International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.3520

(09/2015)

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

Cloud Computing

Cloud computing framework for end to end resource management

Recommendation ITU-T Y.3520

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
IPTV over NGN	Y.1900–Y.1999
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Enhancements to NGN	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Packet-based Networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999
FUTURE NETWORKS	Y.3000–Y.3499
CLOUD COMPUTING	**Y.3500–Y.3999**

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T Y.3520

Cloud computing framework for end to end resource management

Summary

Recommendation ITU-T Y.3520 presents general concepts of end to end resource management in cloud computing; a vision for adoption of cloud resource management in a telecommunication-rich environment; and multi-cloud, end to end resource management for cloud services, i.e., management of any hardware and software used in support of the delivery of cloud services.

History

Edition	Recommendation	Approval	Study Group	Unique ID[*]
1.0	ITU-T Y.3520	2013-06-22	13	11.1002/1000/11919
2.0	ITU-T Y.3520	2015-09-29	13	11.1002/1000/12585

Keywords

Cloud computing, cloud service, framework, requirement, resource management.

[*] To access the Recommendation, type the URL http://handle.itu.int/ in the address field of your web browser, followed by the Recommendation's unique ID. For example, http://handle.itu.int/11.1002/1000/11830-en.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation 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 members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at http://www.itu.int/ITU-T/ipr/.

Table of Contents

Recommendation ITU-T Y.3520

Cloud computing framework for end to end resource management

1 Scope

This revised Recommendation provides a framework for end to end resource management in cloud computing. It includes:

– general concepts of resource management for end to end cloud computing resource management;

– a vision for adoption of resource management for cloud computing in a telecommunication-rich environment;

– multi-cloud, end to end management of cloud computing resources and services, e.g., management of any hardware and software used in support of the delivery of cloud services.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T X.1601]	Recommendation ITU-T X.1601 (2014), *Security framework for cloud computing*.
[ITU-T Y.3500]	Recommendation ITU-T Y.3500 (2014), *Information technology – Cloud computing – Overview and vocabulary*.
[ITU-T Y.3501]	Recommendation ITU-T Y.3501 (2013), *Cloud computing framework and high-level requirements*.
[ITU-T Y.3502]	Recommendation ITU-T Y.3502 (2014), *Information technology – Cloud computing – Reference architecture*.
[ITU-T Y.3511]	Recommendation ITU-T Y.3511 (2014), *Framework of inter-cloud computing*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 cloud computing [ITU-T Y.3500]: Paradigm for enabling network access to a scalable and elastic pool of shareable physical or virtual resources with self-service provisioning and administration on-demand.

NOTE – Examples of resources include servers, operating systems, networks, software, applications and storage equipment.

3.1.2 cloud deployment model [ITU-T Y.3500]: Way in which cloud computing can be organized based on the control and sharing of physical or virtual resources.

NOTE – The cloud deployment models include community cloud, hybrid cloud, private cloud and public cloud.

3.1.3 cloud service [ITU-T Y.3500]: One or more capabilities offered via cloud computing invoked using a defined interface.

3.1.4 cloud service category [ITU-T Y.3500]: Group of cloud services that possess some common set of qualities.

3.1.5 cloud service customer [ITU-T Y.3500]: Party which is in a business relationship for the purpose of using cloud services.

NOTE – A business relationship does not necessarily imply financial agreements.

3.1.6 cloud service provider [ITU-T Y.3500]: Party which makes cloud services available.

3.1.7 cloud service user [ITU-T Y.3500]: Natural person, or entity acting on their behalf, associated with a cloud service customer that uses cloud services.

NOTE – Examples of such entities include devices and applications.

3.1.8 emergency telecommunications (ET) [b-ITU-T Y.2205]: ET means any emergency related service that requires special handling from the NGN relative to other services. This includes government authorized emergency services and public safety services.

3.1.9 emergency telecommunication service (ETS) [b-ITU-T E.107]: A national service providing priority telecommunications to the ETS authorized users in times of disaster and emergencies.

3.1.10 inter-cloud computing [ITU-T Y.3511]: The paradigm for enabling the interworking between two or more cloud service providers.

3.1.11 management system [b-ITU-T M.60]: A system with the capability and authority to exercise control over and/or collect management information from another system.

3.1.12 service level agreement [ITU-T Y.3500]: Documented agreement between the service provider and customer that identifies services and service targets.

NOTE 1 – A service level agreement can also be established between the service provider and a supplier, an internal group or a customer acting as a supplier.

NOTE 2 – A service level agreement can be included in a contract or another type of documented agreement.

3.2 Terms defined in this Recommendation

This Recommendation defines the following term:

3.2.1 resource management: The most efficient and effective way to access, control, manage, deploy, schedule and bind resources when they are provided by service providers and requested by customers.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

3G	Third Generation
4G	Fourth Generation
BSS	Business Support System
CDN	Content Delivery Network
CRM	Customer Relationship Management
CSC	Cloud Service Customer
CSP	Cloud Service Provider
ET	Emergency Telecommunications

ETS	Emergency Telecommunication Service
FI	Functional Interface
IP	Internet Protocol
IT	Information Technology
LAN	Local Area Network
LTE	Long Term Evolution
MPLS	Multi-Protocol Label Switching
NGN	Next Generation Network
OAM	Operations, Administration and Maintenance
OSS	Operations Support System
PaaS	Platform as a Service
PHP	Hypertext Pre-processor
QoS	Quality of Service
SES	Software Enabled Services
SLA	Service Level Agreement
SMI	Service Management Interface
SNMP	Simple Network Management Protocol
VM	Virtual Machine
VoIP	Voice over IP
WAN	Wide Area Network
WiFi	Wireless Fidelity

5 Conventions

In this Recommendation:

The keywords "**is required to**" indicate a requirement which must be strictly followed and from which no deviation is permitted if conformance to this document is to be claimed.

