International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

FG Cloud TR

Version 1.0
(02/2012)

Focus Group on Cloud Computing
Technical Report

Part 4: Cloud Resource Management Gap Analysis

ITU-T

FOREWORD

The procedures for establishment of focus groups are defined in Recommendation ITU-T A.7. The ITU-T Focus Group on Cloud Computing (FG Cloud) was established further to ITU-T TSAG agreement at its meeting in Geneva, 8-11 February 2010, followed by ITU-T study group and membership consultation.

Even though focus groups have a parent organization, they are organized independently from the usual operating procedures of the ITU, and are financially independent. Texts approved by focus groups (including Technical Reports) do not have the same status as ITU-T Recommendations.

INTELLECTUAL PROPERTY RIGHTS

The ITU draws attention to the possibility that the practice or implementation of this Technical Report may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU-T Focus Group participants or others outside of the Technical Report development process.

Table of Contents

1. Scope

This Technical Report provides an analysis of major standards gaps in cloud resource management. The gaps are based on an analysis of SDOs that are involved in resource management and what remains to be accomplished. The intent is to identify the resource management standards gaps that exist, rather than solve the gaps that are identified.

2. References

None.

3. Terms and definitions

3.1 Terms defined elsewhere

3.1.1 Terms defined in ecosystem deliverable

Resource: Any kind of resource to be shared to compose cloud services, including computing power, storage, network, database, and applications.

3.1.2 Terms defined in DMTF (DSP0243)

Virtual hardware: The hardware (including the CPU, controllers, Ethernet devices, and disks) that is seen by the guest software.

Virtual machine: The complete environment that supports the execution of guest software. A virtual machine is a full encapsulation of the virtual hardware, virtual disks, and the metadata associated with it. Virtual machines allow multiplexing of the underlying physical machine through a software layer called a hypervisor.

3.2 Terms defined in this Technical Report

Business support system (BSS): A system dealing with customers and supporting processes such as taking orders, SLA, processing bills, and collecting payments.

Operational Support System (OSS): Computer systems used by telecommunications service providers that deal with the telecom network itself and support processes such as maintaining network inventory, provisioning services, configuring network components, monitoring, auditing, logging, and managing faults.

Resource management: An efficient and effective way to access, control, manage, deploy, schedule, and bind resources when they are provided by resource providers and requested by consumers.

Cloud resource audit: The process of examining and verifying cloud resource records, and supporting records.

4. Abbreviations and acronyms

This Technical Report uses the following abbreviations and acronyms.

CDMI Cloud Data Management Interface

CIM Common Information Model

CSP Cloud Service Provider

OCCI Open Cloud Computing Interface

OVF Open Virtualization Format

QoS Quality of Service

VM Virtual Machine

5. Overview of activities on SDOs for cloud resource management

The following is a list of activities on SDOs for cloud resource management.

5.1 DMTF

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. One of the key standards it maintained is the common information model (CIM).

5.1.1 CIM system virtualization model (DSP2013)

The DMTF common information model (CIM) is a conceptual information model for describing computing and business entities in Internet, enterprise, and service-provider environments. CIM uses object-oriented techniques to provide a consistent definition of, and structure for, data. The CIM schema establishes a common conceptual framework that describes the managed environment. CIM defines a conceptual model that describes how managed resources (for example, computers or storage area networks) may be represented as a set of objects and relationships among them. The model is described using UML and is also available in XML Schema and CIM-XML formats.

The open virtualization format (OVF) specification was prepared by the System Virtualization, Partitioning, and Clustering Working Group of the DMTF.

5.1.2 OVF-Open virtualization format (DSP0243)

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

The key properties of the format are as follows:

5.1.2.1 Optimized for distribution

OVF supports content verification and integrity checking based on industry-standard public key infrastructure, and it provides a basic scheme for management of software licensing.

5.1.2.2 Optimized for a simple, automated user experience

OVF supports validation of the entire package and each virtual machine or metadata component of the OVF during the installation phases of the virtual machine (VM) lifecycle management process. It also packages with the package-relevant user-readable descriptive information that a virtualization platform can use to streamline the installation experience.

5.1.2.3 Supports both single VM and multiple-VM configurations

OVF supports both standard single VM packages and packages containing complex, multi-tier services consisting of multiple interdependent VMs.

5.1.2.4 Portable VM packaging

OVF is virtualization platform neutral, while also enabling platform-specific enhancements to be captured. It supports the full range of virtual hard disk formats used for hypervisors today, and it is

extensible, which allows it to accommodate formats that may arise in the future. Virtual machine properties are captured concisely and accurately.

5.1.2.5 Vendor and platform independent

OVF does not rely on the use of a specific host platform, virtualization platform, or guest operating system.

