

OpenSAF in the Cloud. Why an HA Middleware is still needed

Anders Widell

Ericsson

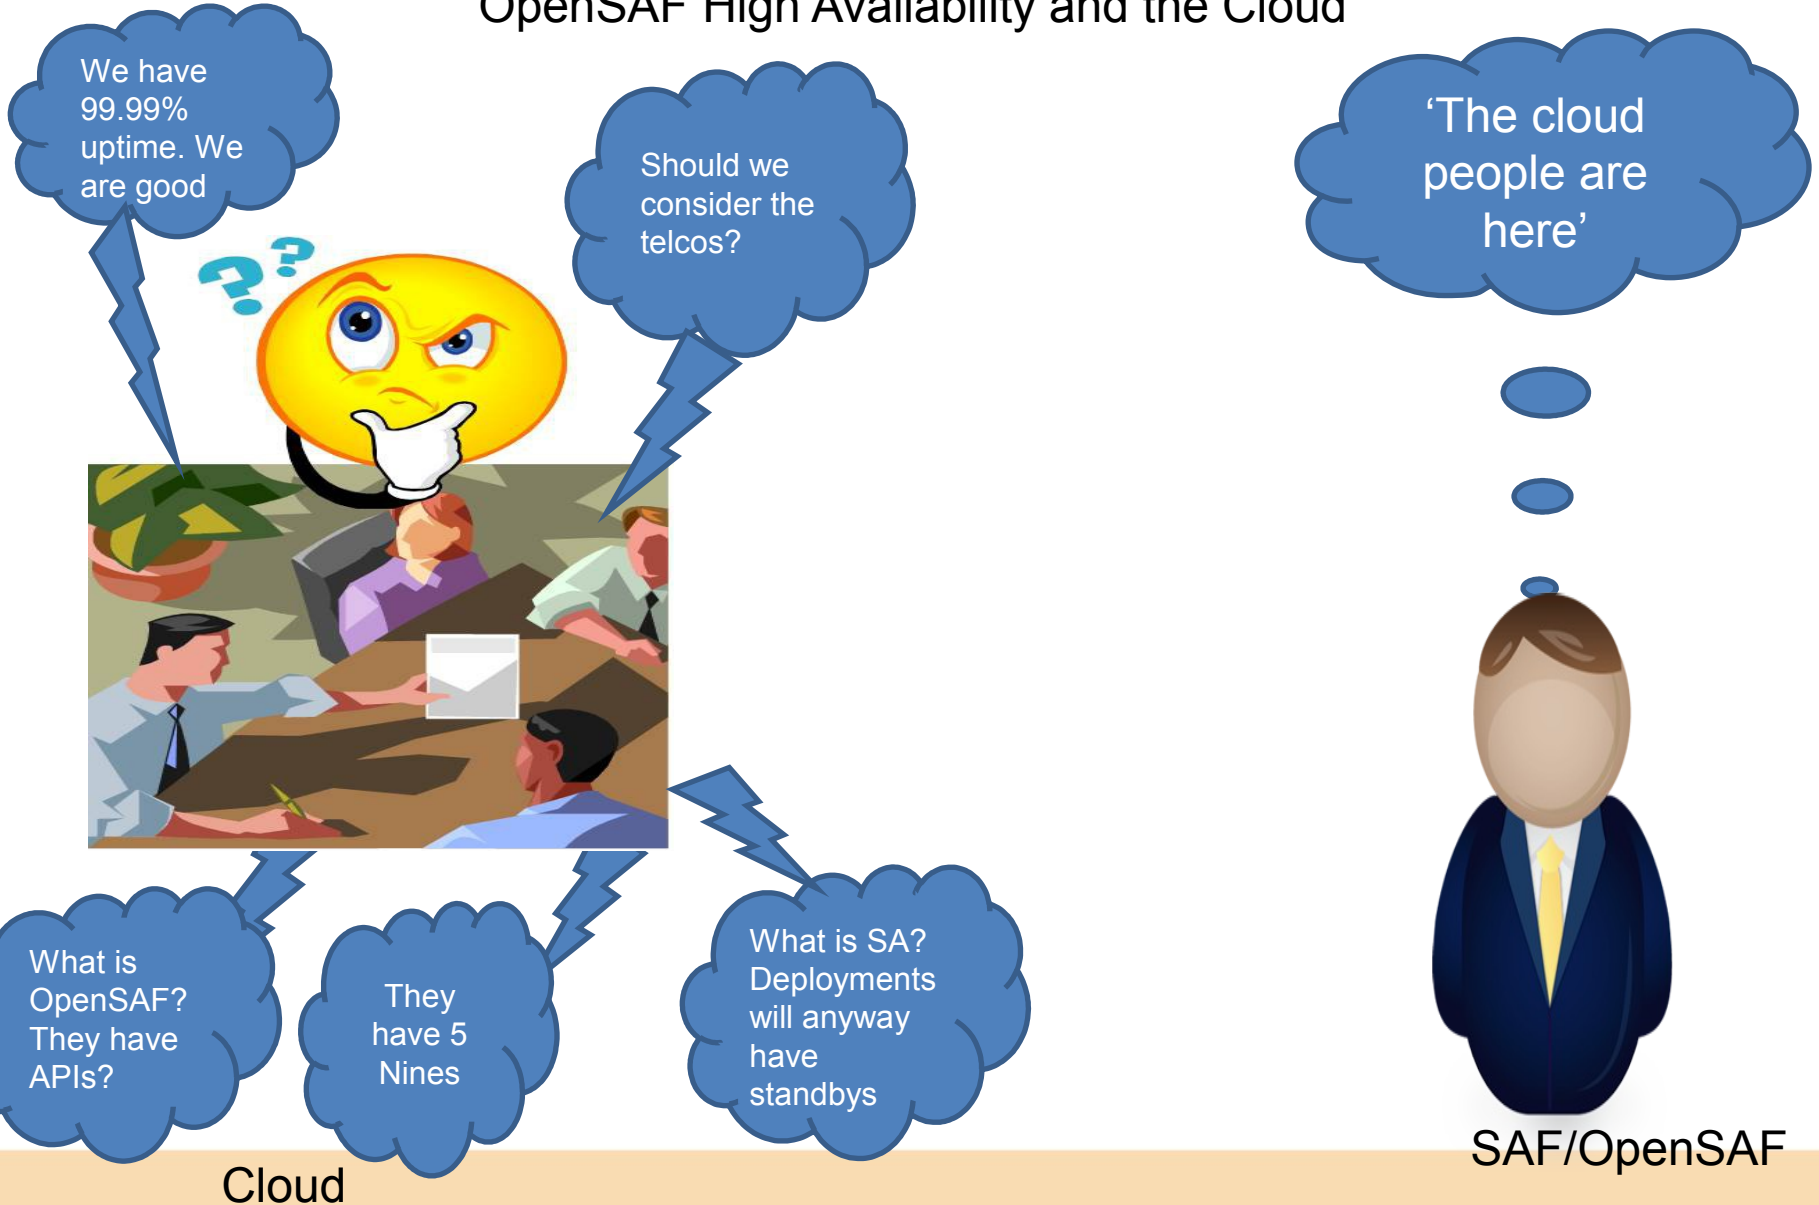
Mathivanan NP

Oracle

Agenda

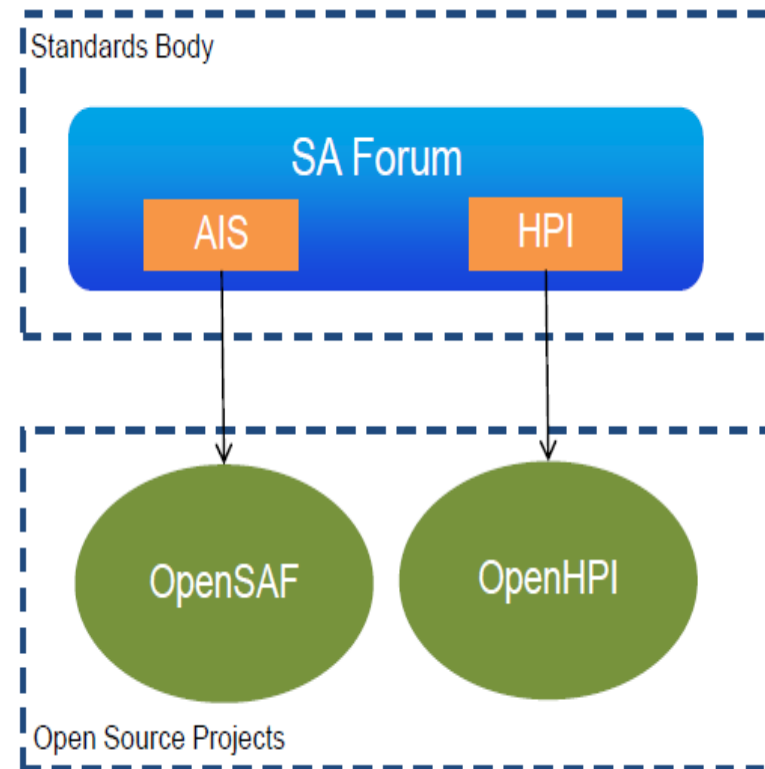
- The OpenSAF Project
- High Availability and Service Availability
- Why Application HA is necessary in the cloud
- OpenSAF HA capabilities
- Proposal to leverage OpenSAF HA with existing cloud solutions for unified availability management
- OpenSAF roadmap

OpenSAF High Availability and the Cloud



The OpenSAF project