The keywords "**is recommended**" indicate a requirement which is recommended but which is not absolutely required. Thus this requirement need not be present to claim conformance.

In the body of this Recommendation and its appendices, the words shall, shall not, should, and may sometimes appear, in which case they are to be interpreted, respectively, as is required to, is prohibited from, is recommended, and can optionally. The appearance of such phrases or keywords in an appendix or in material explicitly marked as informative are to be interpreted as having no normative intent.

6 End to end cloud resource management overview

The following clauses provide an overview of the general concepts of end to end cloud computing resource management in a telecommunication rich environment.

6.1 Introduction

One significant value of cloud service providers will most likely be the rapid design, development, deployment and management of cloud services. With the adoption of cloud computing service delivery capabilities, multiple service providers will provide more cloud services as composite or mash-up services. Service providers will increasingly have as their objective the rapid delivery of more customized, composite cloud-based services tailored to various customer scenarios [b-FGCC Part 4].

In this Recommendation, the term multi-cloud refers to usage scenarios involving the use of various cloud services implemented by more than one cloud service provider (CSP), though this multiplicity of CSPs may not be visible to the cloud service customer (CSC). This is not to be confused with the multi-platform cloud computing environment, which is a characteristic of cloud service providers that have chosen to offer a variety of programming and runtime execution facilities to assist in the development and execution of cloud applications. Nor should it be confused with the term "inter-cloud" which refers to the relationship and interconnection between CSPs and not to the overall end to end system.

Cloud applications (also known as cloud workloads) are applications (i.e., software programs designed for a specific purpose) that require execution in the cloud service provider's data centres in order for cloud services to be instantiated and become available for use by the cloud service users. In other words, a cloud application needs to be executed to make one or more cloud services available.

Cloud service providers need to increasingly offer multi-cloud platform solutions to support the above scenarios. Such solutions will need to be flexible and effective in managing resources across multiple cloud service providers [ITU-T Y.3501].

These solutions can be realized using cloud services, delivered through cloud computing capabilities with reusable services. Cloud service providers need to develop a deep insight into, and understanding of, the run-time aspects of service delivery as well as the management of these services and the resources required to deliver them.

Therefore, there is a need for a common concept for end to end resource management across multiple cloud service providers.

Complex, media-rich, composite services use a variety of both telecommunication and information technology (IT) infrastructures and are composed of individual service components that may be acquired from, or exposed to, third parties.

6.2 Service delivery management structure

The framework described in this Recommendation can be used to enable the delivery of cloud services, independent of the underlying software or network technologies. This framework, which is a service delivery management structure, needs to address the full cloud services lifecycle, covering such important use cases as service composition, aggregation and service catalogues.

Management of cloud services needs to provide a framework for the essential building blocks required to manage the delivery of cloud services and foster the basis for detailed service delivery management.

One objective is to provide a means to allow consistent end to end management, including accounting, of services exposed by and across, domains and platforms of different cloud service providers. A standard framework and best practices are needed to support business practices associated with multiple provider cooperation throughout the lifecycle of the service and to foster wide adoption of the standard artefacts in any architecture, technology environment and service domain.

Achieving consistent maintenance of cloud services sourced from different domains is a challenging task. To address this challenge, an approach that enables and supports consistent management access to the cloud services is desired. Such an approach is desired to complement the service capabilities exposed by the software component's interfaces with additional lifecycle management operations. This approach should also enable reusability of services in different environments, especially in cloud computing.

Frameworks, architecture, design patterns and best practices are required to realize the above objectives for the cloud service providers. The interfaces of individual service components are not the primary focus as the actual interfaces may vary across different implementations, vendor technologies and operator requirements. Standard design principles and frameworks are required to allow for the rapid development, deployment and management of composite multi-cloud services provided by the telecommunication industry.

This provides a framework to guide architects and developers of cloud services regarding the end to end management of cloud computing resources.

6.3 Difference between cloud computing and the traditional form of computing

There are two principal differences between cloud computing and the traditional form of computing that make the problem of managing resources associated with cloud services more difficult. One difference is the virtualization of the computing and network resources in the cloud computing reference architecture [ITU-T Y.3502]. The other difference is that multiple cloud service provider domains are increasingly involved in the delivery of cloud services and this environment greatly complicates end to end resource management.

6.4 Resource management for a single cloud service provider

The overall resource management should be viewed 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 management system functions such as administration, provisioning, configuration, service assurance and charging.

As shown in Figure 1, in the simpler case of an application that resides on a single cloud computing system, it becomes dependent on two distinct categories of virtualized resources. The dotted arrows depict the active coordinated relationship that must be maintained between resources at each level.

