aspect of portability in the cloud environment is that of data (and metadata) portability between storage clouds. Although a standard functional interface for cloud storage helps move individual data objects from cloud to cloud, much of the actual data movement needs to be done in bulk moves of massive numbers of objects, retaining the organization (into containers for example) and retaining the associated metadata that may be driving the data services provided by the cloud. Interoperability here can be achieved by standardizing a canonical format for data and it is associated metadata to be packaged up and moved in bulk.

Because cloud storage offerings typically charge for the bandwidth in and out of the storage cloud as well as because of the prohibitive time it takes given the sheer size of the data transfer over even the fastest internet connections, most cloud storage offerings also offer to ship hard

disks containing the data for cloud import and export. Without a standard for the format of the data and metadata package on the disk (independent of the file system laid out on that disk), import of the data exported from a different cloud is problematic. CDMI 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.

5.6 Cloud Computing Standards for Security

Editor's Note: Should we use or refer to the Guide to Security for Full Virtualization Technologies, NIST Special Publication 800-125 or the DRAFT Guidelines on Security and Privacy Issues in Public Cloud Computing, NIST Special Publication 800-144 ?

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.

Editor's Note: NIST is suggesting that the following areas also be documented.

5.7 Cloud Computing Standards for Maintainability

5.8 Cloud Computing Standards for Usability

5.9 Cloud Computing Standards for Reliability

5.10 Cloud Computing Standards for Resilience

Editor's Note: As a result of the January 27, 2011 WG meeting, it was agreed to collect information on relevant standards via the TWiki page, Inventory of Standards Relevant to Cloud Computing. WG members have been requested to provide contributions by February 18, 2011. The collected information will then be mapped into the following sections.

6. Cloud Computing Standards Gaps, Overlaps

Editor's Note 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.