5.1.2.6 Extensible

OVF is immediately useful and extensible. It is designed to be extended as the industry moves forward with virtual appliance technology. It also supports and permits the encoding of vendor-specific metadata to support specific vertical markets.

5.1.2.7 Localizable

OVF supports user-visible descriptions in multiple locales, and it supports localization of the interactive processes during installation of an appliance. This capability allows a single packaged appliance to serve multiple market opportunities.

5.1.3 DMTF Open cloud standards incubator

This incubator was started in 2009. The goal of the incubator is to define a set of architectural semantics that unify the interoperable management of enterprise and cloud computing. In July 2010, the incubator delivered two important documents: *Use Cases and Interactions for Managing Clouds* (DSP-IS0103) and *Architecture for Managing Clouds* (DSP-IS0102).

These two documents together describe how standardized interfaces and data formats can be used to manage clouds. The first document focuses on the overall architecture, including requirements for the architected interfaces in general (e.g., requirements on resource model). The second document focuses on interactions and data formats. The use cases involved resources include service resources provision, changing and monitoring etc.

5.1.4 The Cloud Management Working Group (CMWG)

This Working Group evolved from DMTF Open Cloud Standards Incubator .The scope of CMWG will mainly be on cloud resource management aspects of Infrastructure as a Service (IaaS) and will include policy, SLAs, QoS, provisioning, and monitoring considerations. Current work items cover protocol requirements, data models, and use cases.

5.1.5 Cloud Auditing Data Federation Working Group (CADF)

The Cloud Auditing Data Federation WG is developing specifications for audit event data and interface models and a compatible interaction model that will describe interactions between IT resources for cloud deployment models. This working group's intent is to accept member use cases as input to profile development in order to make the data and interface models specified by the working group consumable by different customer scenarios and implementations.

5.2 OGF

The Open Cloud Computing Interface Working Group (OCCI-WG) of the OGF was established in 2009. The purpose of this group is the creation of a practical solution to interface with cloud infrastructures exposed as a service (IaaS). It focuses on the creation of an API for interfacing "IaaS" cloud computing facilities, which is sufficiently complete as to allow for the creation of interoperable implementations.

The document *Open Cloud Computing Interface - Use cases and requirements for a Cloud API* (GFD-I.162) is an informal description of use cases and requirements for the OOCI Cloud API.

This document records the needs of IaaS cloud computing managers and administrators in the form of use cases. The use cases serve as the primary guide for the development of API requirements (e.g., manage cloud resources from a centralized dashboard).

5.3 Storage Networking Industry Association (SNIA)

The common goal of the Storage Networking Industry Association (SNIA) is to promote acceptance, deployment, and confidence in storage-related architectures, systems, services, and technologies, across IT and business communities.

5.3.1 Storage Management Initiative Specification (SMI-S)

This Technical Specification defines an interface for the secure, extensible, and interoperable management of a distributed and heterogeneous storage system. This interface uses an object-oriented, XML-based, messaging-based protocol designed to support the specific requirements of managing devices and subsystems in this storage environment. This standard includes following parts: common architecture, common profiles, block devices, file systems, fabric, host elements, and media libraries.

5.3.2 Cloud data management interface (CDMI)

This specification defines an interface for interoperable transfer and management of data in a cloud storage environment. This interface provides the means to access cloud storage and to manage the data stored there.

5.4 ITU-T SG13/Q4

SG13/Q4 is currently developing a draft Recommendation on resource control and management for virtual networks for cloud services (Y.VNC).

The scope of this Recommendation is to address resource control and management issues in virtual networks for cloud services (VNCs) which represent a network aspect of cloud service infrastructure (e.g., data centres). This Recommendation will address the following topics:

- To cope with highly-frequent network changes in the VNC environment;
- To support various customer policies (access control, logging, QoS, etc.);
- To enable self-configuration and automated network bootstrapping;
- To support the hierarchical service-offering model among customers;
- To support logically isolated resources for VNCs;
- To manage large scale numbers of resources that may be federated among cloud providers.

To address such issues, this draft Recommendation specifies functional requirements and architectural frameworks which allow for resource control and management capabilities for VNCs (in progress).

5.5 GICTF

The common goal of the GICTF is to develop an inter-cloud world, in which multiple cloud systems running with different policies, interwork with each other to share resources so that SLAs to users can be maintained even in the event of large fluctuations in the computing load that cannot be handled by a single cloud system. In order to promote interworking between different cloud systems, GICTF wants to establish procedures for resource matching, monitoring, provisioning, discovering and management between cloud systems.