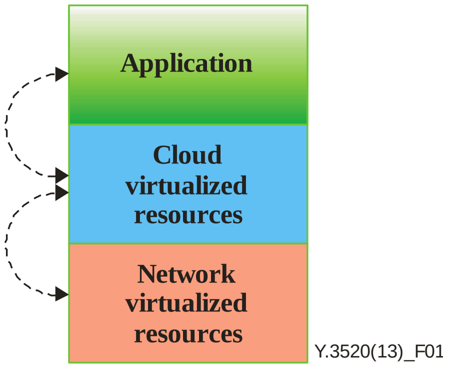

Figure 1 – Applications residing on a single cloud computing system

NOTE – Although Figure 1 divides virtualized resources into "cloud" and "network", cloud computing considers all resources at the same level [ITU-T Y.3502].

A resource management 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.

Due to the rapid elasticity and scalability characteristic of cloud computing [ITU-T Y.3500], the cloud computing system 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 various components of the cloud computing system. This issue arises both within the internal network fabric of large cloud computing data centres, between the interconnecting networks in hybrid scenarios, and across transport networks and content delivery networks.

Another issue that arises is the division of responsibility between an internal cloud computing virtualization management system and an external management system. Although the cloud computing virtualization function can typically manage its own physical and logical resource allocations for supported applications, an external management system may be desired to dynamically reallocate resources in a coordinated fashion across the three levels shown in Figure 1 or to track and have knowledge of those changing relationships.

As shown in Figure 2, the capability of a management system to both manage resource allocations and track their instantaneous state could enable that management system to provide the information necessary to display the status of a given service and all of the underlying relevant resources, at any given point in time.

From the point of view of the quality of service of resource management, the issue is how to ensure that the service assurance systems are receiving relevant telemetry from the cloud computing or network resources actually involved in delivering a particular instance of a service. The issue is less concerned with what telemetry data needs to be managed, as each dataset is often unique to a given management system implementation, but is more concerned with how to use the cloud computing system to do so effectively.

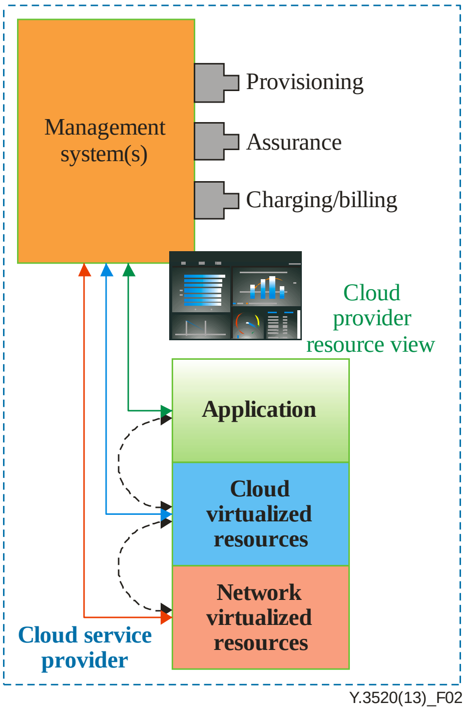

Figure 2 – Cloud resource management system (OSS and BSS)

6.4.1 Software enabled services

The software enabled services (SES) management approach enables both traditional service providers as well as Internet content and media service providers to leverage the opportunities and service marketplace that are presented by the convergence of networking and IT. Specifically, the SES management approach provides a means to allow consistent end to end management and metering of services exposed by and across different service providers' domains and technologies.

Operations, administration and management interfaces for cloud services today are structured in silos per technology, standardized by specific standards development organizations or implemented by vendors as proprietary implementations. This presents a challenge in the rendering of consistent management of services sourced from different domains.

The SES management approach proposes a mechanism to allow consistent access to the software components information as well as management operations. This consistent access is achieved by incorporating the management interface in addition to the functional interface (FI) definition that is part of software component creation. The SES approach enables reusability of services in different environments, including that of cloud computing by manipulating the SES lifecycle management metadata which is supporting the service management interface (SMI) operations.

For further information about SES and SMI concepts, please refer to Appendix III.

The SES pattern is defined to also handle those cases where the composed service is not able to manage all of the management dependencies by means of the logic which is triggered by the SMI operations. In this case, the lifecycle management metadata associated with the SMI is providing a recipe which describes how to manage the composition members.

Protocol neutral interface information models and class models of the SMI, along with corresponding statements of interface information requirements and interface information use cases, can then be defined [b-TMF TR198].

In order for implementation of SESs to be as useful as possible, the following requirements should be addressed:

- It is recommended that service design be as efficient as possible, requiring only the needed information for both input and output parameters, without being verbose. There should not be many arguments on the different management system operations.
- It is recommended that service design be simple, allowing for easy implementation in legacy as well as in new services. There should be no complex dependencies between the arguments of the different management information system operations.
- It is recommended that service implementation rely on industry standards in order to guarantee that it will be interoperable between different platforms.
- It is recommended that the SMI be extensible and generic to accommodate all SES scenarios.
- It is recommended that the SMI be easy to extend and that this interface be adapted to support additional management aspects of a specific domain or of a specific vendor.
- It is recommended that the SMI be agnostic to implementation, architecture, or business processes, to ensure adoption by many industry sectors.
- It is recommended that a non-"well designed service" be wrapped by a façade service in order to make it a "well designed service".

The SMI in the software enabled services reference architecture can be used to describe the management capabilities of SES. Examples of these capabilities are in the areas of invocation, provisioning, status, history, usage and health monitoring and associated alerts, management lifecycle state configuration as well as decommissioning of a given software enabled service [b-TMF TR198].

The SES management approach was designed to address the management of a single cloud service provider. In the next clause, it will be explained how the same concept can be applied to address a multiple cloud service provider scenario.

