



EGI-Engage

VM snapshot support: OCCI extension, final specification

D4.2

Date	25/11/15
Activity	WP4
Lead Partner	CSIC
Document Status	FINAL
Document Link	https://documents.egi.eu/document/2643

Abstract

This Report deals with the effort to standardize the process of saving the current state of a virtual machine to allow creation of additional instances based on that state. The standardization of virtual machine snapshotting with OCCI, originally envisioned as a separate extension to the standard, was finally achieved by augmenting the existing standard during preparation of the OCCI 1.2 release. This paves the way to future implementation of this new functionality on EGI infrastructure via EGI-Engage.



This material by Parties of the EGI-Engage Consortium is licensed under a .
The EGI-Engage project is co-funded by the European Union (EU) Horizon 2020 program under Grant number 654142

COPYRIGHT NOTICE



This work by Parties of the EGI-Engage Consortium is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). The EGI-Engage project is co-funded by the European Union Horizon 2020 programme under grant number 654142.

DELIVERY SLIP

	<i>Name</i>	<i>Partner/Activity</i>	<i>Date</i>
From:	Zdeněk Šustr, Boris Parák	CESNET	25.11.2015
Moderated by:	Kostas Koumantaros	GRNET/PMB	25.11.2015
Reviewed by	Diego Scardaci Enol Frenandez	INFN/WP3 IFCA/WP6.2Task Leader	25.11.2015
Approved by:	AMB and PMB		30.11.2015

DOCUMENT LOG

<i>Issue</i>	<i>Date</i>	<i>Comment</i>	<i>Author/Partner</i>
0.1	22/10/2015	Document created	B. Parak / CESNET
0.2	03/11/2015	Internal review changes	A. Lopez / CSIC
0.3	24/11/2015	External review changes	Z. Šustr / CESNET A. Lopez / CSIC
FINAL	25/11/2015	Final version	A. Lopez / CSIC

TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

Contents

1	Introduction	4
1.1	Use cases	4
2	Virtual Machine Snapshotting	6
2.1	Changes introduced into the standard	7
3	Roadmap and future plans	8
	Appendix I. OCCI Infrastructure Specifications	9

1 Introduction

The OCCI (Open Cloud Computing Interface)¹ standard is a set of community driven specifications delivered by the Open Grid Forum (OGF)², allowing for the development and deployment of interoperable clouds. The OCCI specification³ consists of several complimentary documents defining the core abstract model, its renderings and extensions to the core model. Currently, the Infrastructure extension contains the required resource types, attributes and actions needed to manage Infrastructure as a Service (IaaS) resources.

The current version of the standard at the time of this deliverable is OCCI v. 1.1. The new version 1.2 has been developed in the last months, passing through the public comment phase that ended in July 2015. Currently the v. 1.2 version of the standard is going through the OGF processes to be officially released to the public.

A common use case and requirement from IaaS customers is creating copies or snapshots of running virtual machines. This is highly desirable when the virtual machine has undergone major changes of its operating system environment, when a long running service wants to be preserved, or when a given virtual machine needs to be duplicated. It is worth to note that for the purpose of this document, the term “snapshotting” refers to the process of deriving a new VM image from an existing Virtual Machine (Compute) instance, so that new instances can be spawned from the “snapshot” created. This is contrary to another usual use of the term, where snapshots are only the recording of the current state of an instance to allow the same instance to be returned to that state. Therefore the definition of “snapshotting” should support the creation of “snapshots” as VM images with cleanly bootable Operating Systems.

The OCCI standard lacked a clear support for this functionality, as there was no snapshotting or save action available, even if this functionality was already available in most of the cloud backends being used. EGI Federated Cloud users perceived this lack of support in the OCCI standard as a considerable drawback to embrace its usage as they had to circumvent this issue by contacting each particular resource provider that created the snapshots on their behalf, using the cloud management framework native interfaces.

1.1 Use cases

Virtual machine snapshotting is a common functionality in the current cloud resource providers and management frameworks, and it is commonly used to satisfy two use cases: virtual machine images customization, and backup and archival.

1 <http://occi-wg.org/>

2 <https://www.ogf.org/>

3 <http://occi-wg.org/about/specification/>

Regarding the virtual machine image customization, Infrastructure as Service users tend to create a virtual machine starting from a simple and minimal virtual machine image. This basic virtual machine is then customized and configured with additional software to fit the user requirements. For the simpler applications and customizations, this task can be done automatically at boot time via contextualization mechanisms, but in the case of more complex software this needs to be done manually by the virtual machine operator, consuming a considerable amount of work. Therefore, whenever a virtual machine is configured to satisfy their requirements, users tend to create snapshots of their running machines so that additional instances can be created based on that new image. This way, it is possible for a group of users to rely on the same baseline virtual machine image that has been customized and configured with the needed software and applications.