5.6 Cloud Computing Interoperability Forum (CCIF)

The Cloud Computing Interoperability Forum (CCIF) was formed for the purposes of wider industry adoption of cloud-computing technology and related services. CCIF is an open, vendor neutral, not-for-profit community of technology advocates, and consumers dedicated to driving the rapid adoption of global cloud computing services. CCIF shall accomplish this by working with open forums focused on building community consensus, exploring emerging trends, and advocating best practices / reference architectures for the purposes of standardized cloud computing. One output of this forum is a unified cloud interface (UCI) requirement which is an attempt to create an open and standardized cloud interface for the unification of cloud API. The unified cloud interface (UCI) is focusing on inter-cloud interoperability. This requirements document specifies the implementation of a semantic process that can broker access and represent multiple-cloud providers that are cloud-platform or cloud-infrastructure designs. The concept is to provide a single interface that can be used to retrieve a unified representation of all inter-cloud resources and to control these resources as needed. In this unified cloud interface the use of the resource description framework (RDF) is a method to describe a semantic cloud data model.

5.7 Cloud Computing Use Case Discussion Group

An important effort in the search for standardization comes from the Cloud Computing Use Case Discussion Group, which aims to discuss a way to produce open standards for cloud computing. Their major efforts are collaboration and coordination of efforts on the standardization, adoption of open standards wherever appropriate, and the development of standards based on customer requirements. The Group has produced a white paper highlighting the requirements that need to be standardized in a cloud resource to ensure interoperability in the most typical scenarios in cloud computing.

5.8 Summary and analysis

The following table provides an overview of the scope and range of cloud resource management standardization across the industry.

General requirements	Standardized aspects			SDOs
	Use cases and requirements	**Model and architecture**	**Interface**	
Physical and virtual RM	•*Use cases and interactions for managing clouds* Describes how standardized interfaces and data formats can be used to manage clouds and focusing on use cases, interactions, and data formats	•*Architecture for managing clouds* Describes the reference architecture as it relates to the interfaces between a cloud service provider and a cloud service consumer	•*Use cases and interactions for managing clouds (partially)*	DTMF
	•*Resource control and management for VNCs* Specifies functional requirements and architectural framework which allow for resource control and management capabilities for virtual networks for cloud services (VNCs)			ITU-T SG13/Q4
		•*SMI-S* Defines a cloud storage reference model	•*CDMI* Defines an interface for interoperable transfer and management of data in a cloud storage environment	SNIA
Portability		•*CIM system virtualization model* Defines a conceptual model that describes how managed resources may be represented as a set of objects and relationships among them	•*OVF* Defines an open and extensible format for the packaging and distribution of software	DTMF

			to be run in virtual machines	
Interoperability	•*OCCI* Describes use cases and requirements for the open cloud computing interface cloud (OCCI) API			OGF and SINA
	•*Use cases and functional requirements for inter-cloud computing* Establishes procedures for resource matching, monitoring, provisioning, discovering and management between cloud systems			GICTF
	•*UCI requirements* Creates an open and standardized cloud interface for the unification of cloud API		•*UCI requirements (partially)*	CCIF
	•*Cloud computing use cases* Highlights the capabilities and requirements that need to be standardized in a cloud resource to ensure interoperability in the most typical scenarios			CCUCG

In summing up the above SDO activities, it is noted that there are many SDOs which are developing interoperability standards for cloud resource management. Although some working groups have completed their work, others remain in the process of completing their work. Standards have been an important part of ensuring interoperability among different CSPs. Comprehensive interface standards are lacking for interoperability between cloud platforms built by different providers.

In addition, it is necessary to provide the ability of migration from one CSP to another one, including data. There is a lack of standardization in the portability management in the use case and requirements aspects, such as data storage, etc.

It is notable that many of the standards are being developed rather than completed. The ITU-T may contribute to the following important standardization activities with a view to reducing overlap and duplication:

- cloud resource interoperability, especially in the interface aspect;
- cloud resource portability, especially in the use case and requirement aspect.

6. Cloud resource management vision

One significant value of next generation networks will most likely be delivered via rapid design, development, deployment, and management of cloud services. With the adoption of the cloud computing platform, which realizes the delivery of SOA services, they will increasingly be delivered as composite apps, or mash-ups from multiple providers. Those providers that are able to deliver such composite cloud-based services rapidly and at competitive pricing, in a customizable fashion to tailor to various customer scenarios, will be the most successful in the future NGN clouds.

Cloud-based service delivery platforms need to increasingly deliver multi-platform, multi-cloud solutions to support the above scenarios. Such solutions will also need to be flexible and cost effective across multiple providers to support flexible value chains that will in turn support more open markets for cloud-based service components.

