NFV ISG PoC Proposal

A.1 NFV ISG PoC Proposal

A.1.1 PoC Team Members

- PoC Project Name: Availability Management with Stateful Fault Tolerance.
- Network Operator/ Service Provider A: AT&T
  Contact: Percy Tarapore, pt5947@att.com; Al Morton, acmorton@att.com
- Network Operator/ Service Provider B: iBasis – part of KPN
  Contact: Richard Xu, RXu@ibasis.net
- Network Operator/ Service Provider C: NTT
  Contact: Eriko Iwasa iwasa.eriko@lab.ntt.co.jp; Atsuyoshi Shirato shirato.atsuyoshi@lab.ntt.co.jp
- Manufacturer A: Stratus Technologies, Inc.
  Contact: Ali Kafel, Ali.Kafel@Stratus.com; Pasi Vaananen, pasi.vaananen@Stratus.com (technical)
- Manufacturer B: Aeroflex
  Contact: Mark Lambe, Mark.Lambe@aeroflex.com
- Manufacturer C: Brocade
  Contact: Yue Chen, cheny@brocade.com
- Manufacturer D: Allot
  Contact: Adi Mendel amendel@allot.com

Above companies will be part of the PoC. For the purposes of this PoC, Stratus will act as the system integrator and is primarily responsible for the coordination of the implementation of the PoC project goals and demonstrations.

A.1.2 PoC Project Goals

This PoC focuses on demonstration of how multiple VNFs from multiple vendors can be deployed in an NFVI+MANO environment that enables deployment, monitoring and control of these VNFs in a variety of availability modes including Fault Tolerant (FT), High Availability (HA) and General Availability (GA).

The salient characteristics of the availability modes in terms of the key fault management cycle phases are as follows:

FT (Fault Tolerant): redundant VM instances, full-VM state checkpointing, NFVI fault detection, fault recovery with fully stateful VM failover to redundant entity. Repair (Redundancy Restoration) through re-instantiation of redundant instance & warm-up (state replication), followed by a move of the new redundant instance to standby state.

HA (High Availability): no VM instance redundancy, no infrastructure provided VM state checkpointing, NFVI fault detection, fault recovery with re-instantiation. Repair is not applicable (recovery & repair operations are the same).

GA (General Availability): no VM instance redundancy, no no infrastructure provided VM state checkpointing, NFVI fault detection, fault recovery through clean up only (i.e. no automatic recovery, will require external

intervention). Repair is not applicable (requires either external SW agent request or manual intervention to re-instantiate).

The list of the specific, enumerated PoC goals is given below.

- PoC Project Goal #1: Demonstrate instantiation of VNFs from multiple vendors onto a simple static NFV Forwarding Graph (emulating the configuration that is analogous to instantiation of corresponding physical NFs onto functionally equivalent service chain).
- PoC Project Goal #2: Demonstrate the correct (i.e. constraints compliant) placement of redundant VNF (or VNFc) instances based on specified deployment anti-affinity attributes by MANO/VIM services.
- PoC Project Goal #3: Demonstrate simultaneous instantiation of multiple VNFs in a mix of three different availability modes: FT, HA and GA (as described in the introductory portion of this section).
- PoC Project Goal #4: Demonstrate the feasibility of support for VNF sparing / redundancy and state checkpointing as NFVI + MANO provided services (transparent state checkpointing for FT availability mode only, storage-based state replication for any availability mode).
- PoC Project Goal #5: Demonstrate automated detection and recovery of failed VNFs, including associated NFVI network reconfiguration mechanisms as NFVI + MANO services (for FT and HA availability modes).
- PoC Project Goal #6: Quantify the service availability performance metrics demonstrating both service accessibility and service continuity aspects for each of the different availability modes (downtime performance metrics per event, measured from service client's perspective).
- PoC Project Goal #7: Demonstrate the qualitative service availability performance aspects with different availability modes through the use of suitable end-to-end applications over the instantiated service chains upon injection of failures (e.g. failure event impact to streaming server-client audio/video application for each availability mode).

A.1.3 PoC Demonstration

Examples include: PoC Team's member's labs, industry trade shows, research networks, etc.

- Venue for the demonstration of the PoC: NFV World Congress May 6-8 in San Jose, CA

A.1.4 (optional) Publication

- None targeted at this time beyond the PoC report.

A.1.5 PoC Project Timeline

- What is the PoC start date? March 2015
- (First) Demonstration target date May 6-8th, 2015 (NFV World Congress, San Jose, CA)
- PoC Report target date July 31th, 2015

When is the PoC considered completed? After the PoC demonstration has been publicly executed, associated availability performance data has been captured and the PoC Report has been issued.

A.2 NFV PoC Technical Details (optional)

A.2.1 PoC Overview

The focus of the PoC is on demonstration of the NFVI/MANO implementation of the high availability mechanisms and assessing the key availability performance metrics of such mechanisms. Fault Tolerant availability mode can be considered an example of the implementation where all phases (HA aware function placement, application state protection, fault detection, fault recovery and repair) of the high availability related functions are supplied by the infrastructure services. It should be understood that while PoC FT availability mode combines all of these mechanisms to achieve the highest availability levels, it is not necessary that all the FM related functionality for VNFs are provided by the NFV platform services, but the underlying mechanisms can also be provided as their individual constituent component availability services. PoC FT mode will demonstrate the feasibility of providing all of the availability related mechanisms as a set of services provided by the infrastructure functions without any application awareness, which implicitly demonstrates that any of the required underlying constituent mechanisms can be provided as an infrastructure level service.

