



Technical Report:

Using Cloud Standards for Interoperability of Cloud Frameworks

Thijs Metsch (SUN Microsystems), Andy Edmonds (Intel) and Victor Bayon (Intel)

This document is licensed under the Creative Commons Attribution Share Alike license (cc by-sa) 3.0 unported. For further information on this license and its obligations please see:



<http://creativecommons.org/licenses/by-sa/3.0/legalcode>

Abstract

Today's need for Open Frameworks, especially in the Cloud community, demands for standards. These standards are needed to ensure portability, interoperability and integration of Cloud sites. This technical report describes how an emerging Cloud standard, the Open Cloud Computing Interface, can be used to integrate two major Cloud Frameworks (SLA@SOI and RESERVOIR). Both of the Cloud Frameworks are developed with a focus on elasticity. Within the project different management layers are used to control the life-cycle of the Cloud Services. Still both frameworks can interoperate by using the Open Cloud Computing Interface API.

Introduction

What today is called Cloud computing comes a long way down the road. It had other names before and many technologies are involved in it. Virtualization, utility computing, and grid technologies are among the most representative.

Cloud offerings can be classified according to the resources they offer 'as a Service' (XaaS): Infrastructure as a Service (IaaS) that allows to allocate virtual machines and storage capacity; Platform as a Service (PaaS) where users are provided with remote software platforms to run their services; and Software as a Service (SaaS) where applications are moved to the Internet and accessed through web interfaces.

Cloud frameworks on the other hand can be seen as the software environment in which Cloud services can be deployed. Most of the frameworks have automatic and elastic management solutions inherited which control the life cycle, placement and Service Level Agreements (SLAs) of a Service.

Different frameworks have come up over the past time, including commercial proprietary and open frameworks like those developed in the European Union Framework 7 Programme Projects RESERVOIR [1] and SLA@SOI [2]. They consist of similar modules and layers but have different basic architectures. They can be messaging based or client/server, have a focus on IaaS or SLA management. The different foci leads to different architectures. SLA@SOI focuses on SLA-driven management and monitoring the life-cycle of services such as software and infrastructure services but can be applied to human-based services. RESERVOIR focuses on federated clouds with focus upon the management, interoperability and optimisation within such types of clouds.

Still all Cloud frameworks seem to need a layer to deploy virtual workloads on infrastructure. Having the frameworks interoperate even makes more sense with the background of using each others abilities and functionalities. In this manner services could be moved to the Cloud framework which would best fit their needs and their characteristics.

Having the need to interoperate comes with the demand for standards. Only if the frameworks support certain standardized interfaces, can interoperability be achieved. This technical report tries to show the overall setup and ideas behind two Cloud frameworks and includes the description of an upcoming Cloud standard. A architecture which combines all three aspects is proposed as well.

Cloud Frameworks

The two Cloud frameworks described in this technical report are part of research and development efforts of the

