OpenSAF and VMware from the Perspective of High Availability

Ali Nikzad al_nik@encs.concordia.ca **Ferhat Khendek** ferhat.khendek@concordia.ca **Maria Toeroe** maria.toeroe@ericsson.com

Engineering and Computer Science, Concordia University, Montreal, Canada

Core & IMS, Ericsson, Montreal, Canada

Abstract

Cloud services are becoming one of the most popular means of delivering computational services to users who demand services with higher availability. Virtualization is one of the key enablers of the cloud infrastructure. Availability of the virtual machines along with the availability of the hosted software components are the fundamental ingredients for achieving highly available services in the cloud. There are some availability solutions introduced by virtualization vendors like VMware HA and VMware FT. At the same time the SAForum specifications and OpenSAF as a compliant implementation offer a standard based open solution for service high availability. In this poster, we investigate these solutions for availability through experiments, compare them according to metrics and based on the results propose architectures that combine them to provide highly available applications in virtualized environments.

Introduction

The term virtualization broadly describes the separation of a resource or request for a service from the underlying physical delivery of that service. Several virtualization products exist. Among these solutions VMware is one of the virtualization solution providers which has tackled the problem of availability.
VMware has introduced two solutions for providing availability, VMware HA and VMware FT. Both solutions are available in VMware vSphere.

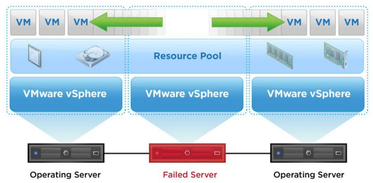

Figure 1 – VMware HA

The Application Interface Specification (AIS) is a set of middleware services defined by the Service Availability Forum (SAForum) to enable the development of highly available applications. OpenSAF is an open source SAForum compliant middleware implementation.
The Availability Management Framework (AMF), one of the most important AIS services, plays the key role in keeping an application's services highly available by coordinating its redundant resources, and performing recovery/repair actions in the case of a failure. AMF manages the application components and recovers their services according to the configuration provided with the application. This configuration represents the architecture of the application from the AMF perspective and describes the different entities composing it and their relations.

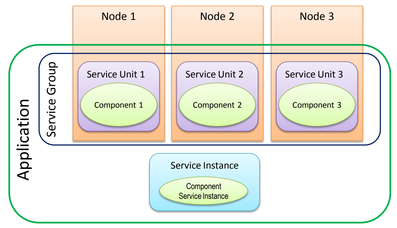

Figure 2 – Sample AMF configuration

To evaluate the two solutions and their combinations from the perspective of availability, we defined a set of metrics

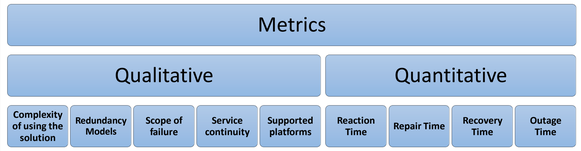

Figure 5 – Defined set of metrics

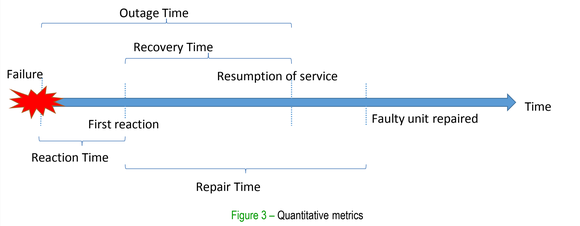

Figure 3 – Quantitative metrics

Baseline architectures

For the case study we selected the VLC media player application as it has been already modified to work with OpenSAF as a SA-Aware component.

Figure 4 – VLC media player as our case study application

OpenSAF on physical nodes

The first baseline architecture is the deployment of OpenSAF on the physical nodes. Since our case-study is configured with the 2N redundancy model we selected two nodes to host our experiments.

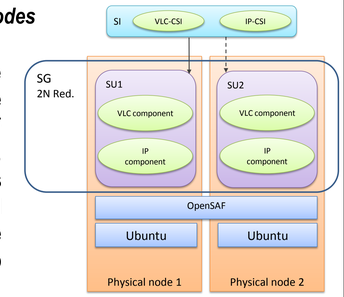

Figure 5 – First Baseline architecture

VMware HA baseline architecture

We created a vSphere cluster using 2 ESXi nodes and enabled VMware HA on the cluster using VMware vCenter. We also added one VM with Ubuntu Linux and VLC installed on it. We put the VM image on an NFS shared storage so that it is accessible from all cluster nodes.

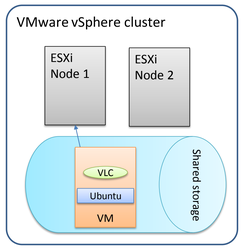

Figure 6 – Baseline architecture based on VMware HA

OpenSAF on virtual nodes

To take advantage of the virtualization which VMware provides and the service high availability management of OpenSAF we combined these two solutions. In this first combination we deployed the OpenSAF cluster on virtual nodes rather than on physical nodes.