The above can be realized with the increasingly wide uptake of SOA-based services, delivered though cloud-computing platforms with commodity, re-usable services at competitive cost. This has been the ultimate promise of componentized, SOA-based offerings.

In order to realize this promise, the telecommunication industry needs to develop deep insight and understanding in both the run-time aspects of service delivery, as well as the management of these services.

Today's complex, media-rich, composite services use a variety of network infrastructures, both telecom and IT, and are composed of individual service elements that may be acquired from, or exposed to, third parties. To be profitable in this complex environment, service providers must standardize and automate service delivery processes across service life cycles so that services can be created and delivered as rapidly as the market demands.

Cloud computing work at ITU-T needs to aim at minimizing the cost and cycle time to translate service ideas into market offerings, reduce cost by repurposing existing content and applications, and adapt swiftly to market changes and customer preferences.

What is needed is a conceptual management structure enabling the delivery of next generation services — independent of the underlying software or network technologies used to implement those services. This service-delivery management structure addresses the full-services lifecycle, covering such important use cases as concept-to-cash, service marketplace, service composition or aggregation, and service catalogues.

ITU-T's cloud resource management effort needs to address the service-delivery management landscape and provide a reference model and reference architecture for the essential building blocks needed to manage the delivery of next-generation services. The reference model and architecture form the basis for detailed service delivery management requirements.

One of the objectives should be to provide a means to allow consistent end-to-end management and metering of services exposed by and across different cloud service providers' domains and technologies. The envisioned standards frameworks and best practices are needed to support business practices associated with multi-provider cooperation throughout the lifecycle of the service, and by means of lightweight design to foster wide adoption of the standard artefact in any architecture, technology environment, and service domain.

Avoiding siloed technologies and vendors' proprietary implementations to get consistent maintenance of services sourced from different domains is a challenging task. To address this challenge, ITU-T RM effort should put together an approach to enable and support consistent access to the software-enabled services. Such effort is needed to complement the functional capabilities exposed by the software component's API with additional lifecycle management operations. This should also enable reusability of services in different environments, especially the cloud.

The highest priority of the development of cloud resource management standards in ITU-T needs to be centred on the development of frameworks, architecture, design patterns, and best practices that help realize the above requirements for the telecommunication industry, as it adopts cloud computing. Such frameworks, architecture and design patterns predominantly assume API implementations using web services over SOAP or REST. The API interfaces of individual service components are not the primary focus, as the actual interfaces may vary across different implementations, vendor technologies and operator requirements. The need is for standard design principles and frameworks that allow for rapid development, deployment and management of composite, inter-cloud services by the telecommunication industry, at low cost.

Such standards are aimed at providing guidance to SOA architects and developers in the telecommunication industry who want to build service mash-ups that are manageable, end-to-end. New standards should include multi-cloud service creation, delivery, and end-to-end management scenarios. It has to meet telecommunication requirements, seamlessly using cloud ecosystem.

7. Cloud resource management capabilities

For a general overview of resource management capabilities, refer to Clause 10 of the Technical Report on *Requirements and framework architecture of cloud infrastructure*.

The following sub-clauses describe capabilities that are specifically related to inter-cloud aspects of resource management.

7.1 Resource management

This capability deals with the management of resource configurations in a secure way for each service across multiple cloud infrastructures. .

The general requirements for this capability are:

- It should be possible to describe resource information (e.g., resource type, resource status, etc.) in a standard manner, in order to be able to manage resources across multiple-cloud infrastructures.
- It should be possible to update a cloud-infrastructure's configuration information across multiple cloud infrastructures in synchronization with events (e.g., reserve or release of resources) involving multiple cloud infrastructures.

7.2 Resource and service status monitoring

This capability deals with the collection and monitoring of the various status attributes of cloud infrastructure resources (e.g., usage, performance, service quality etc.) residing in the interworked CSPs.

By monitoring status information about resource availability (e.g., dead/alive status of machines) or service-level performance degradation (e.g., delay or response time degradation), a CSP can initiate actions to maintain the service availability with the help of other CSPs.

The general requirements for this capability are:

- It should be possible to, periodically or on a request basis, collect information about the usage and performance status of the resources of the different CSPs' cloud infrastructures.
- It should be possible to, periodically or on a request basis, collect information about the availability (e.g., dead/alive status of machines) of the different CSPs' cloud infrastructures.
- It should be possible to exchange monitoring information, in commonly defined ways, across different CSPs' cloud infrastructures. In case of mutual monitoring among interworked CSPs, it should be possible to maintain the appropriate level of security.

7.3 Resource performance estimation and selection

This capability deals with the resource selection from those which are candidate and have already been reserved in other CSPs' cloud infrastructures. This capability estimates the achievable performance by available reserved resources and assists the selection of the resources to be effectively used from all the reserved resources.