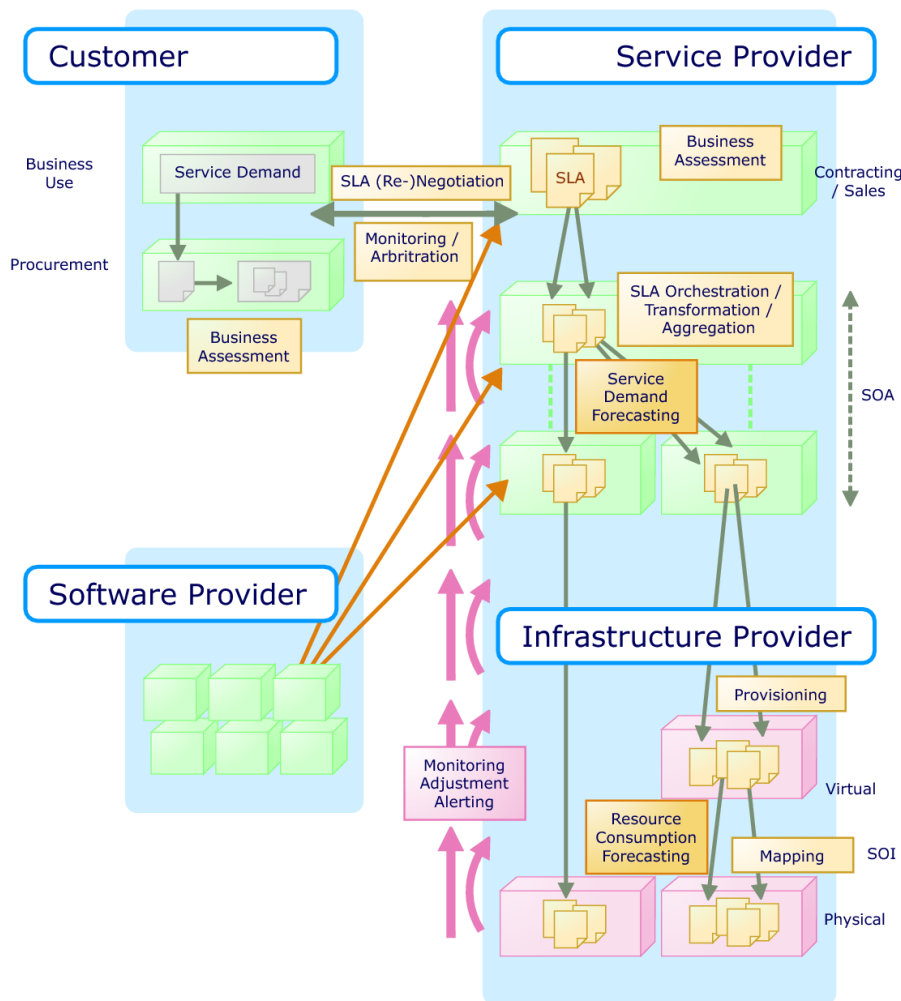


Figure 1: The High-level SLA@SOI Architecture

RESERVOIR and SLA@SOI projects. Although on a high level view they both manage services based on SLAs there are important differences.

The RESERVOIR framework shows a more layered approach in which the focus is on the management of virtual workloads. These workloads can be of different types like virtual machines or even Java services. Also RESERVOIR tries to break the barriers between Cloud sites which include geographical, security and management issues. Cloud services can be moved between sites to gain Interclouds.

While doing so the services have attached SLAs which for example can state that a Cloud service A should stay on the private Cloud so a Service B must be moved to a public Cloud to free resources for Service A.

The SLA@SOI framework provides a multi-layer, multi-domain SLA management framework. It is targeted at providing these management capabilities to service providers, from business service providers all the way down to infrastructure service providers.

The project tries to create an environment in which Cloud services can be traded as economic goods. For example when deploying a service, a set of dependable conditions can be defined.

Still both frameworks are able to manage virtual machines and the according SLAs. Scaling up and down as well as cloud bursting scenarios can be handled.

The SLA@SOI Framework

The research project SLA@SOI aims at providing a major milestone for the further evolution towards a service-oriented economy, where IT-based services can be flexibly traded as economic goods, i.e. under well defined and dependable conditions and with clearly associated costs. Eventually, this will allow for dynamic value networks that can be flexibly instantiated thus driving innovation and competitiveness.

The project SLA@SOI envisions "a business-ready service-oriented infrastructure empowering the service economy in a flexible and dependable way", where business-readiness requires the following three major characteristics:

1. *Predictability and Dependability* The quality characteristics of services can be predicted and enforced at run time.
2. *Transparent SLA management* Service level agreements (SLAs) defining the exact conditions under which services are provided/consumed can be transparently managed across the whole business and IT stack.
3. *Automation* The whole process of negotiating SLAs and provisioning, delivery and monitoring of services will be automated allowing for highly dynamic and scalable service consumption.

A motivating business scenario highlighting the project idea is a service provider who is enabled to offer services with differentiated, dependable and adjustable SLAs, and can negotiate concrete SLAs with (individual or groups of) customers in an automated fashion, adding value to provisioning systems. This business goal imposes additional requirements on software providers (to provide components with predictable non-functional behaviour) and infrastructure providers (to support an SLA aware management of resources).

This vision maps to the overarching challenge for a service-oriented infrastructure (SOI) that supports consistent SLA management across all layers of an IT stack and across the various stake-holder perspectives. Noteworthy, the SLA characteristics may span across multiple non-functional domains such as security, performance, availability, reliability. This technical report focuses on the infrastructural part of the SLA@SOI framework.