1

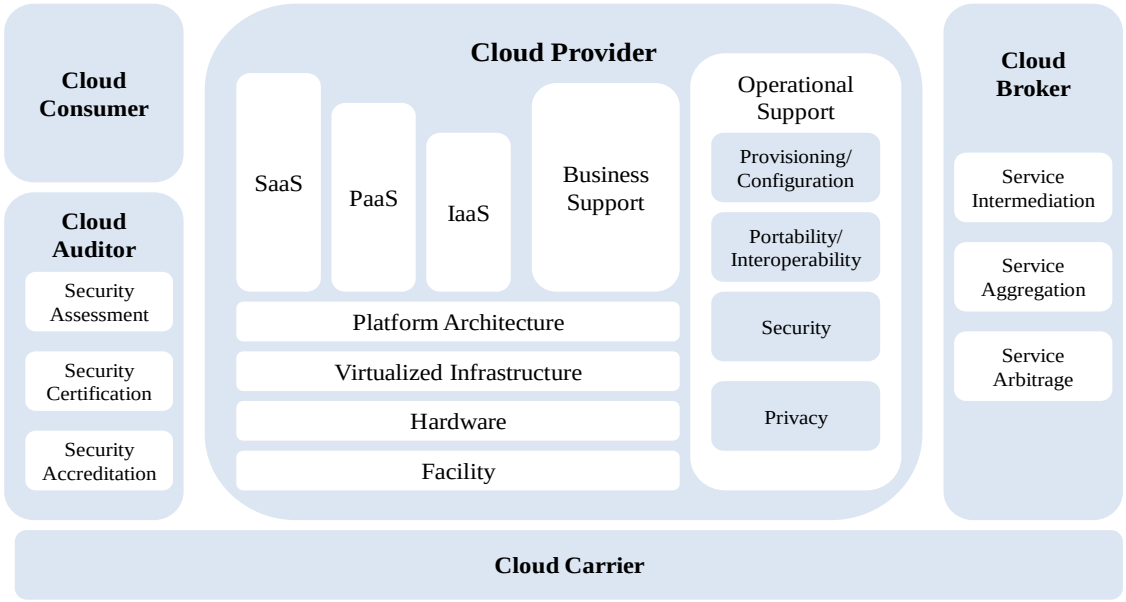

Figure 12 The Combined Conceptual Reference Diagram

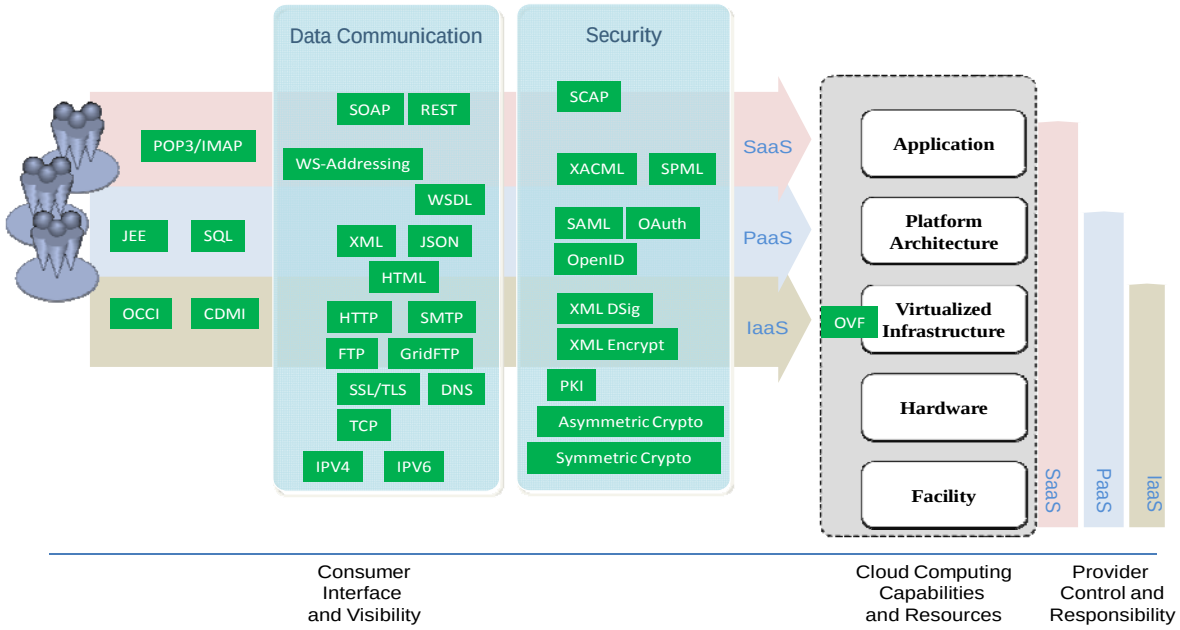

Figure 8 Relationships of Standards to Cloud Computing

6.1. Security

Add table by Fang.

6.2. Interoperability

Add table by Fang.

6.3. Portability

Add table by Fang.

6.4. Maintainability

6.5. Usability

6.6. Reliability

6.7. Resilience

7. Cloud Computing Standards Priorities

8. Conclusions and Recommendations

Editor's Note: The following are some suggestions from the NIST Editors.

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

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

8.2. Tactical Recommendations

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

Editor's Note: the following straw recommendations are from Bob Marcus (March 3, 2011)

- Require Cloud Providers to support DMTF OVF
- Require that Cloud Providers build adaptors to support CDMI interfaces to their Cloud storage
- Recommend acceleration of IaaS API standards e.g. harmonizing DMTF and OCCI activities
- Define core IaaS API functions and create adaptors to Provider Clouds (e.g. Cloud Broker layer)
- Create a Federated Identity capability for government users and map to Cloud Providers authentication
- Evaluate the capabilities of the OGF Data Format Description Language (DFDL) for data exchanges

Editor's Note: the following is excerpted from the NIST CCSRWG Use Case Integration AHG. Summary of results and recommendations for multiple IaaS Cloud Deployments

To achieve portability and interoperability across deployments, it will be necessary to have standardization of the interfaces to IaaS Clouds (e.g. management APIs). There are several efforts underway in this area including OGF OCCI and DMTF's Cloud Management Working Group. Unfortunately neither of these initiatives has yet produced a standard that will be implemented by the government's IaaS suppliers.