The general requirements for resource planning are:

- It should be possible to estimate the achievable performance of available reserved resources (e.g., computing resources, storage resources, input/output capacity between storages, network bandwidth) in other CSPs' cloud infrastructures.

7.4 Resource discovery and reservation

This capability deals with search, discovery, and reservation of the available resources in other CSPs' cloud infrastructures. This capability deals also with reservation acknowledgement for the candidate resources which have been tentatively reserved in other CSPs' cloud infrastructures.

The general requirements for resource discovery and reservation are:

- It should be possible to search resources available in other CSPs' cloud infrastructures.
- It should be possible to reserve the discovered resources in other CSPs' cloud infrastructures.
- It should be possible to provisionally reserve the discovered resources, i.e. to keep the resources to be used (as candidates), for later acknowledgement (for some of them) or release (for others).
- It should be possible to search for resources based on different priorities (e.g., in a different order of searching).

NOTE: Quality requirements may vary from service to service and each resource contribution to the service quality may vary as well. For example, if latency is critical, it should be possible to firstly reserve servers that are near the user, and then network resources. In contrast, if bandwidth is critical, it should be possible to firstly reserve networks that can provide sufficient bandwidth, and then search for available servers that are connected to those networks.

It should be possible to reserve available resources based on different priorities (e.g., early recovery, required quality guarantee, service type, etc.).

7.5 Resource setup and service activation

This capability deals with the setup of the reserved resources in the remote interworked CSPs, and the activation of the middleware and applications for service provision over the remote CSPs. This includes connecting cloud infrastructures via networks, remotely activating (i.e. invoking) application or middleware, and transferring or copying data to enable the use of resources in other CSPs' cloud infrastructures.

The requirements for setup of reserved resources are:

- It should be possible to remotely set up reserved resources in the remote cloud infrastructure, and to access their configuration and policy settings from the requesting cloud infrastructure.

7.6 Alteration and reversion of the user access to the cloud service

This capability deals with alteration cloud service user access from the original cloud infrastructure to other cloud infrastructures to which the services may be delegated, in order to cope with a

disaster or degradation in service performance, and reversion of the destination to the previous one when the original cloud infrastructure becomes able to provide the services again.

The general requirements for alteration and reversion of cloud service user access are:

- It should be possible to alter the destination of cloud service user access to the substitute cloud infrastructure without any cloud service user's operations in order to allow cloud service users to use services similarly to the pre-disaster situation.
- It should be possible to alter cloud service user access without any cloud service user's operations, in case load distribution between interworking cloud infrastructures is no longer needed
- It should be possible to reverse the destination of cloud service user access to that of the original cloud infrastructure, when the affected original cloud infrastructure has recovered from a disaster, or when load distribution between interworking cloud infrastructures is no longer needed.

7.7 Releasing resources

This capability deals with release of resources reserved and used from other cloud infrastructures by judging that cloud interworking is no longer needed based on monitoring results, e.g., that disaster recovery has been completed or load distribution has been adopted.

The general requirements for release of reserved resources are:

- It should be possible to release other resources that were activated when the reserved resources began to be used, to update the remote-cloud infrastructure's configuration information, and to erase and/or transfer back the previously received data.

8. Gap analysis

Cloud resources management (RM) supports the management of virtual and physical computing, storage and network resources. Cloud RM is a core technical issue of cloud computing and is limiting the mainstream adoption of cloud. As a result, it is essential to analyse the RM requirements and identify future areas of study (i.e. standards gaps) that, if addressed, could significantly promote the development of cloud RM.

Based on an analysis of requirements and the activities of SDOs involved in cloud resource management, the following areas of standardization should be developed in ITU-T.

8.1 Cloud resource description and discovery

The move to cloud computing brings with it a number of special challenges when it comes to resource management. One particular area is that of describing cloud resources and how to discover them in a physically-distributed environment in which individuals now have a significant presence in IT systems that are outside of their organization. Consequently, as consumers transit from buying and managing their own resources to renting resources in a cloud environment, methods are needed to identify, and manage computing resource description and discovery. The important area is to

identify methods and best practices to use cloud systems to satisfy the need to deliver a telecom-centric usage model.