We intend to demonstrate two sets of resiliency and availability management lifecycle related scenarios in this POC. They are categorized as follows:

Category 1 – Deployment related lifecycle functions, including description and instantiation of the VNFs in NFVI through the VNFM enabled VNF Descriptors (VNFD) in multiple availability modes (FT, HA and GA, as described in section 1.2 of this document).

Category 2 – Post deployment use case lifecycle demonstration of VNFs with multiple availability modes, including Stateful FT, High Availability (HA) and General Availability (GA). Fully Stateful application fault tolerance (FT) of VNFs means that associated VM state is always redundant such that a failure to one of the nodes will not have negative impact on the QoS or SLA on the service function. HA and GA modes will not perform full VM state replication, and therefore associated service availability impact of faults in these modes is expected to be more intrusive as compared to FT mode.

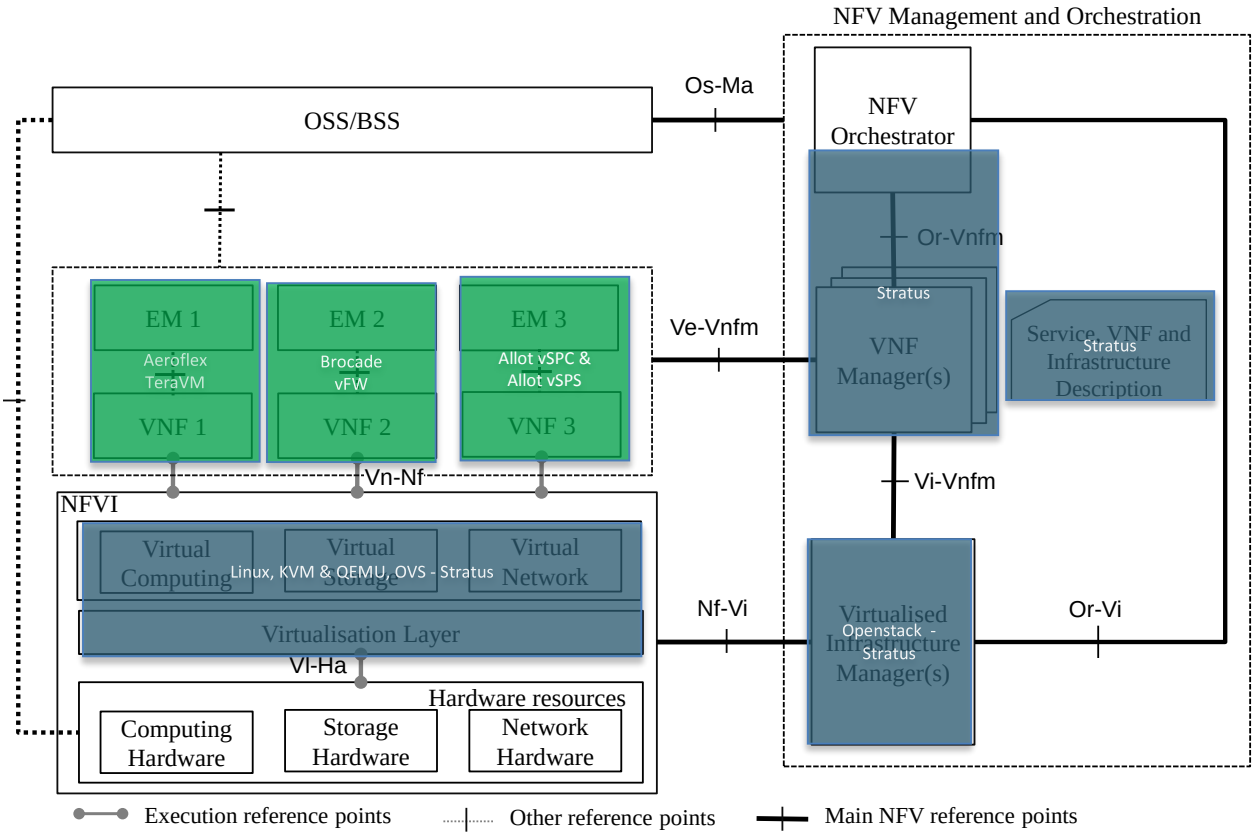

The scope of the PoC demonstration is single NFVI domain (or "NFVI PoP"), which is considered sufficient to demonstrate the functional and performance aspects relevant to the goals of this PoC, and therefore OSS/BSS functionality and multi-domain orchestration elements of the ETSI NFV reference architecture are not required and will not be present in the PoC configuration. In the absence of the full orchestrator, the configuration of the VNFFGs is accomplished through the use of the GUI that exercises the functionality of the underlying software entities, specifically VNFM and VIM through their exposed northbound interface APIs. Effectively, this GUI application acts as a proxy for the full orchestration function, allowing the configuration of the PoC VNF Forwarding Graph connectivity to support the PoC goals as detailed in this document.

The VNFs that will be demonstrated as part of this multi-vendor PoC will be Brocade virtual router / firewall (vFW) configured as stateful inspection firewall VNF and Allot virtual service protection system composed of two VNFs (virtual Service Protection Sensor vSPS and virtual Service Protection Controller vSPC). Aeroflex TeraVM virtual test infrastructure is used as a primary test tool for the datapath functionality provided by the vFW and service protector functions.

Stratus will provide infrastructure SW elements for the PoC, including physical HW (3rd party commodity compute, storage and network hardware), NFVI infrastructure software, OpenStack based VIM, VNFM and support for various descriptors required for the support of lifecycle management of the 3rd party, multi-vendor VNFs and their constituent components (VNFCs) in different availability mode configurations.