6.5 Resource management for multiple cloud service providers

Clause 6.4 describes the managing of resources for a single cloud service provider. However, cloud service delivery scenarios typically involve coordination across multiple cloud service providers residing in different domains.

Figure 3 illustrates the end to end management framework in a multi-cloud service provider domain scenario. Given the way in which customized management interfaces are exposed in a single cloud service provider implementation, the framework enables end to end management of composed services and their underlying dynamically changing resources.

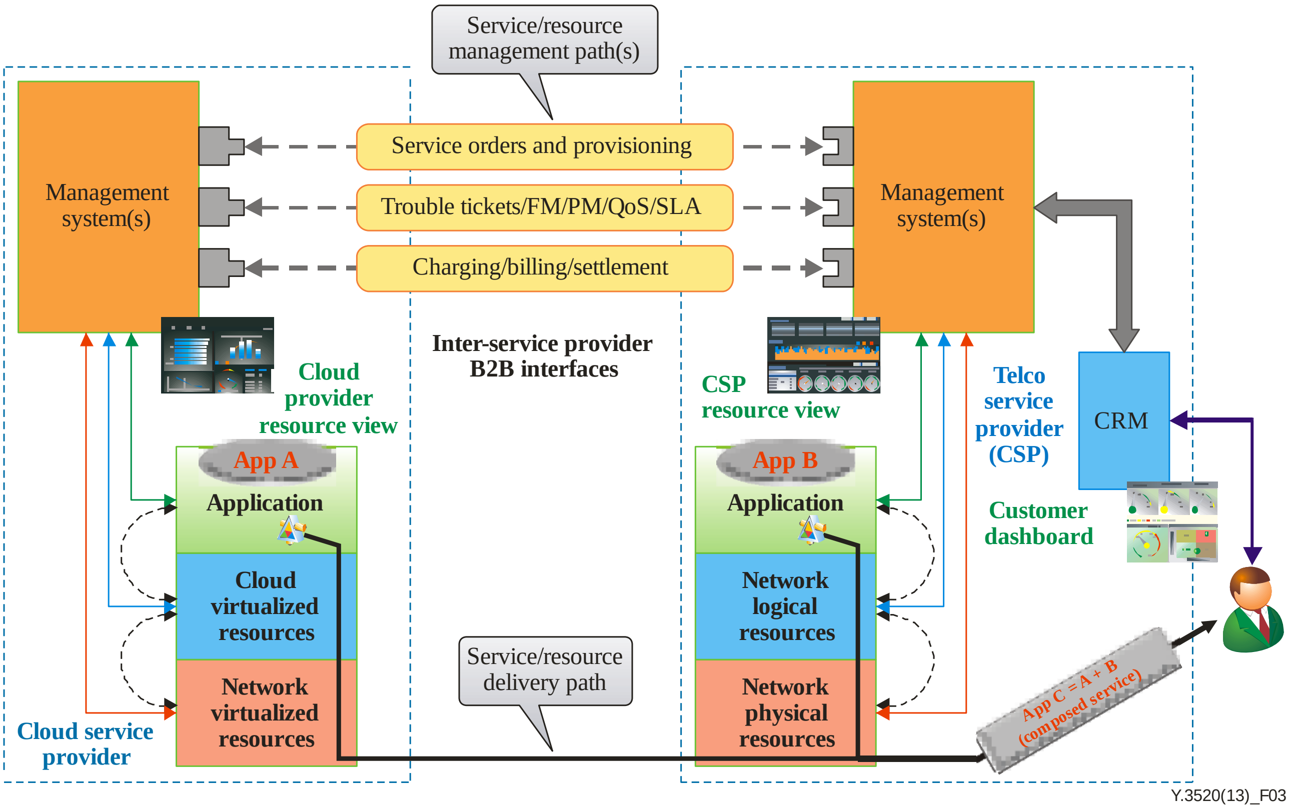

Service/resource management path(s)
Management system(s)
Service orders and provisioning
Trouble tickets/FM/PM/QoS/SLA
Charging/billing/settlement
Management system(s)
Cloud provider resource view
Inter-service provider B2B interfaces
CSP resource view
Telco service provider (CSP)
CRM
App A
Application
Cloud virtualized resources
Network virtualized resources
App B
Application
Network logical resources
Network physical resources
Customer dashboard
Service/resource delivery path
App C = A + B (composed service)
Cloud service provider
Y.3520(13)_F03

Figure 3 – End to end management expectations in a multi-cloud scenario

Similar to the case of a single cloud scenario, service and resource management interfaces need to be able to manage the relevant underlying resources in a coordinated manner that is effectively transparent to the external systems that are interacting with those management interfaces.

Figure 3 depicts a management system architecture 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 or resource management interfaces. In addition, they need to enable a management system to expose one or more of the interfaces so that the management system is tracking the dynamic changes in the underlying resources allocated to support the cloud application being managed.

The framework permits each CSP, as well as the CSC, to have accurate knowledge regarding the actual status of services via metrics retrieved from the underlying relevant resources across a multi-cloud environment. In other words, all three dashboards depicted in Figure 3 need to accurately display the status of the services. In addition, the framework should consider a comprehensive service lifecycle management process from the point of view of CSPs and CSCs, i.e., the stages needed from the time the CSC makes a request until the CSP receives compensation.