Figure 1 gives a simplified overview on how such a systematic SLA management process may look like. As today's business systems typically consist of complex layered systems, user-level SLAs cannot be directly mapped onto the physical infrastructure. Services might be composed of other more fundamental services that could be also provided by external parties. Consequently, a stepwise mapping of higher-level SLA requirements onto lower levels and the aggregation of lower-level capabilities to higher levels is crucial for grounding user-level SLAs to the infrastructure.

This vertical information flow must carefully reflect service interdependencies as well as the originating business context. In addition to SLAs, the vertical information flow also covers monitoring, tracking, and accounting data and must support brokering and negotiation processes at each layer. As shown in the figure, the overall SLA management process may include different stake-holders, namely customers, service and infrastructure providers, and also various business steps such as business assessment, contracting and sales. The overview is intentionally simplified in the sense that no service chains are visualized. Such chains would represent all cases where service providers rely on other external providers.

Infrastructure SLA Management

For the purpose of this report and discussion, we focus on the infrastructure provider's perspective (see figure 1). Today's IaaS offerings tend not to have automated SLA handling, what's provisioned is not easily verified [3], typically static, require manual intervention in the case of claims, making for inefficient processes. Many current infrastructure as a service (IaaS) providers consider SLAs to be "best effort". For example Slicehost [4] considers SLAs to be "silly" [5], Amazon EC2 [6] SLAs [7] are not machine readable and not negotiable.

An infrastructure provider client should get what they asked for from that provider. If you have a guarantee and a machine readable SLA that an associated infrastructure framework understands then you can verify that the provider is holding up its end of the agreement.

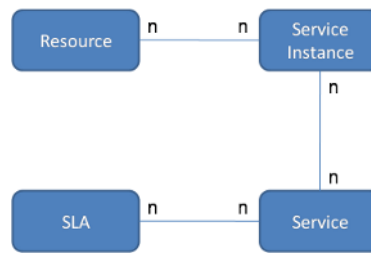


Figure 2: Main SLA@SOI Architectural Entities

Infrastructure SLA model and Resource Model

An Infrastructure SLA (iSLA) within SLA@SOI is conceptually comprised of:

- Set of terms:** Terms in an iSLA are the attributes that define the eventual representation of an instantiated infrastructure resource. These within the SLA@SOI framework are broadly split into two categories, functional and non-functional and are in principle immutable. A change of the value of the term will mean an SLA violation. Functional attributes are those that define attributes that are core to the operation of an infrastructure resource. Examples of these include: number of CPU cores, CPU slice (how much percentage of CPU a instantiated resource is allowed to take), amount of allocated memory, root partition size, network, etc. Non-functional attributes are those that do not directly affect the operation of an executing infrastructure resource. Examples of these include: resource's geographic region, security, green policy, replication, etc.
- Set of SLOs:** SLOs are explicit and possibly implicit attributes that are monitored for changes in their value and are associated at runtime properties of the resource. The SLOs are mutable. An example of an SLO is CPU or I/O utilization. With different levels of granularity SLOs could be for example shared among the service provider and the service consumer. The service consumer will know about how much capacity of the current provided solution is using and the service provider could recommend a better updated configuration with less or more resources as required.
- Management Rules:** Where iSLA terms and/or SLOs change above or below a particular explicit or calculated values then a set of associated actions are executed. For example, if a provisioned resource unexpectedly terminates the iSLA term state changes and an SLA violation action is executed. Another example is if CPU utilization is too high for a pre-determined period of time, an SLA warning action is executed. Both actions are of the interest of the resource provider and consumer. Currently within the SLA@SOI framework, the consumer of the infrastructure can listen to such events by means of subscribing to a private pub/sub channel.

Relationship between iSLA Registry and Infrastructure Landscape

Once an iSLA is agreed and provisioned this action links the associated provisioned resources, reflected in the infrastructure landscape (IL), with the iSLA. iSLAs are persisted in a iSLA registry and if one is to inspect the iSLA registry, a collection of iSLAs can be viewed and their associated resources. This allows management be navigated from the iSLA to the Resource level via the iSLA registry through to the IL.

These iSLAs are managed by the iSLA manager and persisted within an IL, which represents the instantiated resources and immutable SLA terms. The current SLA@SOI infrastructure resource model represent the instantiated SLA plus other running parameters. The main entities considered in the model are Compute, Storage and Network with particular attention paid to the compute entity from a technical and implementation point of view.