Aeroflex TeraVM is a virtualized test product for assessing performance of the wide variety of network elements and protocol implementations. TeraVM will be used primarily to measure the availability related service performance metrics of the VNFs listed below, when instantiated in the different availability modes.

Brocade Vyatta vRouter is a virtual implementation of the comprehensive set of IPv4/IPv6 router functionality and associated protocols for virtualized environments, packaged as VNF. In addition to the router functionality, Brocade Vyatta vRouter VNF also includes the rich set of security functions, including VPN gateway and stateful inspection firewall functionality which are both used in this PoC to demonstrate the secure session-stateful datapath functions.

Allot virtual Service Protection Sensor (vSPS) and virtual Service Protection Controller (vSPC) VNFs provide the defense against Distributed Denial of Service (DDoS) attacks, outbound spam prevention and botnet host containment and cleanup services. vSPS analyses traffic for threat signatures for real-time detection of the threats, and enforces the security threat mitigation rules using packet filtering rules under the control of the virtual Service Protection Controller (vSPC). vSPC monitors the threat detection event indications from vSPS's, determines the enforcement actions based on these incoming events and communicates the associates enforcement rules to vSPSs.

The depiction of the associated functionality in fully stateful fault tolerant configuration is given in the figure below. The green arrows in the figure depict the state replication mechanisms between the redundant virtual machine pairs.

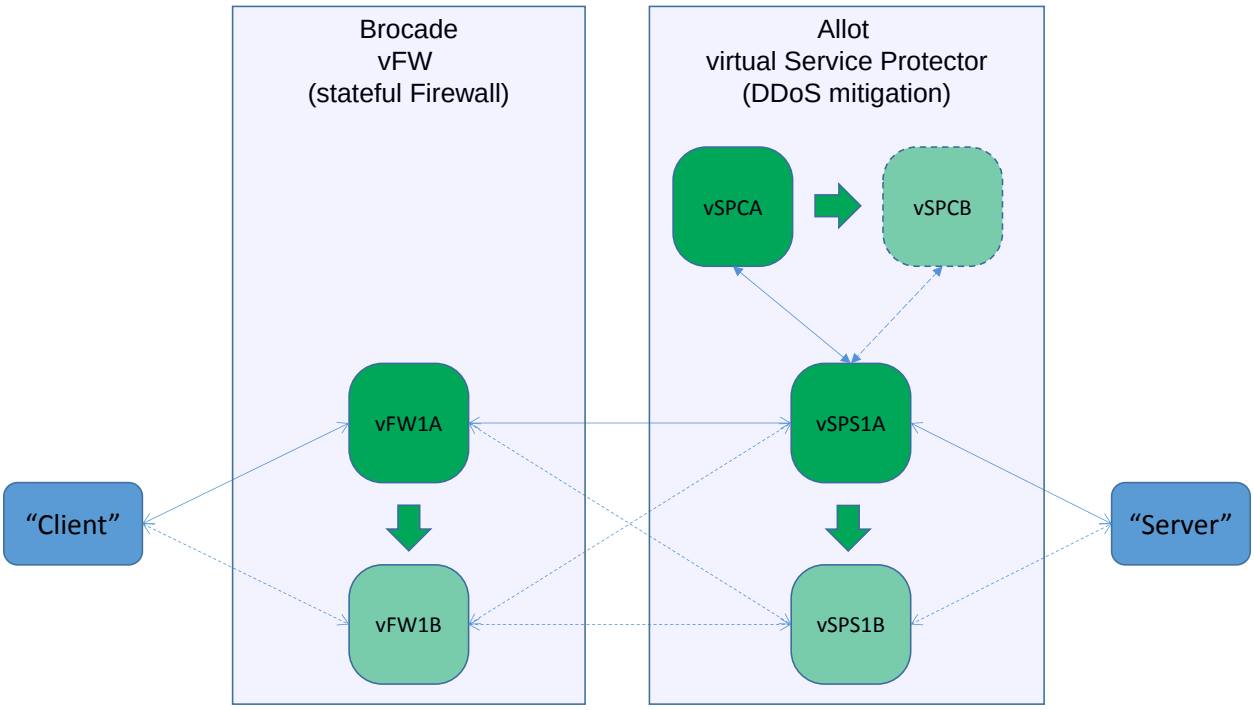

Figure below shows the logical topology of the PoC VNF (redundant) datapath configuration and associated external entities (client and server).