7 Requirements for the resource management involving multi-cloud service providers

7.1 High-level architecture for end to end multi-cloud resource management

Figure 4 shows a high-level architecture for end to end multi-cloud resource management. This architecture depicts the virtual machines containing a software stack consisting of middleware layers containing application servers hosted by the runtime environment of choice, on top of which cloud applications execute.

Figure 4 also shows both functional interfaces (FIs) and service management interfaces (SMIs) being exposed by various cloud applications running on multiple cloud data centres. The information can be consumed from various SMIs that are exported by multiple applications executing in multiple cloud data centres, allowing for a comprehensive end to end multi-cloud resource management and monitoring system to be realized.

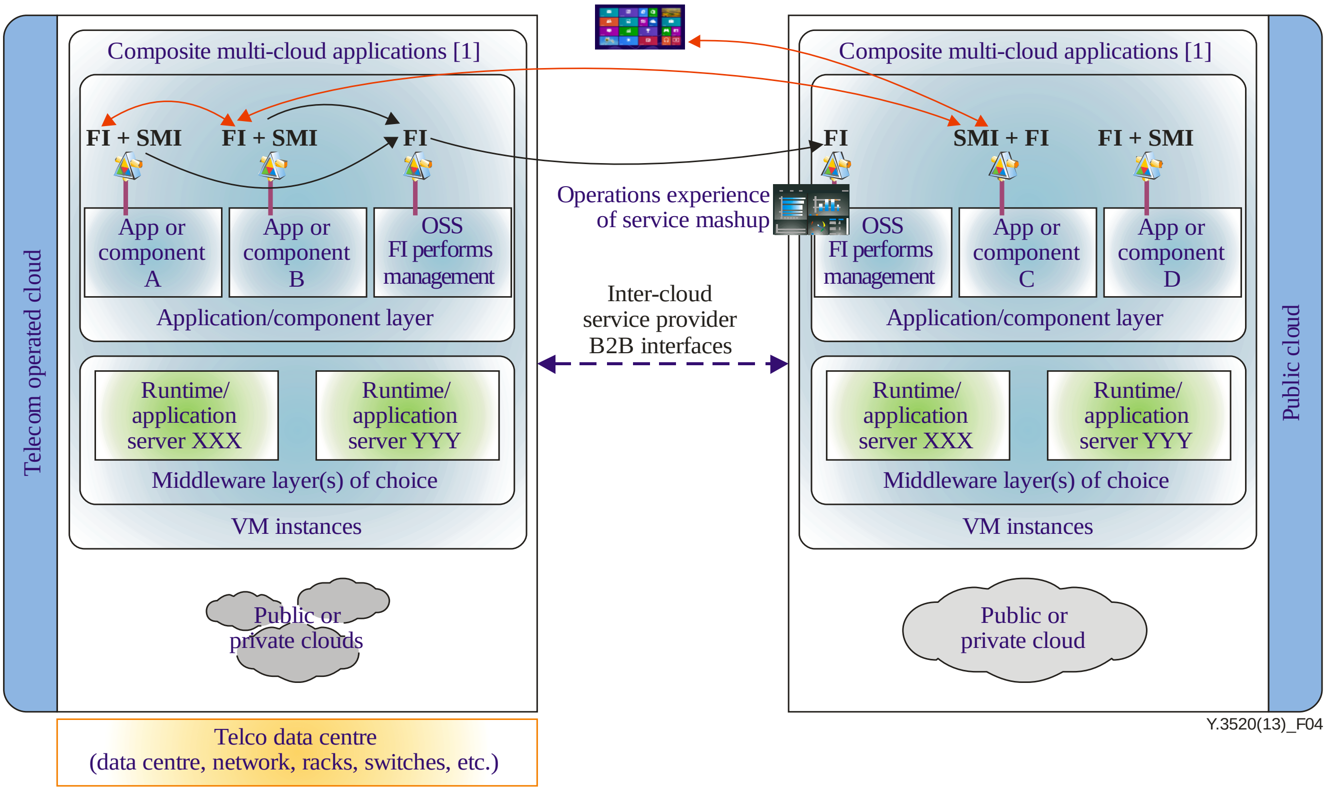
Telecom operated cloud
Composite multi-cloud applications [1]
FI + SMI
FI + SMI
FI
App or component A
App or component B
OSS FI performs management
Application/component layer
Runtime/ application server XXX
Runtime/ application server YYY
Middleware layer(s) of choice
VM instances
Public or private clouds
Telco data centre (data centre, network, racks, switches, etc.)
Operations experience of service mashup
Inter-cloud service provider B2B interfaces
Composite multi-cloud applications [1]
FI
SMI + FI
FI + SMI
OSS FI performs management
App or component C
App or component D
Application/component layer
Runtime/ application server XXX
Runtime/ application server YYY
Middleware layer(s) of choice
VM instances
Public or private cloud
Public cloud
Y.3520(13)_F04

NOTE – Composite multi-cloud applications can be written in a runtime and programming environment of choice, independent of the choice of cloud provider or cloud-deployment model. For example, use of Java, Node.js, PHP or .NET in both private and public clouds.

Figure 4 – Architectural vision for multi-cloud, multi-platform cloud management

In Figure 4, the applications executing in the virtual machines (VMs) could be a composite, distributed application built from various software components. A VM instance could contain all software components that belong to such an application, or only some of them in the case where the application is distributed and executing in more than one VM (hence the references to applications or components in Figure 4).

The architectural vision shown in Figure 4 enables interoperable applications to support cloud burst or hybrid cloud computing scenarios.