When instances of this model are created through the provisioning process, they are registered within the IL. Also registered in the IL are the physical resources on which virtual ones execute upon. The IL is in effect a configuration management database. Free and used resources can be monitored here. When the running state of a resource changes, it is checked against the IL.

The registration and publishing the information against the IL is done by means of the pub/sub messaging pattern. Each physical resource is responsible for updating its status and the status of the provisioned resources (virtual machines) that it manages.

SLA@SOI Architecture

In an ongoing process, the existing architecture [8] has been further refined such that the separation of concerns between Resource management and SLA management has been introduced. This will provide the advantage such that integration with other Resource managers will quite trivial to achieve. The major entities within the SLA@SOI framework are SLA, Service and Resource and their relationships are depicted in figure 2.

An infrastructure SLA manager is linked with the Resource Manager. The Resource manager offers a number of services and manages service instances. The SLA manager accepts SLA offers and advertises service offerings, otherwise known as SLA Templates. SLA offers are requests for infrastructure, defined by the iSLA model and are the basis of the negotiation that happens between clients and the infrastructure SLA manger. Clients can query the infrastructure SLA manager in order to see what SLA Templates the iSLA manager supports. Once the SLA offer has been successfully negotiated the client can

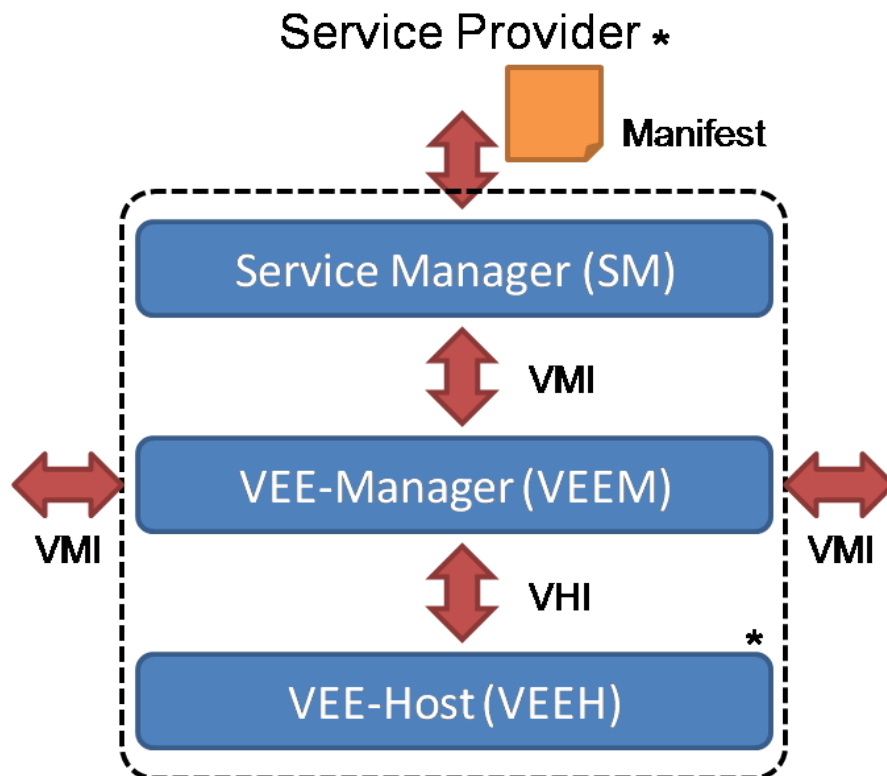


Figure 3: The High-level RESERVOIR Architecture

then provision the agreed SLA. In order to support the architecture scenario outlined, it is necessary for the Resource manager to support a number of critical functional APIs. Required services for infrastructure SLA management are:

- *Services Offer Listing:* This service supplies a listing of the services that the Resource manager can offer. These service offers are known as SLA Templates with the SLA@SOI framework. Examples of these include compute services and storage services. The compute service offers the service to allow for computation to be executed on the provider's infrastructure. This is typically take the form of a virtual machine in the IaaS world. In the example of a storage service, this offers a service to allow storage of data.
- *Service Instance Creation:* For every service offered, an API for creating the related service offers is required. At a minimum, the basic operations of Create, Retrieve, Update and Delete (CRUD) should be offered.
- *Service Instance Monitoring:* This service is in place to guarantee that the terms specified by the iSLA are delivered and that are not changed while the resource is active. In terms of infrastructure, basic parameters such as CPU share allocation, amount of memory, network capacity are monitored such that their agreed guarantees are met and not violated

by the resource provider. For example the resource provider might under-provision the resource during provisioning time and this situation should be reported to the consumer of the resource when appropriate.