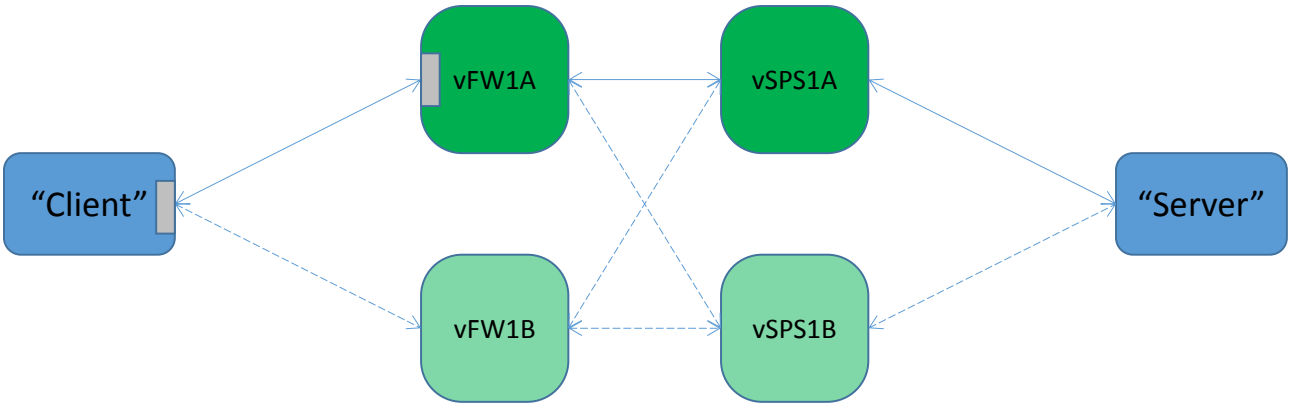

The service chain, which is implemented in the NFVI is composed of two redundant virtual security application instances in series. "vFW1" is stateful inspection firewall network function, and "vSPS1" is DDoS detection/prevention network function, which is also session-stateful. vFW1 instance also provides VPN session termination / origination VPN access point functionality. Together these applications will be sufficient to

demonstrate the availability performance impacts of various redundancy mechanisms to the session-stateful applications. A and B suffixes represent the redundant entities on pair-wise redundancy deployment scenario.

The expected behavior for the stateful inspection firewall and associated VPN termination point with respect to the sessions would be that upon failover without session state protection, sessions (VPN and TCP, respectively) needs to be re-established. The equivalent behavior is expected in the vSPS VNF for TCP sessions. Since "stand-alone" UDP sessions are session-stateless by nature (UDP flows are generally considered to be stateless uni-directional flows), TCP & VPN sessions will be used primarily to monitor the service accessibility (i.e. session re-establishment is required/not required for sessions that were established prior to failure and new session establishment works/does not work). Client-Server UDP round-trip sessions within VPN session are respectively used to monitor the service continuity, where VPN provides a stateful session wrapping between the security association endpoints (i.e. client-vFW1 VPN service access point, which is implemented within the firewall VNF).

Figure below illustrates the key communications sessions from the client's perspective in the above logical topology. Control and management plane sessions are not shown for clarity, but are assumed to exist as required for the associated functionality to work (esp. for the VPN client-server functional pair). UDP session is assumed to be simple client-server session where a simple server application is expected to "turn around" each of the received PDUs for return to client application through separate associated return UDP session.

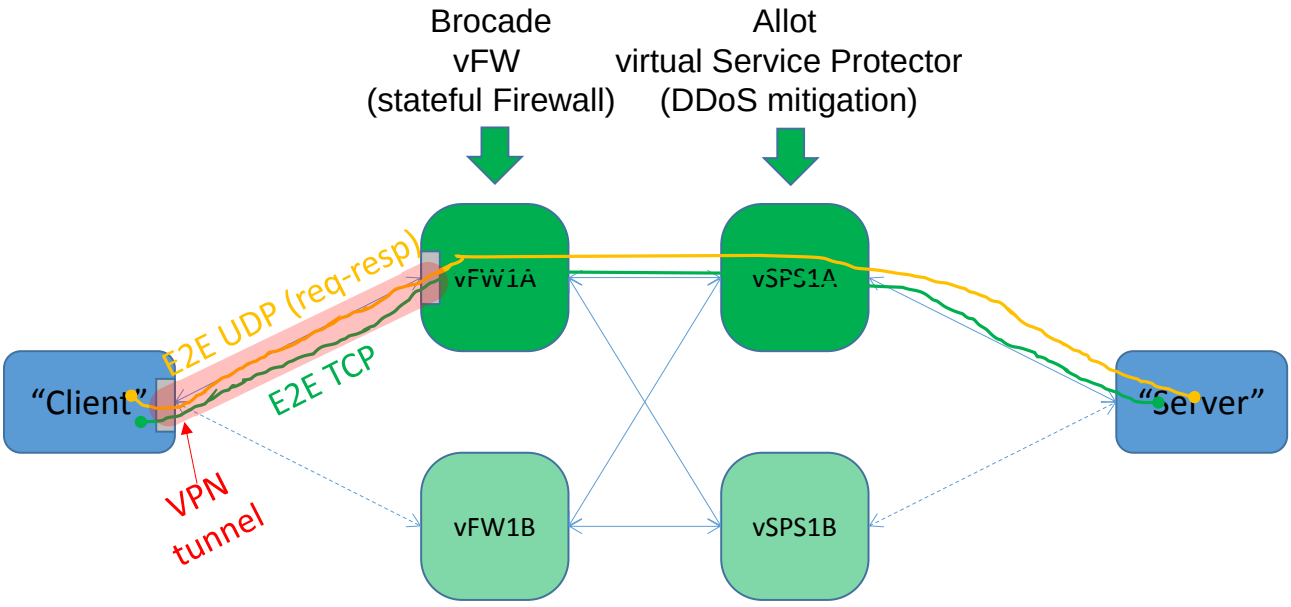

Since the availability related metric are monitored and collected from the Client viewpoint, these metrics have certain limitations, as follows:

a.) the client cannot distinguish what component (or components) within the service chain are failed, and

b.) for simplicity, monitoring is done primarily from client perspective, and therefore no detailed direction specific performance info can be obtained.

Despite of the limitations, the collected service availability performance metrics are expected to be sufficient to characterize the performance and behavioral differences on fault related incidents in different configurations as seen by client, and as such could be used as inputs for estimation of resulting availability performance characteristics in more complex topologies and configurations.

Client and server entities are considered to be external to the NFVI and their availability performance is not a specific target of neither improvement or monitoring in the context of this PoC (although if desired, one or both of these entities could be instantiated within the NFVI with or without redundancy, which would potentially be particularly relevant to server entity, as such instantiation would be considered common in various real-life deployment scenarios).