7.2 Functional requirements for end to end cloud resource management

To meet the high level architecture of end to end cloud resource management described in this Recommendation, a cloud computing platform should conform to the following requirements:

- It is required that the CSP supports the architectural and functional capabilities offered by the SES management approach in order to realize end to end cloud resource management.
- It is recommended that the cloud computing platform offers the cloud deployment model [ITU-T Y.3500] choice and workload portability across multiple CSPs in order to share workloads.
- It is recommended that the cloud computing platform provide the ability to support hybrid cloud applications, where the components of the cloud application run on various cloud data centres managed by different CSPs.
- It is recommended that cloud service provider, irrespective of the cloud deployment model they use, provide the support for multiple application frameworks, programming languages, tools and technology platforms, thereby lowering the potential for lock-in into specific solution or middleware technology.
- It is recommended that the cloud computing platform provides an architecture enabling telecommunications-grade capabilities including reliability, fail-over and monitoring inclusive of choice of middleware, programming language and runtime.
- It is recommended that the cloud computing platform supports workload portability and related management capabilities (e.g., control, operation and monitoring) amongst cloud service providers, supporting various cloud deployment models [ITU-T Y.3500], in a cost effective way.

8 Cloud resource management for emergency telecommunications

Emergency telecommunications (ET) [b-ITU-T Y.2205] are any emergency related service that requires special handling relative to other services (i.e., priority access for authorized users and priority treatment to emergency traffic).

While not always required, if the resources of the CSP are used to support the emergency telecommunications service (ETS) [b-ITU-T E.107], appropriate resource management functions will be needed to allow priority treatment in the use of the cloud computing resources by authorized users. The requirements in [b-ITU-T Y.1271] are relevant.

NOTE – Requirements in [b-ITU-T Y.1271] apply across multiple layers of the cloud computing reference architecture [ITU-T Y.3502].

9 Security considerations

The security framework from cloud computing [ITU-T X.1601], analyses security threats and challenges in the cloud computing environment and describes security capabilities that could mitigate these threats and address security challenges. The single cloud and multi-cloud resource management framework described in this Recommendation is based on the interconnections within a single cloud service provider or between two or more cloud computing systems operated by different service providers. Thus, secure interconnection within and across the systems should be considered. Protection of internal management system interfaces and information against unauthorized access, internally or by an external interconnected entity, should also be considered. Exposed internal and external management interfaces should also be security protected. It is recommended that the applicable X, Y and M series of ITU-T security Recommendations be taken into consideration, including access control, authentication, data confidentiality, communications security, data integrity, availability and privacy.

Appendix I

Comprehensive view of management layers

(This appendix does not form an integral part of this Recommendation.)

Figure I.1 is an attempt at describing the management layers and how the service management interface (SMIs) for each layer correlate with each other to offer a complete picture. The cloud computing data centre implementation layers are depicted by large circles parallel to each other. The SMIs are depicted by small blue circles containing the information required by the management system in order to achieve a holistic view of the entire operation. The straight lines between each plane show the flow of information and relationship between what is happening at each layer, depicting how each layer is related to and affected by, the neighbouring one. Looking at the diagram in its entirety helps the viewer to realize why it makes sense to expose consistent SMIs from each layer and to expose management information and telemetry in a consistent matter that can then be rolled up into a comprehensive diagnostics and management solution that can be used by a telecom operator offering and consuming cloud services and products.