8.1.1 Study areas: Resource discovery examples

The following are four examples of resource discovery standards that could be adapted for Internet-scale cloud applications:

- KANTARA's ID-WSF which defines both a discovery information model and discovery services that cover federated identity and access management. LDAP is an existing standard that has been used to build catalogue and discovery services. (http://projectliberty.org/resource_center/specifications/liberty_alliance_id_wsf_2_0_specifications)
- Domain name system (DNS) is a hierarchical distributed naming system for computers, services, or any resource connected to the Internet or a private network. DNS associates various kinds of information with domain names assigned to each of the participating entities. Most importantly, it translates domain names meaningful to humans into the numerical identifiers associated with networking equipment for the purpose of locating and addressing these devices worldwide.
- X.discovery is being developed by ITU-T, Q10/17. This draft ITU-T Recommendation provides a framework for discovering identity-related information and the generic mechanisms that enable this. Thus, the framework covers how to discover identity-related information in this context.
- Cloud infrastructure management interface (CIMI) is being developed by DMTF Cloud Management WG. This draft provides a logical mode and REST over http interface to manage resources in IaaS. The mode and interface cover how to describe and operate resources in the cloud infrastructure framework.

Existing standards need to be evaluated for their suitability for cataloguing and discovery of cloud resources.

8.1.2 Differences

There are two principal differences within cloud computing that make the problem of managing resources associated with cloud services more difficult. One difference is the virtualization at the compute and network layers. The other difference is that multiple clouds and multiple enterprise domains are increasingly involved in the delivery of cloud services and this environment greatly complicates RM.

8.1.3 Virtualization

It is useful to look at the overall RM problem from the point of view of the lifecycle management of a cloud application. The application, as it passes through its lifecycle, must be acted upon by traditional business processes associated with administration, provisioning/configuration, service assurance, and charging/billing/settlement.

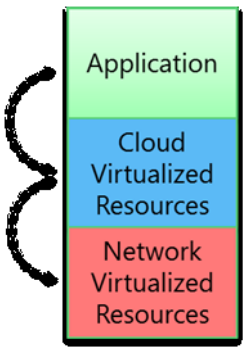

As shown in the adjacent figure, in the simpler case of an application that resides on a single-cloud infrastructure, it becomes dependent upon two distinct layers of virtualized resources. The dotted lines depict the active coordinated relationship that must be maintained between resources at each layer.

An RM issue requiring further work is how to use existing cloud management systems to maintain awareness of which logical and physical resources are actually relevant to a specific instance of a specific application at any given point in time.

Although the elastic cloud infrastructure provided by IaaS and PaaS can configure additional resources to handle changing application demands, there are additional requirements needing further analysis for dynamically reconfiguring the underlying network configurations in response to the changing resources at the cloud layer. This issue arises both within the internal network fabric of large cloud data centres, between two clouds and the interconnecting networks in hybrid scenarios, and externally across transport networks and content delivery networks.

Another issue that arises is the division of responsibility between an internal cloud virtualization management layer (IaaS and PaaS) and an external OSS. Although the cloud virtualization layer can typically manage its own physical and logical resource allocations for supported applications, an external OSS may be required to dynamically reallocate resources in a coordinated fashion across all three layers, or to track and have knowledge of those changing relationships.

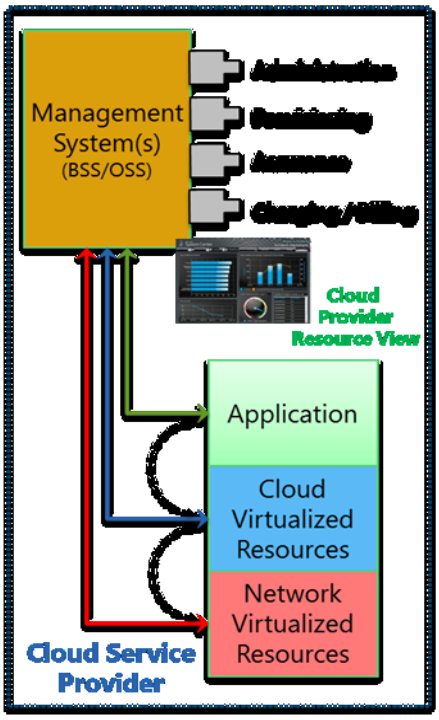

The capability of an OSS to both manage resource allocations and track their instantaneous state could enable the OSS to provide the information necessary to display a dashboard of the health and welfare of a given service and all of the underlying relevant resources, at any given point in time. From a resource-management quality of service point of view, the issue is how to ensure that the service assurance systems are receiving relevant telemetry from the cloud or network resources actually involved in delivering a particular instance of a service. The issue is less to do with what telemetry data need to be managed, as each dataset is often unique to a given telecom cloud implementation of OSS, but how to use the cloud infrastructure to do so effectively and economically.

8.1.4 Inter-cloud resource management

Up until now the discussion has been about the managing of resources across the application, cloud, and network resource stack within one logical cloud. Actual cloud service delivery scenarios typically involve coordination across inter-clouds residing in different enterprise domains.