A high level **example** of the POC mapped onto minimalistic physical infrastructure that can be used to demonstrate the associated availability performance metrics is depicted in figure below. VIM (OpenStack) and other infrastructure control related nodes (e.g. VIM, VNFM, storage nodes, etc.) and associated networks are not shown here for clarity, but they will be present in the PoC configuration.

In this example, the VNFs with differing availability modes (and their constituent VNFCs) are mapped to the same hardware entity (i.e. compute node). This is not necessary from the underlying infrastructure point of view, but for the demonstration purposes this allows simultaneous fault injection (e.g. hypervisor / server fault) affecting the VNFs with differing availability modes and simultaneous observation of the associated end system / application behaviors through VNFs with differing availability modes.

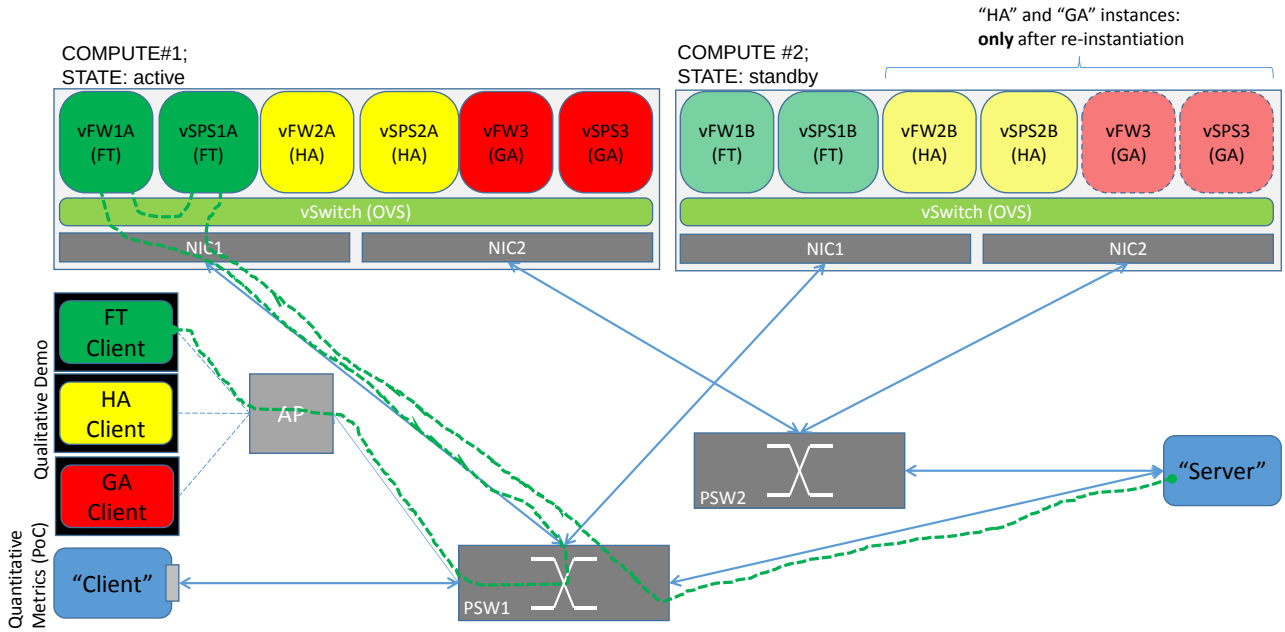

It is not a goal to demonstrate availability improvement of the server component in this PoC, but Server **may** also be instantiated as VNF if desired (and at any availability level if virtualized). The decision on whether to use physical server or virtualized server function will be left to be made in implementation phase. Likewise, it is not the goal of this PoC to demonstrate availability improvements for the Aeroflex virtualized test solution, but this is used as virtual test infrastructure to help characterize the availability performance metrics of the datapath functionality instantiated in the different availability modes.

A.2.2 PoC Scenarios

2.2.1 Category 1 (deployment) Scenarios

In these scenarios we demonstrate how we can reduce human errors and complexity while providing overall availability management that automates the deployment, monitoring and control of applications running in Fault-Tolerant (FT), High Availability (HA) and/or General Availability (GA) modes.

Scenario #1.1 – VNF description at different availability levels

In this scenario, VNF descriptions are created to describe the VNF on-boarding related requirements and constraints, including Availability Modes. This involves generation of the VNFD which includes the information required to onboard the VNFs and their constituent components in the instantiation phase (see following scenario).

VNFD's that are otherwise equivalent in functionality will be created for the FT, HA and GA availability modes. In the instantiation step (next scenario) these VNFDs are used to instantiate VNFs on the NFVI resources.

Scenario #1.2 – VNF instantiation at different availability levels

In this scenario, VNFD's created in the scenario #1.1 above are used to place and start the VNFs (and their associated services) within NFVI infrastructure. The VNFs are placed according to their deployment requirements, including the constraints and requirements associated with the support of the availability modes.

Scenario #1.3 – VNF configuration and configuration store to externalized persistent storage

In this scenario, VNF instances are configured sufficiently to provide the associated VNF service (i.e. vFW or vSPS) through its VNF specific element management interface(s). VNF configuration is written to persistent storage, which is associated with the VNF instance, and this configuration information is expected to be available and used when the VNF is re-started or re-instantiated (i.e. restart should not require re-configuration). This demonstrates partial VNF state (i.e. configuration / management state only) persistence through externalized state storage (in disk).

2.2.2 Category 2 (mixed availability) Scenarios

Scenario #2.1– Resiliency of the network service in a Fault Tolerant (FT) VNF deployment model with fully stateful VNF VM redundancy. This means keeping a duplicate copy of full virtual machine state (and implicitly all session state) so that if a failure happens, the secondary VNFC will have a replica of the full VM state and will continue providing the service as if a failure never happened.