- *Service Instance Adjustment*: Again, for every service offered, a means to change the service instances' configuration should be offered. This service, in practical terms, should be simply a proxy to the service instance creation API, where by it utilises the Update functionality of the service instance creation API.

The RESERVOIR framework

The RESERVOIR project [1] tries to create an environment in which Cloud Services can easily be deployed and managed. This environment provides an abstraction layer from specific resources, platforms and geographical locations.

Beside the major idea of creating abstractions, the project tries to create an environment in which services can be migrated. Migration should be possible across network and storage boundaries. In this case, both a "live" and a "suspend and resume" method of migration are supported.

RESERVOIR encapsulates all of these ideas of an environment, in which different kind of services can be managed. It also tries to create a basic model of how Cloud frameworks can look like when overcoming political, security, geographical and migration boundaries.

The first abstraction the RESERVOIR project uses is the encapsulation of services. A service is the collection of infrastructure resources, specified as Virtual Execution Environments (VEEs), which are deployed by RESERVOIR for the delivery of RESERVOIR's customer services (e.g. PaaS, SaaS).

Several VEEs can then run on a Virtual Execution Environment Host (VEEH); meaning one physical resource. Each host has a virtualization technology installed in which the VEEs are deployed and hosted. A small component for management and monitoring is also available on all VEEHs.

The overall management of all VEEHs is realized in the Virtual Execution Environment Management (VEEM) system. It is in charge of the deployment of VEEs on top of VEEHs. It can bootstrap and unload VEEHs. VEEs can be moved around and placed according to the VEEM setup.

To complete the overall system, a Service Manager is needed. It is responsible for the instantiation of service applications. It, therefore, requests the VEEM to create and manage VEEs. Beside this the Service Manager also tracks and manages the SLAs. It ensures that all SLA policies are satisfied at all times. To do so, it monitors the current state of the overall system and executes the elastic rules. [9]

RESERVOIR can be seen as the perfect environment in which Cloud services can be deployed which have different needs for resources at different times. Overall multi-tier cloud service can be deployed which have a high demand for compute resources can be managed easily.

Virtual Execution Environment Management

Within the RESERVOIR project OpenNebula [10] is used as a VEEM. The Service Manager interface with the VEEM through the VEE Management Interface (VMI). The VEEM itself use a VEE Host Interface (VHI) to communicate with the VEEHs.

OpenNebula is an open source virtual infrastructure engine which provides features to control, deploy and monitor virtual machines. A modular and open architecture is used which support the integration of all kind of hypervisors.

Three central components are present within OpenNebula. The Core is a centralized component which manages the life-cycles of the virtual machines. Next to the core a Capacity Manager adjusts the placement of the virtual workloads. To support multiple hypervisors a third component is used for abstraction. The Virtualizer Access Layer exposes the features of the hypervisors to the Core. [11]

To interface with an OpenNebula system the OpenNebula Cloud API (OCA) supports the development of Cloud interfaces. This interface is used to implement the VMI inside of the RESERVOIR architecture. Next to this OpenNebula also support a subset of Amazon's EC2 queries.

To place the virtual workloads on the resources requires the VMI to be as precise as possible. This means that requests need to be complete and information regarding the placement need to included. The Open Cloud Computing Interface will form the base of the VMI.