Through the process previously discussed scenario-specific service/resource management interfaces need to be able to manage the relevant underlying resources in a coordinated manner that is

effectively transparent to whatever external systems are interacting with service/resource management interfaces. In the adjacent figure, a BSS/OSS enterprise architecture is depicted providing the needed management interfaces (again, the interfaces themselves are not at issue, as each implementation may have fine-tuned its own). The best practices should provide the flexibility for the cloud application itself to expose its service/resource management interfaces in addition to enabling a BSS or OSS to expose one or more of the interfaces because that OSS is tracking the dynamic changes in the underlying resources allocated to support the cloud application being managed.

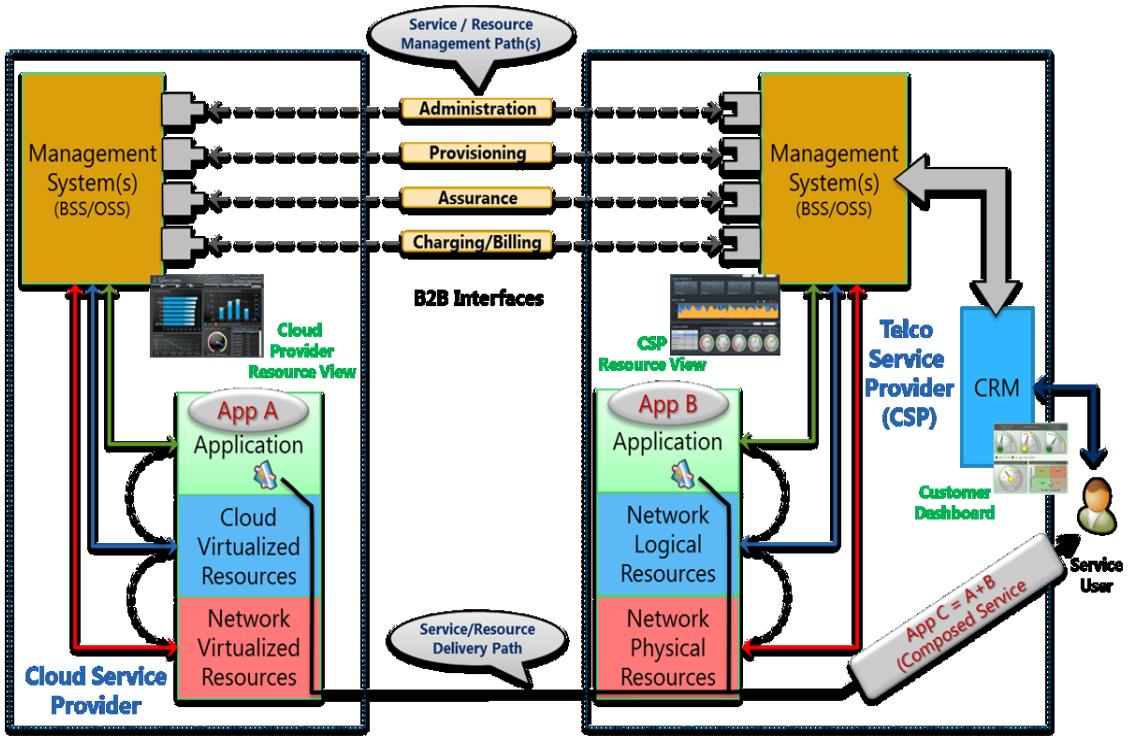

The figure above illustrates the end-to-end resource management requirements in an inter-cloud, multi-enterprise domain scenario. Given a useful way to expose customized resource management interfaces in a single cloud resource stack, the focus of the additional analysis can shift to enabling end-to-end management of composed services and their underlying dynamically changing resources.

The ultimate test of whether the standards are adequate is whether the three dashboards depicted can accurately display the actual status of services via metrics retrieved from the underlying relevant resources across an inter-cloud environment. It is important that the RM issues address the problem holistically, taking into consideration the service-creation lifecycle-management processes from the point of view of a product manager and service developer, as well as the order-to-cash processes associated with a service user.

8.2 Security consideration for resource management

Security requirements for cloud-resource management include user authentication, identity and access management, data protection, multi-tenancy resource isolation, monitoring and auditing for

compliance, incident response, user and customer privacy, and the underlying portability and interoperability of security components.

8.2.1 Cryptography

The cloud service provider should allow its customers to control and manage the cryptographic materials and methods used to protect its applications and data, whether in motion or at rest, when entering, exiting, or residing within the cloud service provider's infrastructure. The minimum functionality would allow an administrator to enter this information manually, with more functionality allowing secure synchronization. In the case of the enterprise user, it is necessary for the enterprise to synchronize identity information (of internal users) with the cloud service provider in order to allow the users to access internal IT service and cloud services with the same identity.