- In this scenario two VNFs are deployed in FT mode and associated VNF VM state is maintained at all times
- Resiliency is demonstrated in three different scenarios
    a. Kill the individual VM process (VNFs/VNFCs)
    b. Kill the OS and/or hypervisor (all instantiated VNFs/VNFCs)
    c. Kill the hardware server (all instantiated VNFs/VNFCs)
- Availability metrics (accessibility and continuity metrics) are automatically monitored by test application through the injected fault scenarios and associated performance parameters are recorded.
- For qualitative availability performance monitoring, the quality of the application performance is monitored and any observed performance impacts (including no noticeable impact) are recorded.

Scenario# 2.2 – Service resiliency upon VNF failure using hypervisor HA. When the failure is detected we immediately respin the VM. This is defined as Hypervisor High Availability (HA)

- In this scenario two VNFs are deployed in HA mode, without state replication (recovery is based on restart)
- Resiliency is demonstrated in three different scenarios
    a. Kill the individual VM process (VNFs/VNFCs)
    b. Kill the OS and/or hypervisor (all instantiated VNFs/VNFCs)
    c. Kill the hardware server (all instantiated VNFs/VNFCs)
- Availability metrics (accessibility and continuity metrics) are automatically monitored by test application through the injected fault scenarios and associated performance parameters are recorded.
- For qualitative availability performance monitoring, the quality of the application performance is monitored and any observed performance impacts (including no noticeable impact) are recorded.

Scenario# 2.3 – Service resiliency upon VNF failure in GA mode. When the failure is detected, operator is notified but operator (or NFVI external agent) intervention is required to respin the VM to restore the service. This is defined as General Availability (GA)

- In this scenario two VNFs are deployed in GA mode, without state replication (recovery is based on restart).
- Resiliency is demonstrated in three different scenarios
    a. Kill the individual VM process (VNFs/VNFCs)
    b. Kill the OS and/or hypervisor (all instantiated VNFs/VNFCs)
    c. Kill the hardware server (all instantiated VNFs/VNFCs)
- Notification of the failure to operator is demonstrated, followed by the use of MANO interfaces to re-instantiate the failed VNF to restore service.
- Availability metrics (accessibility and continuity metrics) are automatically monitored by test application through the injected fault scenarios and associated performance parameters are recorded.
- For qualitative availability performance monitoring, the quality of the application performance is monitored and any observed performance impacts (including no noticeable impact) are recorded.

2.2.3 End-to-end service availability performance metrics

For simplicity, the service availability performance is monitored externally from the "client" endpoint, and therefore is generally representative of the whole service chain availability performance, rather than the availability performance of the independent functions on the service chain (i.e. endpoint cannot be fully aware of the nature of the constituent service components causing the service outage or degradation).

For quantitative service availability performance monitoring, we will use separate mechanisms to demonstrate the service availability and service continuity metrics.

For service accessibility metrics, we use IPSEC and TCP sessions to demonstrate the impact of the "session stateful" applications for the service accessibility. This means that services are accessible when the client is able to establish VPN session with the associated VNF for the infrastructure VPN service component in the service chain, and for TCP sessions the service (and implicitly the whole service chain) is accessible when the TCP session can be established between the client and server. The associated performance metrics are session establishment time, and service is considered to be un-available when the session establishment is not successful within typical timeframe (established by monitoring the times in absence of failures), or when the service times out.

For the service continuity metrics, we use sufficiently high packet rate UDP sessions between the client and server to quantify the service continuity performance. These UDP sessions are tunneled over the VPN tunnels, making them stateful from the perspective of the associated VNFs in the service chain. The reason for using UDP is that it allows us simpler and more accurate measurement of the service continuity metrics than TCP sessions due to non-interference from the congestion control mechanisms of TCP.

The observed metrics for the service accessibility are time for the successful VPN and TCP session establishment, and associated impairments upon the simulated failures of the components in the service chain.

The observed metrics for the service continuity are the round-trip time of the UDP based "ping" style interactions between the client and server, and associated impairments upon the simulated failures of the components in the service chain. The key parameter to be monitored is the observance length of the outage associated with failure. We will also monitor and report the packet loss rates through the chain using the simple sequence number based scheme.

The same metrics will be used for monitoring of all availability configuration scenarios, allowing direct comparison of the associated results between them.

Figure below illustrates the scope of the availability related metrics in the context of the proposed PoC logical topology and their high level dependencies.