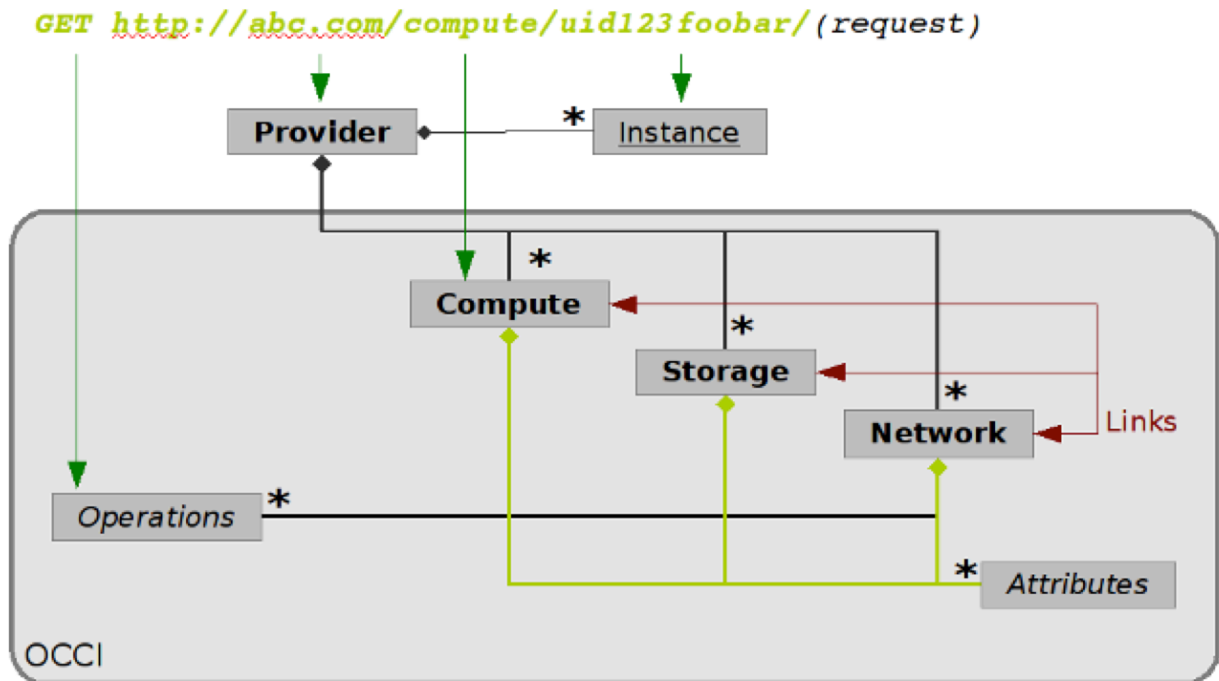


Figure 4: An Overview of the OCCI Model

Gaining Interoperability

Interoperability can be reached when two systems use the same interface and so can exchange information, defined by that standard's data model. A common interface for both previously described frameworks would be an API which supports the deployment, management and monitoring of virtual workloads like virtual machines.

The Open Cloud Computing Interface

The Open Cloud Computing Interface (OCCI) is a working group organized within the Open Grid Forum [16]. Motivation for founding this group was the lack of a standard for Infrastructure as a Service model based clouds. The open standardization process is driven by the following topics:

- *Interoperability* describes the ability to let different systems integrate with each other. Most known use case for this is the creation of Interclouds.
- *Portability* defines the need for easy code reuse in end-user application like cloud clients or portals.
- *Integration* is the idea of wiring up Cloud computing services with legacy re- sources.

With the focus of providing Infrastructure as a Cloud Service the OCCI group will define a RESTful [18] protocol. The goal is to create a small interface which can be easily extended. The RESTful approach supports this.

Based on use cases and requirements [15] which define how the end-users interfaces with the Service provider a first protocol is defined. The overall specification consists of four parts which shows the modularity and extensibility.

The OCCI protocol highly relies on the HTTP [17] protocol. Each resource is identified by an unique URL. Still OCCI can present an machine readable interface (HTTP rendering) and a human-readable interface (XHTML5). Still OCCI itself is not dealing with how data is represented but focuses on the portability, interoperability and integration aspects.

While using HTTP as underlying protocol OCCI share the same manner for addressing, authentication (OAuth) and versioning. Entry-Points can be the definite resource as well as collections of resources. Figure 4 gives a rough overview of the OCCI protocol.

Operations on resource are model using the CRUD model: Create, Retrieve, Update and Delete. Each of these operations can be mapped to an HTTP operation: POST, GET, PUT, DELETE. These operations can trigger actions. Each actions is defined by a link. This allows clients to easily identify actions even from different cloud providers. Actions can be for example be resize, start and stop of virtual machines.

So if a service provider offers an offering for the management of Virtual Machines it offers kinds which have a set of attributes on which actions can perform operations. A kind of a resource is defined within the OCCI namespace and can described compute, storage or network resources. Attributes are expressed in key-value pairs. For example a storage kind would have the attribute `occi.storage.size` and support the actions: offline, online.