8.2.2 Identity services

The cloud service provider should provide identity services that permit customers to provision and manage identities that can be authenticated and authorized to access cloud services using standardized mechanisms. These services should consider the need to provide identities, either manually or by automated means, and to synchronize or collaborate with external identity provider services.

Virtualized operating systems should support layered security controls and reduce dependency on the platform provider alone.

8.2.3 Resource audit

The cloud service provider should assure that all access or changes to resources, and the data, produce auditable records, and that these audit records can be compiled into a consistent audit trail that can correlate to end user- or service-initiated actions. Audit records should include clear indication of any delegations of identity or authorizations.

9. Future study areas

Based upon the above gap analysis, the following questions should be addressed in future areas of study:

- How can RM be accomplished holistically across all applicable business processes?
- How to build and maintain new, multi-cloud based OSS/BSS systems to dynamically reconfigure underlying network configurations in response to the changing resources at the cloud layer?
- How to address the division of responsibility between an internal cloud virtualization management layer (IaaS and PaaS) and an external enterprise architecture OSS?
- How to ensure that the service assurance systems are receiving relevant telemetry from the cloud or network resources actually involved in delivering a particular instance of a service, in a flexible and cost-effective way using new cloud-based implementations.?
- How to develop best practices, architectural guidelines and frameworks to further expose diverse, application defined service/resource management interfaces for each

service/resource in business-to-business interfaces, supporting heterogeneous, multi-vendor, multi-cloud environments?

- How to provide the flexibility for the cloud application to expose desired service/resource management interfaces in addition to enabling the enterprise architecture BSS or OSS?
- How to use the cloud computing environment to enable flexible, end-to-end management of composed services and their underlying dynamic changing resources?
- How to holistically take into consideration the service-creation lifecycle-management processes from the point of view of a product manager and service developer, as well as the order-to-cash processes associated with a service user when addressing RM issues?
- How to evaluate the security of resource provided by a cloud provider?
- How to audit that the security controls are implemented correctly, and operating as intended?
- How to assess whether the security controls are producing the desired outcome with respect to meeting the security requirements for the cloud system?

Annex I
Resource Management Requirement Summary

The following Table provides a summary of the cloud functional RM requirements that are described in the Draft Technical Report on *Requirements and framework architecture of cloud infrastructure*.

Clause	Title	Brief Description
10.1	General	
10.1.1	Physical & Virtual RM	Management support for both physical and virtual resources is needed
10.1.2	Heterogeneity shielding	Unified user interface for all physical or virtual resources
10.1.3	Dynamic shielding	Capability to evaluate the dynamic performance of resources
10.2	Functional	
10.2.1	Encapsulation	Heterogeneous resources are integrated in a manner that provides a unified interface for creating, locating ,deploying, provisioning, recovering and deleting resources and the attributes (status, capacity, execution, error, interrupt) can be measured and searched.
10.2.2	Orchestration and Provisioning	All resources should be elastic, dynamic, on-demand and provisioned based on service requirements that trigger resource actions to include real time monitoring of applications and SLAs
10.2.3	Assessment	Unified hardware (racks, servers, storage devices, network equipment, VMs) and software (hypervisors, OS's, middleware, DBs, applications) attributes should be automatically up-dated when the physical or virtual device is changed
10.2.4	Template	Capability to provide life cycle management of each template, including creation, publication, activation, renovation, deletion
10.2.5	Monitoring	Capability to provide multi-layers resource monitoring, including service instance monitoring, physical resources monitoring, resource pool monitoring, user connection monitoring, software monitoring etc
10.2.6	User Resource Environment Management	Capability to ensure the secure isolation between different user resource environments; to provide appropriate control of user resource environment
10.3	Management model	
10.3.1	Resource allocation	Capability to allocate recourses, including resource modelling and description, resource offering and treatment, resource discovery and monitoring, and resource selection
10.3.2	Access and control	Capability to provide an efficient and effective way to access, control, manage, deploy, schedule and bind resources when they are provided by resource providers and requested by resource clients

Bibliography

[1] US Government Cloud Computing Technology Roadmap Volume II Release 1.0 (Draft),

[2] Useful Information for Cloud Adopters, NIST Special Publication 500-293 (Draft), November, 2011

[3] TM Forum, The Enhanced Cloud Service Management Catalyst, (http://www.tmforum.org/EnhancedCloudService/10353/home.html)

[4] IETF Draft: Virtual Resource Management in Cloud, July 11, 2011 (http://www.ietf.org/id/draft-junsheng-opsawg-virtual-resource-management-00.txt)

Printed in Switzerland
Geneva, 2012