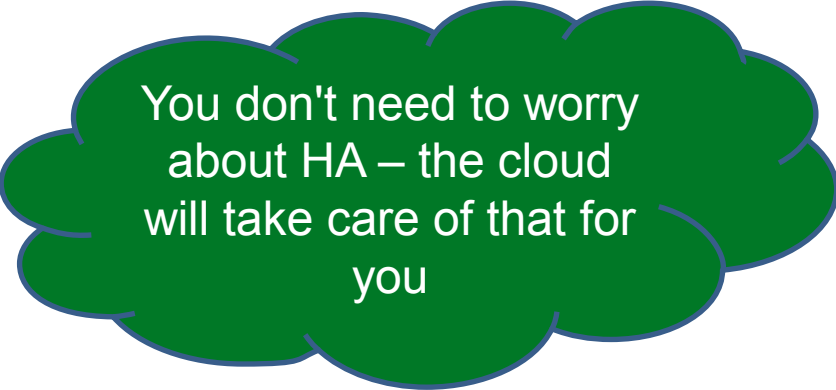
- Most comprehensive Service Availability middleware providing availability, manageability and platform services for developing HA available applications
- Interface APIs in C with support for Java and Python bindings
- LGPL v2.1 license
- Implements SA Forum AIS specification
- Supported by the OpenSAF foundation



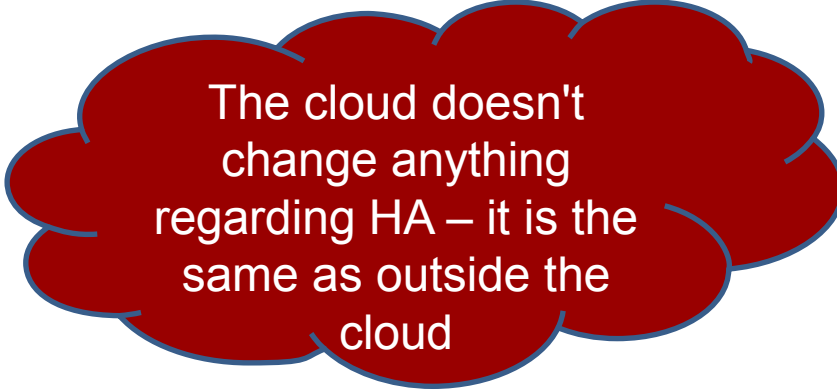
High Availability and Service Availability

- The probability that a service is available to its users at a random point in time
- In telecom, 99.999% availability (five nines) is often required
- HA and SA are essentially the same, but SA enables more – for example planned updates of hardware and software

Two Opinions about Application HA in the Cloud



You don't need to worry about HA – the cloud will take care of that for you



The cloud doesn't change anything regarding HA – it is the same as outside the cloud



High Availability and Service Availability

WALTER'S MOMENTS

by S.A. Forum

"THERE'S NO UPSIDE TO DOWN TIME."