Resources are be organized into categories which can be extended by cloud providers. For example a storage resource would be defined by a category which has the following attributes:

```
Category: storage;
        scheme="http://purl.org/occi/kind#";
        label="Storage Resource"
```

The scheme for the category are defined with the OCCI namespace. In this example the definition of the `#kind`.

Whenever a operation could return a list of resources (Collection) the results can limited using pages. E.g. stating the client wants to have the results 10 results and the Next 10. To model associations between resources the concept of linking is used. In the case that a virtual machine would like to link to a storage resource it would look like:

```
Link: </storage/disk1>;  
rel="http://purl.org/occi/kind#storage";  
title="Quorum Disk"; device="sda"
```

Overall the OCCI specification tries to address the need for interoperability and portability. The usage of namespaces and the easiness to extend them with own features, descriptions, actions, etc. allow Cloud providers to implement an interface which matches their own existing Cloud API but in the same way ensures interoperability.

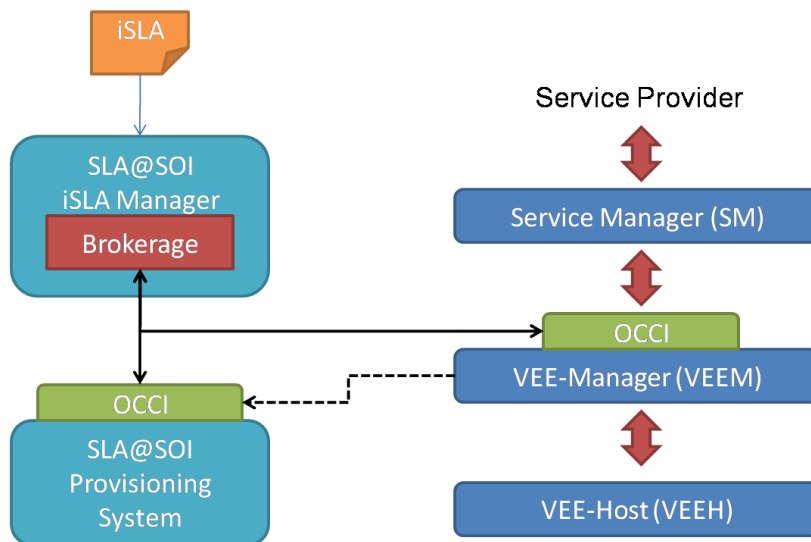


Figure 5: The Use-case Scenario Topology

A Hybrid Cloud framework

In order to allow for IT-services to be traded as a commodities there is a requirement for different infrastructure service providers to interoperate via standardized interfaces. These offer both access to infrastructure provisioning and iSLAs management capabilities. These aspects can be managed from the consumer's perspective across domains and in multi-providers settings in a transparent and unified way.

As more services are moved to the Cloud, there will be also a need for the service consumer to retrieve information about all stages of service provisioning and service usage during runtime. Service providers should offer different layers of functionality and granularity, such as monitoring. The service consumer will require this level access in order to perform its own monitoring of the provisioned services.

With the cloud frameworks of SLA@SOI and RESERVOIR presented we now detail a use case scenario that illustrates the benefit that OCCl brings to the scenario. In this scenario the SLA@SOI framework is configured to only provision and manage iSLAs. There are two infrastructure provisioning systems; one that is the "native" SLA@SOI provisioning system and another 3rd party provisioning system, the RESERVOIR VEEM. The topology of such a scenario is illustrated in figure 5. Naturally the SLA management layer could know about more providers but for the purposes of simplicity we have selected these two.

The SLA@SOI framework has the responsibility of acting as a broker. Requests, in the form of iSLAs, are supplied to the SLA@SOI SLA management layer. Within the SLA management layer, a number of provisioning systems are known. The iSLAs presented to the SLA manager which then selects the most appropriate provider based on the terms (functional and non-functional parameters) within the iSLA. Once selection is accomplished the SLA manager then negotiates for the required amount of resources, provisions and then begins to monitor the resources based on the iSLA's SLOs and conditional rules.

In this particular use case scenario the iSLA presented to the SLA management is one that requires support for a federated multi-tier Cloud services. This requirement cannot be satisfied by the internal SLA@SOI provisioning system but can be through RESERVOIR. The SLA management layer knows this as it can retrieve the services offered by each provider through the Service Listing Interface.