The second use case is the backup of running virtual machines. This is not only useful for backing up long-running executions (if the cloud management framework supports *hot* snapshots) but also to ensure that a given version of a virtual machine can be preserved along the time. Whenever an actualization is done, users can easily preserve the original state, so that it is possible to roll-back to it if the new update is found to be problematic.

As it was already stated, OCCl 1.1 lacked this functionality; therefore the EGI Federated Cloud users were unable to manage snapshots on their own. Instead, they had to perform individual requests to each of the particular resource providers where their virtual machines were running asking for a snapshot to be made. Then, the resource provider would create the snapshot and would make it available for the user or group. This produced a negative impact on their experience using the infrastructure, as the users were unable to self-service their own snapshots on-demand as they had to rely on an intervention by the resource provider.

2 Virtual Machine Snapshotting

In the envisioned reference scenario for Virtual Machine snapshotting, a user can have a running Virtual Machine (VM), which they want to use as a starting point for further work. They should be able to derive a usable Virtual Machine Image from that machine, and spawn further instances based on that template. Using OCCl terminology Virtual Machines are modelled as Compute Resources, and VM images as Operating System (OS) Templates.

Up until, and including, OCCl v. 1.1, the OCCl Infrastructure specification only recognized a snapshot action associated with Storage type resources. This snapshot action could not be applied to virtual machines (i.e. Compute type instances) therefore it could not be applied for taking a snapshot of a running virtual machine.

Therefore, the OCCl Infrastructure definition for OCCl v. 1.2 has been updated based on EGI-Engage comments and discussions carried out during the public comment stage to introduce a save action, which, when called for an existing compute instance, creates a new Operating System (OS) Template.

The specification makes it possible for the user to optionally specify a method (hot or deferred) for the save action. The chosen method will, in case of hot, instruct the implementation to make an immediate (potentially inconsistent) copy without interfering with the running instance or, in case of deferred, to gracefully stop the instance, safely copy its data and start it again. An optional name attribute will be available for the user to specify their desired name of the OS Template, and the server may choose to use that name if technically possible.

The save action will then produce an OS template based on the content of the disks of the originating compute instance and, on success, return an identifier of the newly created OS Template. The action target state over the Compute instance will be “active”, via a “stop” and “start” chain if needed. Once the snapshot is created, it will be discoverable as any other virtual machine image available at the resource provider, but it is out of the scope of OCCl how to export that image to outside the resource provider domain.

As with any other OCCl action, if the functionality is not supported in the underlying backend, the OCCl implementation MUST return any of the HTTP error codes, according to the OCCl HTTP rendering.

As a reference, we attach the OCCl 1.2 Infrastructure Extension drafts as it stands as of the date of this deliverable. However, it is worth notice that these are not the final documents, as they are subject to the OGF processes that are still ongoing.

2.1 Changes introduced into the standard

The Section 3.1 “Compute” from the OCCl Infrastructure Extension has been augmented as follows. An ellipsis (...) means that unmodified parts of the standard have been omitted.

3.1 Compute

(...)

Action term	Target state	Attributes
(...)		
save	active (via stop and start chain)	method={hot, deferred}, name=String

(...)

Action “save” is expected to create an OS Template referencing an independent copy of the current state of the Compute instance. The provider MAY choose to respect the “name” given by the client or override it according to its internal policies. A successful execution of this action MUST lead to a response containing the rendering of the newly created OS Template as defined by the chosen rendering and transport protocol. The provider MAY choose to include a reference to the original Compute instance in the Mixin attributes of the newly created OS Template.

3 Roadmap and future plans

The EGI-Engage proposed changes for standardizing the process of saving the current state of a virtual machine have been accepted in the public comment phase for the elaboration of the next OCCI 1.2 standard by the OGF. The preparation and release of the new version of the standard is currently undergoing the internal OGF processes, but no substantial changes are expected to the proposed extensions.

Although outside of the scope of this deliverable, the implementation of this functionality on EGI infrastructure within the scope of EGI-Engage depends on OCCI 1.2 support being introduced for each of the OCCI interfaces for the available cloud management frameworks supported in the EGI Federated Cloud. The first OCCI 1.2 preview release is planned for March 2016. The final OCCI 1.2 release will depend on the delivery by the OGF of the final OCCI 1.2 specifications, but it is expected for the 2nd quarter of 2016.

Regarding the adoption and deployment of OCCI 1.2 by the EGI Federated Cloud sites, it is expected that resource providers deploy the new interface as soon as it is available. OCCI 1.2 solves several usability shortcomings (such as the lack of a JSON rendering) that are present in the OCCI 1.1 implementation.

Appendix I. OCCI Infrastructure Specifications

As a reference, we attach the OCCI 1.2 Infrastructure Extension draft as it stands as of the date of this deliverable. The latest version of the draft document is provided in the following link:

<https://documents.egi.eu/document/2643>