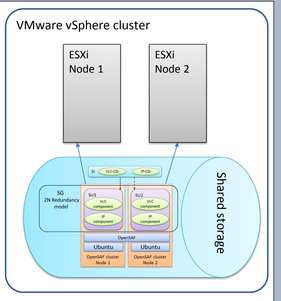

Figure 7 – OpenSAF on virtual nodes

Architectures	VLC failure	VM failure	Node Failure
OpenSAF on physical nodes with SA-Aware VLC component	√	Not applicable	√
OpenSAF on physical nodes with Non-SA-Aware VLC component	√	Not applicable	√
OpenSAF on virtual nodes with SA-Aware VLC component (with/without VMware HA enabled)	√	√	√
OpenSAF on virtual nodes with Non-SA-Aware VLC component (with/without VMware HA enabled)	√	√	√
VMware HA	Not detectable	√	√

Table 1 – Baseline architectures and experimented failures

Measurements and analysis

- No component failure detection in VMware HA
- SA-Aware components have better performance than non-SA-Aware components
- VMware HA unlikely can provide service high availability which is 99.999% availability (5.26 minutes of down time per year)
- Long repair time in the third baseline architecture (VM restart time)

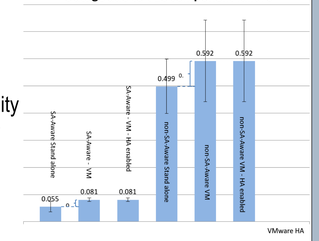

Figure 8 – Outage due to VLC component failure

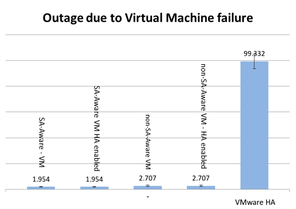

Figure 9 – Outage due to VM failure

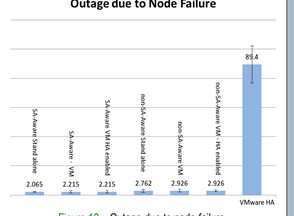

Figure 10 – Outage due to node failure

	Repair of		
	Failed component	Failed VM	Failed Node
OpenSAF on Standalone machine	Yes	-	No
OpenSAF in VM	Yes	No	No
VMware HA	No	Yes	Yes*
OpenSAF in VM + HA	***Yes***	***Yes***	***Yes****

* By restarting the VM on another host

Table 2 – Baseline architectures and repair of the failed unit

Proposed architectures combining OpenSAF and virtualization

VM availability management with non-bare-metal hypervisor

- Two OpenSAF clusters for VM availability management and service availability within VMs
- Using VMware CLI commands for starting and stopping VMs
- VM failure detection by OpenSAF passive monitoring
- Deployed on VMware workstation

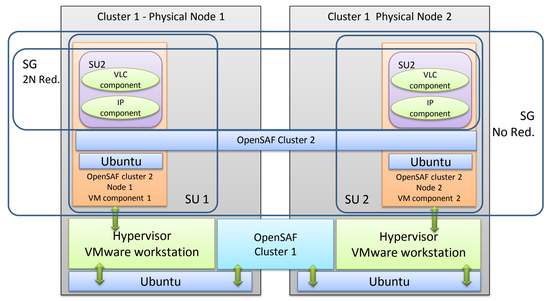

Figure 11 – VM availability management in non-bare-metal hypervisor

VM repair time is reduced drastically.
Without OpenSAF it took VMware HA about 60 seconds to detect the failure and 30 to repair

	Repair	Outage
ESXi without OpenSAF (VMware HA manages the VMs)	27.166	99.332
OpenSAF with ESXi (VMware HA manages the VMs)	*107.90*	1.953
OpenSAF with VMware Player (OpenSAF manages the VMs)	*3.73*	3.505

Table 3 – Comparison of measurements for VM failure in different architectures

The drawback: increased service outage because of the additional layer of the host operating system and the difference between the hypervisors used (ESXi vs. Workstation)

	Repair	Outage
OpenSAF with no virtualization	0.136	0.055
OpenSAF with ESXi (VMware HA manages the VMs)	0.243	0.081
OpenSAF with VMware Workstation(OpenSAF manages the VMs)	0.848	0.592

Table 4 – Comparison of measurements for failure of the SA-Aware VLC component in different architectures

VM availability management with bare-metal hypervisor

- Intended to fix the delay in service outage.
- Two other VMs called manager VMs added to manage the availability of the service VMs on the shared storage
- Avoiding single point of failure by having 2N redundancy for the manager VMs
- Starting and stopping the VMs the libvirt's "virsh" command
- Service VMs health checking done by external active monitor

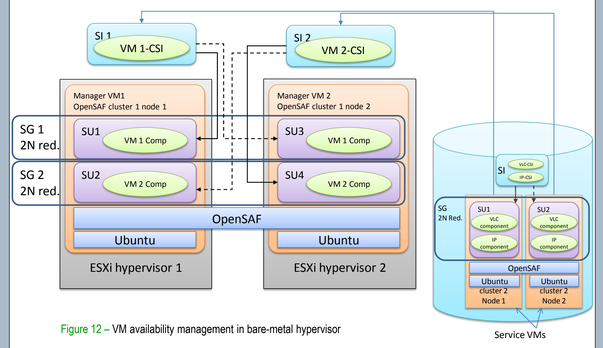

Figure 12 – VM availability management in bare-metal hypervisor

Acknowledgement

This work has been partially supported by Natural Sciences and Engineering Research Council of Canada (NSERC) and Ericsson Research.