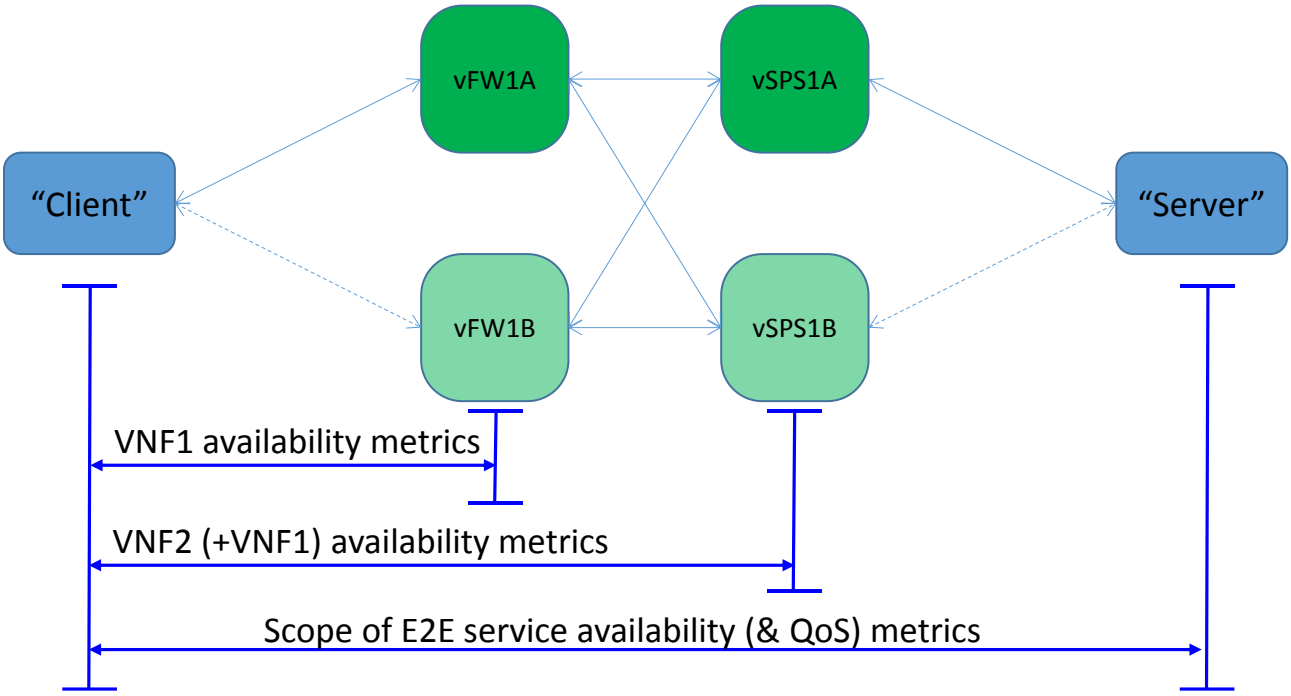

We will also be able to demonstrate qualitative impacts with simulated applications, such as web and/or video server and associated clients. This will allow observation of the impact for the ongoing sessions on per availability mode level. Association of the different availability mode functionality to same physical HW entity (as discussed in "Physical Topology" section) allows simple injection of simultaneous faults for each parallel availability mode path.

A.2.3 Mapping to NFV ISG Work

Key aspects of this PoC as well as the relevant key demonstration objectives are related to the NFV reliability work. Specifically, the objectives related to NFV REL are quantitative characterization of the service availability metrics from the end-user application perspective for different types of VNF availability categories, and include baseline metrics for both service accessibility and service continuity in the context of the provided service (which is essentially a form of secured IP network access service in the context of this PoC).

As a new capability, the PoC will demonstrate the feasibility of the implementation of the key high-availability mechanisms as NFV Infrastructure provided services (i.e. without specific availability awareness and state checkpointing functionality implemented in application level), which can be used instead of or as a complementary set of mechanisms to application level availability implementations. Because the PoC demonstrates the feasibility of arguably the most demanding possible availability mode (i.e. fully stateful VM Fault Tolerance with infrastructure provided state protection and rapid recovery), this gives a good indication that all phases of the fault management cycle from detection to repair (and including state replication) can be candidates for implementation with NFVI/MANO level mechanisms.

In the context of the MANO, PoC demonstrates enhanced NFVI level deployment capabilities and implementation of service resiliency mechanisms beyond the capabilities implemented by baseline OpenStack based NFV/MANO services.

PoC also necessarily implicitly incorporates and demonstrates many key aspects of NFVI as required for its implementation, however specific measurement and implementation related objectives and discussions with respect

to these other NFVI aspects are limited to mechanisms and services required to support availability improvement and demonstration related objectives of this PoC only. For example, while the application example VNFs used in this PoC are mostly related to security, it is not objective of this PoC to focus on demonstration of security related aspects of functionality.

The key NFV ISG end-to-end mechanisms from the NFV Use Cases, Requirements, and Architectural Framework functional blocks or reference points which are addressed by the different PoC scenarios are outlined in the table below. While the implementation demonstrates and utilizes various other mechanisms and requirements of the NFV, only the items related to resiliency mechanisms are listed below for brevity and emphasize the specific focus of this PoC project intended to be on such mechanisms. The referenced documents in the table are in brackets [], with associated specific document numbers and versions listed in the section A.2.6 (References) of this document.