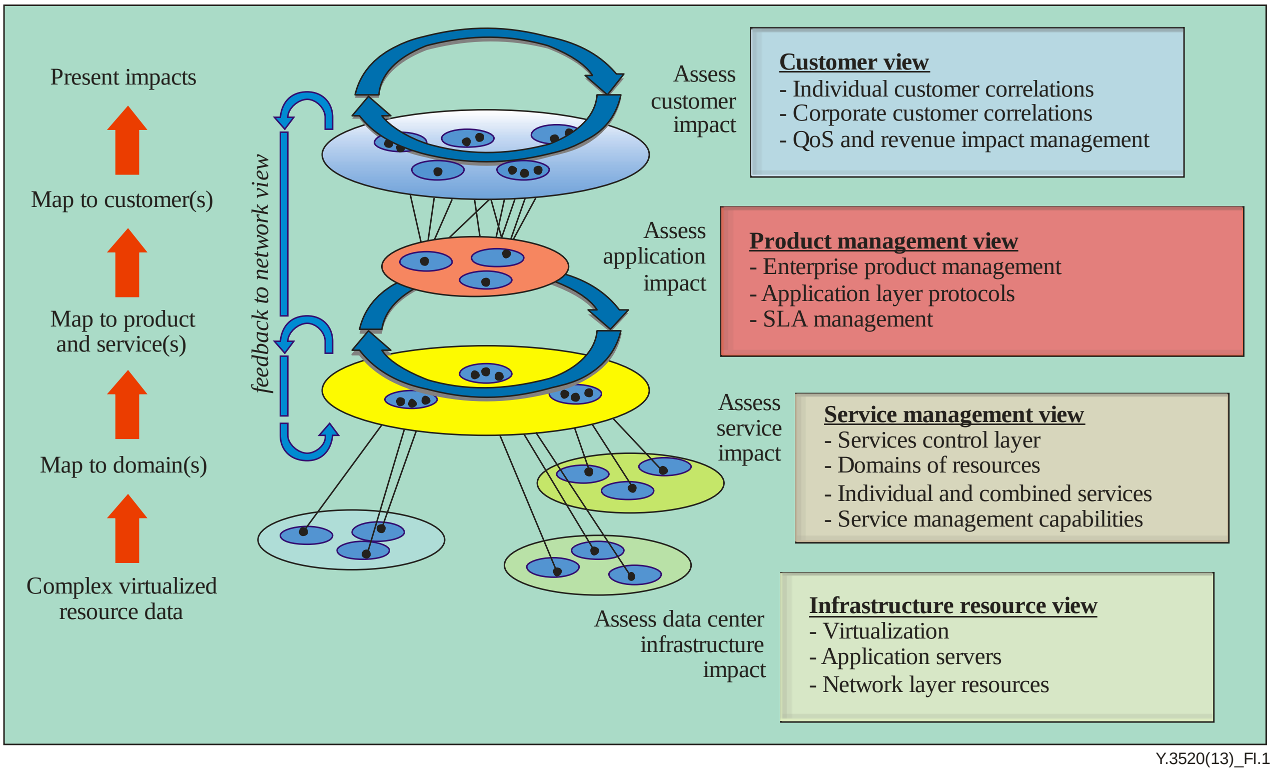

Present impacts
Map to customer(s)
Map to product and service(s)
Map to domain(s)
Complex virtualized resource data
feedback to network view
Assess customer impact
Assess application impact
Assess service impact
Assess data center infrastructure impact
Customer view
- Individual customer correlations
- Corporate customer correlations
- QoS and revenue impact management
Product management view
- Enterprise product management
- Application layer protocols
- SLA management
Service management view
- Services control layer
- Domains of resources
- Individual and combined services
- Service management capabilities
Infrastructure resource view
- Virtualization
- Application servers
- Network layer resources
Y.3520(13)_FI.1

NOTE – Domains are related to CSP.

Figure I.1 – Comprehensive view of management layers

Appendix II

Multi-cloud end to end service management

(This appendix does not form an integral part of this Recommendation.)

The following use case describes the challenges associated with multi-cloud end to end service management.

Figure II.1 illustrates an example where a cloud VoIP service is provided by CSP1 to CSP2 that is bundling it with other services and reselling a package to a cloud service customer (CSC). Even if CSP1 runs network services such as a content delivery network (CDN), CSP2 provides network connectivity services to the end user through its own core networks (e.g., IP/MPLS) and access networks (e.g., the backhaul, WiFi, 3G/4G/LTE and enterprise LAN/WAN infrastructures).

When a CSC, such as an IT department, has a problem with the cloud VoIP quality of service, it contacts CSP2 using a customer relations management (CRM) system. The CSP2 support agent should have the capability to see the health and welfare of the VoIP service from a holistic (end to end) perspective. This requires visibility into the VoIP and network resource management systems of both CSP1 and CSP2.

As shown in Figure II.1, there are the two types of connection paths:

1. **Service delivery path** – used by the functional interfaces of the services to deliver the combined service value to the customer. In this use case, cloud VoIP and IP/MPLS are combined to create a premium ICT bundle.
2. **Service management path(s)** – all of the logical management paths that perform operations and maintenance functions such as provisioning, service assurance and charging/billing of the relevant services to this bundle.

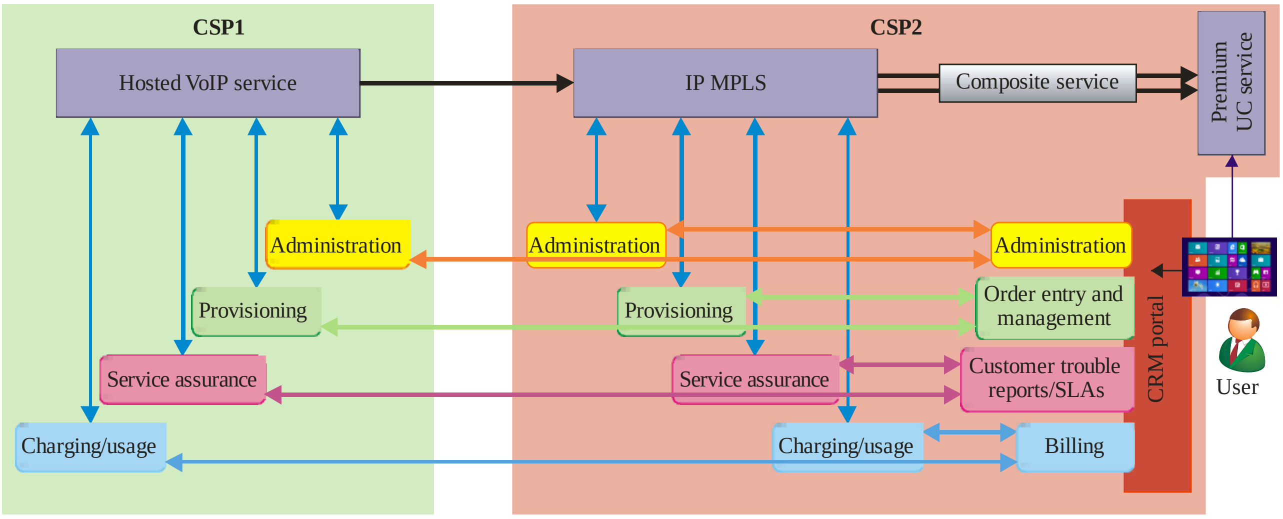

Y.3520(13)_FII.1

Figure II.1 – Managing multi-cloud services end to end

The delivery path for the service, via their functional interfaces, is not addressed by this use case.