Hardware Faults

- The cloud infrastructure can handle hardware faults for you – all the application sees is a node reboot
- With a hot standby VM, even a reboot may be avoided
- Problem with co-located VMs – we don't want to have active and standby app on the same physical node

Software Faults

- Applications currently have no or limited HA support from cloud infrastructure
- Using HA middleware, we can also get shorter fail-over time in the event of a hardware fault

The Cloud Gives You More Faults

- Hypervisor and cloud infrastructure are also subject to faults
- Hardware used in cloud may be less reliable (not carrier grade)
- Geographic distribution may decrease the risk of total outage, at the cost of network latency and increased risk for split-brain

The cloud way – pets vs. cattle

- Pets: few powerful nodes, scale-up
- Cattle: many cheap nodes, scale-out
- “architecting for failure” vs “architecting for scale”

The cloud way – Standardized Service Level Agreement

Provide service throughout the year

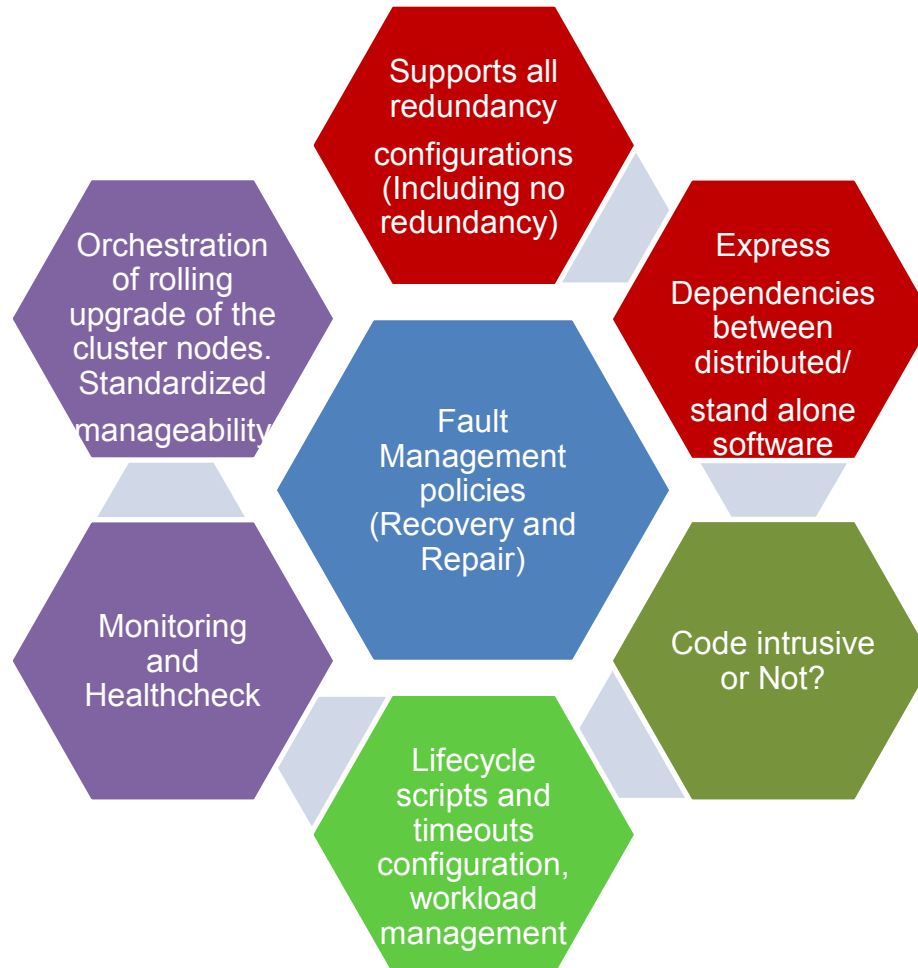
Your problem was triggered by some other vendor/service inside the cloud