	Use Case	Requirement	E2E Arch	Comments
Scenario 1.1 (descriptions)	*[NFV-UC] UC#1 (NFVIaaS)*	[NFV-REQ] Res.1, Res.2, Res.3, Res.4 [NFV-REL] Req.4.2.2, REQ 4.2.3, REQ 4.2.4, REQ 4.2.15, REQ 5.4.11, REQ 7.3.1	VNFDs, VNFM, VIM and NFVI	Availability mode descriptions for VNFs/VNFc's
Scenario 1.2 (instantiation)	*[NFV-UC] UC#1 (NFVIaaS)*	[NFV-REQ] Res.1, Res.2, Res.3, Res.4 [NFV-REL] Req.4.2.2, REQ 4.2.3, REQ 4.2.4, REQ 4.2.15, REQ 5.4.11, Behaviour.2, REQ 7.3.1, REQ 9.5.1, REQ 9.5.2, REQ 9.5.3	VNFs/VNFCs, VNFM, VIM and NFVI	Availability mode aware instantiation of VNFC's
Scenario 1.3 (externalized storage)	*[NFV-UC] UC#1 (NFVIaaS)*	[NFV-REQ] Res.1, Res.2, Res.3, Res.44 [NFV-REL] Req.4.2.9	VNFs/VNFCs, VNFM, VIM and NFVI	Support externalized state storage via persistent disk based state storage services (can be used for management & configuration as well as more dynamic i.e. signaled state). This is aligned with a set of mechanisms required to support "VNFC w/ externalized state", as described in [SWA], section 5.1.2
Scenario 2.1 (FT stateful recovery)	*[NFV-UC] UC#1 (NFVIaaS)*	[NFV-REQ] Res.1, Res.2, Res.3, Res.4 4 [NFV-REL] Req.4.2.2, REQ 4.2.3, REQ 4.2.4, REQ 4.2.4, REQ.4.2.6, REQ 4.2.7, REQ 4.2.12, REQ 5.4.1, REQ 5.4.2, REQ 5.4.4, REQ 5.4.6, REQ 5.4.9, Behaviour.2, REQ 7.3.7, REQ 10.8.1, REQ 10.8.4, REQ 10.8.5, REQ 10.8.16	VNFs/VNFCs, VNFM, VIM and NFVI	Fully stateful & fast VM (VNFC) recovery provided as infrastructure service by NFVI+MANO entities
Scenario 2.2 (HA re-instantiation recovery)	*[NFV-UC] UC#1 (NFVIaaS)*	[NFV-REQ] Res.1, Res.2, Res.3, Res.4[NFV-REL] Req.4.2.2, REQ 4.2.3, REQ 4.2.4, REQ 4.2.4, REQ.4.2.6, REQ 4.2.7, REQ 4.2.12, REQ 5.4.1, REQ 5.4.2, REQ 5.4.4,	VNFs/VNFCs, VNFM, VIM and NFVI	Re-instantiation based stateless or externalized state based VM (VNFC) recovery provided as infrastructure

		REQ 5.4.6, Behaviour.2, REQ 7.3.7, REQ 10.8.1, REQ 10.8.4, REQ 10.8.5, REQ 10.8.16		service by NFVI+MANO entities
Scenario 2.3 (GA, standard availability)	*[NFV-UC] UC#1 (NFVIaaS)*	[NFV-REQ] Res.1, Res.2, Res.3, Res.4 [NFV-REL] Req.4.2.2, REQ 4.2.3, REQ 4.2.4, REQ 4.2.4, REQ.4.2.6, REQ 4.2.7, REQ 10.8.1, REQ 10.8.4, REQ 10.8.5	VNFs/VNFCs, VNFM, VIM and NFVI	Clean-up on failure of state based VM (VNFC) provided as infrastructure service by NFVI+MANO entities; recovery actions need to be taken either manually or through external SW agents (e.g.. VNF or Orchestrator level mechanisms).

A.2.4 PoC Success Criteria

This proof-of-concept is considered to be successful when all of the objectives stated in section 1.2 of this document have been successfully completed and demonstrated by executing the associated test scenarios as outlined in section 2.2, and the findings and measured availability performance metrics have been documented and published in the PoC report.

A.2.5 Expected PoC Contribution

One of the intended goals of the NFV PoC activity is to support the various groups within the NFV ISG. The PoC Team expects to submit contributions relevant to the NFV ISG work items as a result of the PoC Project, particularly towards the ongoing work items in REL and IFA working groups.

Areas of contributions towards specific NFV ISG WIs which are expected to result from this PoC Project are listed below.

- PoC Project Contribution #1: confirmation of the feasibility of implementation of various NFV REL mechanisms and interfaces established in ETSI NFV ISG phase 1 documents. Establish and document the baseline performance values for quantitative availability performance service accessibility and service continuity metrics of the different infrastructure provided availability mechanisms. Quantitative performance data is expected to be useful for realistic availability performance modelling work. NFV WI DGS/REL-003.

- PoC Project Contribution #2: Demonstrate and document the set of availability methods and associated underlying mechanisms in NFV REL Work Item REL003, availability mechanisms and end-to-end availability modelling. NFV WI DGS/REL-003.

We also plan to determine the specific additional contribution opportunities as the PoC and NFV ISG REL / IFA work progresses, especially for items necessary to enable support for the specific availability related mechanisms through reference interfaces and associated data models being considered in IFA work items. However, it is too early to positively identify all the specific / detailed IFA contribution opportunities at this stage, before the majority of the PoC work has been completed.

Currently ongoing IFA WIs which are relevant for the availability / fault management functionality for related architecture and functionality include NFV WI DGS/IFA-009 and NFV WI DGS/IFA-010, and for interfaces and data elements NFV WI DGS/IFA-005, NFV WI DGS/IFA-006, NFV WI DGS/IFA-007, and NFV WI DGS/IFA-011. PoC team prefers to contribute to IFA work items through REL working group joint contributions.

A.2.6 References

[NFV-UC] ETSI GS NFV 001; Network Functions Virtualization (NFV); Use Cases; V1.1.1 (2013-10):
http://www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf

[NFV-REQ] ETSI GS NFV 004; NFV Virtualization Requirements; V1.1.2 (2013-10):
http://www.etsi.org/deliver/etsi_gs/NFV/001_099/004/01.01.01_60/gs_NFV004v010101p.pdf

[NFV-REL] ETSI GS NFV-REL 001; NFV Resiliency Requirements;V1.1.1 (2015-01):
http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV-REL001v010101p%20-%20Resiliency%20Requirements.pdf

[NFV-SWA] ETSI GS NFV-SWA 001; NFV Virtual Network Functions Architecture; V1.1.1 (2014-12):
http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV-SWA001v010101p%20-%20Virtual%20Network%20Functions%20Architecture.pdf