What is addressed is efficient implementation of all of the resource management functions depicted by the lines between the customer relationship management (CRM) portal and the administrative, provisioning, service assurance and charging functions for each component (VoIP, etc.) that makes up a complete service. This challenge, associated with effective cloud resource management, is a major technical issue and can be a limiting factor for the adoption of cloud computing based solutions. In order for the composite cloud computing services to work effectively, all the prerequisite services of both CSP1 and CSP2 must function properly.

When either of the two CSPs becomes aware of a VoIP problem, tools are needed so that they can quickly resolve the problem in an effective manner. This includes being able to see via a service dashboard or CRM portal what has occurred relative to the VoIP service and to investigate in order to obtain greater details concerning any significant item. Additionally, the customer service agent should also be able to initiate an order for new or changed service configurations. However, if the agent lacks access to useful end to end cloud resource management tools and can only create a trouble ticket and then pass the problem off to another agent for action, the cloud service customer will be unsatisfied and this could potentially result in excessive operational expenses.

Appendix III

Summary of SES and SMI concepts

(This appendix does not form an integral part of this Recommendation.)

This Recommendation makes reference to various software enabled service (SES) and service management interface (SMI) concepts as developed in the TM Forum. This appendix is intended to provide a further informative introduction to these concepts; however reference should be made to the relevant TM Forum specifications for all technical details.

III.1 Software enabled service (SES)

A software enabled service is a service that exposes a management interface in addition to its functional interface (FI). Just like a router or switch exposes a simple network management protocol (SNMP) or other style management interface here we are referring to a digital service. However, since a digital service, such as represented by a Platform as a Service (PaaS) workload, is hosted on a virtualized cloud infrastructure, the cloud platform must enable the SMI for each instance of a given service/role at a given point in time. An SES may represent a physical device, a "software" instance, or indeed a distributed function that has no single location or instance.

The TM Forum developed the concept of SES to provide a means to allow consistent end to end management and metering of services exposed by and across different service provider's domains and technologies, such as communication or Web 2.0 services. The TM Forum specifications are intended to support business practices associated with multi-provider cooperation throughout the lifecycle of the service and this by means of lightweight design to foster wide adoption of the standard artefacts in any architecture, technology environment and service domain.

Management interfaces today are siloed per technology, standardized by specific SDOs or implemented by vendors as proprietary implementations. This renders the consistent management of services sourced from different domains as challenging.

The SES management approach proposes a way to allow consistent access to the software components for operations, administration and maintenance (OAM) tasks. This consistent access is achieved by incorporating the service management interface (SMI) in addition to the functional interface (FI) definition that is part of software component creation.

III.2 Service management interface (SMI)

In the context of managing an SES, the SMI concept delivers the ability to configure, activate or suspend a service instance and to receive or be notified of any kind of metrics, health state and detailed information about eventual failures, independent of the underlying technology or architecture.

Perhaps the best way to think of the SMI is as a simple "base class" in object oriented software development that defines the core management interface that can then be inherited by specific interface classes for specific purposes. The base SMI provides the set of operations supported by management objects, which can then be implemented using various management protocols.

The following operations are exposed on the SMI:

- Activation of an SES: Making the SES available for a particular context (deploying the SES)
- Provisioning of an SES: Configuring the settings of an SES or an SES instance
- State monitoring of an SES: Querying the history and current status in terms of life cycle management (for a specific instance of the SES) and listening for status updates
- Usage monitoring of an SES: Querying for usage metrics from the SES instance or listening for usage metrics reports or alarms (e.g., if metrics conditions imply notifications)

- Health monitoring of an SES: Querying for health metrics from the SES instance or listening to alarms from the resource
- Update of an SES: Modification of the setting or life cycle management status of an SES instance
- De-activation of an SES: making the SES unavailable in a particular context

III.3 SMI interface

The SMI support a set of simple operations to allow SES components to interact with management systems in a consistent way:

- getExecutionState
- getManagementReport
- getServiceConfiguration
- setExecutionState
- setServiceConfiguration

For further information about SMI concepts, refer to [b-TMF TR198].

Bibliography

[b-ITU-T E.107] Recommendation ITU-T E.107 (2007), *Emergency Telecommunications Service (ETS) and interconnection framework for national implementations of ETS*.

[b-ITU-T M.60] Recommendation ITU-T M.60 (1993), *Maintenance terminology and definitions*.

[b-ITU-T Y.1271] Recommendation ITU-T Y.1271 (2014), *Framework(s) on network requirements and capabilities to support emergency telecommunications over evolving circuit-switched and packet-switched networks*.

[b-ITU-T Y.2205] Recommendation ITU-T Y.2205 (2011), *Next Generation Networks – Emergency telecommunications – Technical considerations*.

[b-FGCC Part 4] ITU-T Focus Group on Cloud Computing – Technical Report (2012), *Part 4: Cloud Resource Management Gap Analysis.*
www.itu.int/dms_pub/itu-t/opb/fg/T-FG-CLOUD-2012-P4-PDF-E.pdf

[b-TMF TR198] TM Forum TR198, *Multi-Cloud Service Management Pack – Simple Management API (SMI) Developer Primer and Code Pack, Release 2.2*.
www.tmforum.org/?s=TR198

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Terminals and subjective and objective assessment methods
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	**Global information infrastructure, Internet protocol aspects and next-generation networks**
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2015