The government should consider alternatives to deal with this IaaS interface standards maturity gap. There have been several suggestions including accelerating the standards activities, having suppliers address future integration problems, and creating government Cloud brokering layers. To choose the best approach, it is recommended that the government set up an IaaS Interfaces Study Group with members from government, SDOs, and industry. The final report from this Group should provide guidelines that will support efficient Cloud deployments while avoiding costly future system integration problems (e.g. from vendor lock-ins).

One of the key issues for a Cloud Standards Roadmap is recommendations for deployments while waiting for standard to mature. This issue is especially significant for large enterprises

(e.g. US Government) that will deploy multiple Cloud applications often interfaced to existing systems.

In the current state of enterprise Cloud Computing, standards are immature and eventual dominant products have not been decided. The emphasis at this time should be on maintaining flexibility. The general principle is to design deployments to minimize the impact of change without significantly compromising functionality. There are key Cloud architecture design decisions where this principle can be use to avoid vendor lock-in.

For the "Cloud First" Business Use Case of multiple IaaS deployments, the use of tightly coupled non-standard interfaces from enterprise to Clouds can increase the effort required for future portability. Unfortunately mature standard IaaS interfaces are not available at this time. This will increase costs when applications have to be migrated across Clouds in the future. There are several possible ways to address this problem. Note that different aspects of the interfaces (e.g. APIs, VM movement, Federated Identity, data exchange formats) can be handled separately as standards become available. The alternative approaches can be grouped into three categories.

* Not require any constraints on IaaS interfaces now or in the future
  - Let Cloud Providers optimize their current IaaS interfaces
  - Let agencies choose the approved Cloud Provider that best meets their needs
  - Use system integrators or Cloud Providers to implement migrations when necessary

* Use formal standards in the future with no current IaaS interface constraints
  - Accelerate the creation and adoption of standards
  - Require Cloud Providers to natively support standards when they become available
  - Require Cloud Providers to create adaptors to support standards when available

* Use government IaaS interface specifications now while moving to future formal standards
  - Develop consensus government specifications that can migrate to future formal standards

- Require Cloud Providers to build adaptors to support government specifications
- Create a modular distributed Cloud Brokering layer to support government specifications

Selecting the correct alternatives will be critical to ensure the cost-effective future government Cloud deployment architectures. These decisions will require a structured process based on empirical data and evaluations.

Recommendation: The Cloud Computing Standards Roadmap Working Group should provide specific recommendations related to interface standards for IaaS Cloud deployments. The recommendations can include the use of existing standards, emerging standards that need to be evaluated, areas where standards should be accelerated, tactical solutions while awaiting for standards to mature, and potential risks in proprietary interfaces. The key question is how to create robust, flexible architectures while deploying applications on diverse IaaS Clouds.

Annex A – Definitions

Standard is a 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 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]

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)]

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)]

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)]

Editor's Note: Privacy is a complex multi-dimensional idea. Suggest that we keep our coverage to data protection or informational privacy in the context of information security. This will be a narrow treatment of privacy. Privacy does not equate to confidentiality.

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)]

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]

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.]

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]

Annex B –Acronyms

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)

Cloud Computing Forum, Korea (CCF)

CloudAudit

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.

- **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.

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. 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.

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

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 Grid Forum (OGF)

The OGF is an open community committed to driving the rapid evolution and adoption of applied distributed computing. Applied Distributed Computing is critical to developing new, innovative and scalable applications and infrastructures that are essential to productivity in the enterprise and within the science community. OGF accomplishes its work through open forums that build the community, explore trends, share best practices and consolidate these best practices into standards.

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

[i] 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