Provisioning and monitoring is then carried out and mediated by the OCCl interface. The advantage in this approach is that there is no custom provider-specific implementations required to enable firstly these actions and secondly, as explained previously, provides a way for new providers to be utilized owing to the ease of implementing OCCl. The selection of providers carried out by the SLA manager between the 2 systems is made easy via OCCl, otherwise the selection is made all the more difficult due to the differences in the two systems provisioning APIs. SLA@SOI uses OCCl to go straight to our two different infrastructure management layers without the need for multiple provider adapters. In this use case SLA@SOI is used as a means to perform infrastructure provisioning negotiation, brokerage and the management of the agreed iSLAs. OCCl provides the capability for two major European FP projects, RESERVOIR and SLA@SOI, to easily integrate by sharing a common infrastructure resource manager interface (provided by OCCl).

Conclusions

The proposed architecture in this technical report showed how two cloud frameworks could be coupled to gain interoperability. The usage of OCCl allows the SLA@SOI framework to deploy workloads to the RESERVOIR environment. All this is done in a way that no SLAs are breached.

OCCl plays a major role in this setup. It describes the interface between the two frameworks. Is one of the first standards which arises in the cloud community. But the current specification can already be used to propose the architecture.

Overall this setup showed that it can be possible to let two Cloud frameworks interoperate even with their different architectures. The overall advantage would be that Cloud service can be deployed and managed in a environment which best fits their need. In this case groups of services which intercommunicate can be managed in the SLA@SOI framework and deployed in the RESERVOIR frame- work.

Notes and Comments. This technical report demonstrates a work in progress approach for interoperability of Cloud frameworks. The Open Cloud Computing Interface specification will be released in near future.

The authors would like to thank the members of the Open Cloud Computing Working Group and the members of both the RESERVOIR and SLA@SOI teams.

The research leading to these results is partially supported by the European Community's Seventh Framework Programme ([FP7/2001-2013]) under grant agreement no.215605 (Reservoir) and no.216556 (SLA@SOI).

References

1. RESERVOIR project website <http://www.reservoir-fp7.eu>
2. SLA@SOI project website <http://sla-at-soi.eu/>
3. Allspaw, J.: Slides from Web2.0 Expo 2009. (and somethin else interestin), retrieved on the 6th August 2009. <http://www.kitchensoap.com/2009/04/03/slides-from-web20-expo-2009-and-somethin-else-interestin/>
4. Slicehost VPS Hosting, retrieved on the 6th of August 2009. <http://www.slicehost.com>
5. Slicehost VPS Hosting SLA, retrieved on the 6th of August 2009. <http://www.slicehost.com/questions/#sla>
6. Amazon EC2, retrieved on the 6th of August 2009 <http://aws.amazon.com/ec2>
7. Amazon EC2 (see Credit Request and Payment Procedures), retrieved on the 6th of August 2009. <http://aws.amazon.com/ec2-sla/>
8. Theilmann,W.,Yahyapour,R.,Butler,J.:Multi-level SLA Management for Service-Oriented Infrastructures, Service Wave (2008)
9. Caceres, J., Montero, R., Rockwerger, B.: RESERVOIR - An Architecture for Services, RESERVOIR project, <http://www.reservoir-fp7.eu/twiki/pub/Reservoir/Year1Deliverables/080531-ReservoirArchitectureSpec-1.0.PDF> (2008)
10. OpenNebula website <http://opennebula.org> 11. Sotomayor, B., Montero, R., Llorente, I., Foster, I.:Capacity Leasing in Cloud Systems using the OpenNebula Engine, Workshop on Cloud Computing and its Applications (CCA08) (2008)
11. Open Grid Forum <http://www.ogf.org>
12. Fielding, R.T.: Architectural Styles and the Design of Network-based Software Architectures, Doctoral dissertation, University of California, Irvine (2000)
13. Open Cloud Computing Interface working group website <http://www.occi-wg.org>
14. Metsch, T (2010) Usecases and requirements for a cloud API. Tech. rep., Open Grid Forum, <http://ogf.org/documents/GFD.162.pdf>
15. OGF (2010) Open grid forum website
16. (1999) Hypertext transfer protocol. <http://tools.ietf.org/html/rfc2616>
- 17.Fielding RT (2000) Architectural styles and the design of network-based software architectures. PhD thesis, University of California, Irvine