OpenSAF based HA

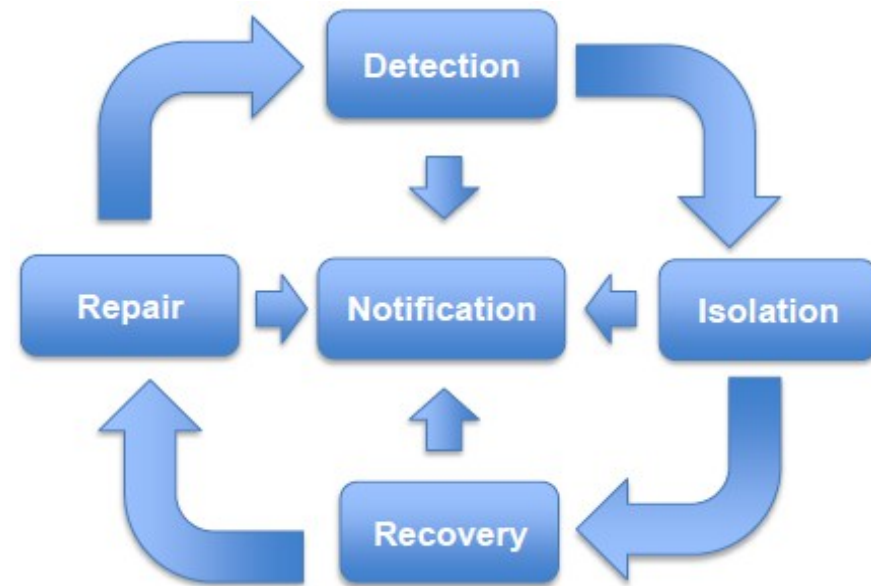
- OpenSAF based HA solutions are applicable across the availability spectrum:
 - Enterprise
 - Telecom and aerospace/defense
 - Millisecond failover

OpenSAF based HA



OpenSAF based HA - Fault Management

- Detection - Component Health Checks, Active/Passive Monitoring, api based error reporting, resource agents
- Isolation - Node Power off or Resource isolation
- Recovery - Failover of role assignments to standby/spare resources
- Repair - Automatic restart of failed resource
- Notifications – Standardized state change notifications (and logging)



OpenSAF HA – Key Advantages

- Provide for Availability as a service in the cloud
- Centralized/Streamlined orchestration of workload management (maintaining affinity)
- Enable cloud software to be more carrier grade
- Ease of Integration – With Both API based and scripts based entities (software, vm, agents, etc)

OpenSAF HA – Key Advantages

- Enables reliability for stateful applications
- Application level failure detection and recovery. Enables fault mitigation and milli second failover
- Support for automated rolling upgrades across the cluster involving application and cluster expansion/shrinking
- Pythonic interface for provisioning, status and management of HA entities. (Java mappings also supported)

Leveraging existing cloud solutions with OpenSAF

OpenSAF and Vmware (A study)

- Outage time measured with/without adding OpenSAF capabilities to existing VMware solutions (FT and HA)
- Outage time measurement by running OpenSAF within and outside the VMs and other combinations
- OpenSAF can detect Hardware, OS and Application failures
- The study concluded that outage time significantly reduced when combining OpenSAF with existing Vmware capabilities

Reference: Ali Nikzad's thesis: 'OpenSAF and Vmware: From the perspective of HA'
http://spectrum.library.concordia.ca/978013/4/Nikzad_MASc_S2014.pdf

Leveraging openstack and OpenSAF

- OpenSAF can provide HighAvailability as a service in openstack – Uniform, centralized, automated availability management across openstack
- Openstack's flexible deployment architectures enables easy integration with OpenSAF for all redundancy configurations for any of the OpenStack infrastructure software (distributed and standalone)
- Monitoring (Intrusive and Non-Intrusive) a basic requirement
 - With/Without Resource agents.
- Provide for a perspective of TRY_AGAIN /TIME_OUT semantics

OpenSAF provides for a Unified HA

Integrated HA architecture
for compute, network,
storage, dashboard



Application HA



VM HA



Unified view
and/of Availability
Management



Provides for openstack
'availability architecture,
hierarchy' and 'standardized
management' (admin, log,
notification, upgrade)
interface

OpenSAF Roadmap

- Enhanced cluster management (quorum/consensus based membership)
- Scaling out even further
- Feature rich CLI
- Container - contained



&
Thank You