LIAISON STATEMENT

Title: **NFV Requirements**

Date: 2014-09-23

From (source): ETSI ISG NFV

Contact(s): Michael Brenner, Michael.Brenner@alcatel-lucent.com
Jan Ellsberger, Jan.Ellsberger@ericssoncom
NFVSupport@etsi.org

To: OpenStack

Copy to:

Response to: Michael Brenner, Michael.Brenner@alcatel-lucent.com
Jan Ellsberger, Jan.Ellsberger@ericssoncom
NFVSupport@etsi.org

Attachments: None.

ABSTRACT: *OpenStack is identified as one main component in the NFV architectural framework. This liaison statement informs OpenStack generally about NFV requirements, together with an initial list of requirements currently not supported by OpenStack.*

1. Overall description:

ETSI ISG NFV is securing industry consensus on business and technical requirements for NFV, and common approaches to meeting these requirements. The outputs are openly published and shared with relevant standards bodies, industry fora and consortia to encourage a wider collaborative effort. The NFV ISG intends to collaborate with other SDOs if any standardization or other approaches to encourage broad adoption are deemed necessary to meet the requirements – since one of the major ETSI NFV goals is interoperability. The NFV ISG also provides an environment for the industry to collaborate on Proof of Concept (PoC) platforms to demonstrate solutions which address the technical challenges for NFV implementation and to encourage growth of an open ecosystem.

OpenStack is invited to take note of the following initial set of NSG NFV documents:

- NFV Use Cases
  document describes initial fields of application. The uses cases have been selected to span the scope of technical challenges being addressed by the NFV ISG, it is not meant to be an exhaustive list.
- NFV Requirements document describes the high level business and technical requirementsfor an NFV framework including service models.
- NFV Architectural Framework document describes the high-level functional architecture and design philosophy for virtualised network functions and the underlying virtualization infrastructure. By delineating the different constituents and outlining the reference points between them, it paves the way for fully interoperable multi-party NFV solutions.
- NFV Terminology document is a common repository for terms used within the NFV ISG documents.
- NFV Proof of Concept Framework document. The NFV ISG has launched a global call for multi-party NFV Proof of Concepts (PoC) to validate NFV approaches and to encourage progress towards interoperability and development of an open ecosystem.
- NFV Performance and Portability Best Practices describes a list of features which the performance and portability templates (Virtual Machine Descriptor and Compute Host Descriptor) should contain for the appropriate deployment of Virtual Machines over a Compute Host (i.e. a "telco datacentre").
- NFV Management and Orchestration (draft) describes the management and orchestration framework required for the provisioning of virtualised network functions (VNF), and the related operations, such as the configuration of the virtualised network functions and the infrastructure these functions run on.
- Other documents may be made available as needed in future

ETSI ISG NFV has analyzed the OpenStack capabilities; while a large number of NFV requirements can be met with currently available OpenStack features, a subset of NFV requirements cannot be met. So far two categories of ISG NFV requirements have been identified, that need more support in OpenStack, and are captured inn the tables below:

- A category focused on workload performance
- A category focused the management and orchestration of virtualized resources

In addition, ISG NFV will continue this analysis and may communicate in the future additional perceived gaps (e.g. related to reliability/availability, etc).

Gaps identified focusing on workload performance are captured in the table below:

NFV Entity	ETSI NFV spec Identifier	Semantics	Reference	Gap Analysis
MANO Virtual Machine Descriptor	*Vi-Vnfm and VDU in VNF Descriptor*	*SRIOV allows the NIC to be shared across multiple VNFs. This features allows the VIM to expose that functionality while allow retaining PCI Pass Through Capability*	*Identified in Annex B3 Gap Analysis Table of ETSI GS NFV-PER 001 V0.0.9*	*OpenStack Blueprint in Flight* ***PCI Passthrough_SR-IOV Support*** *https://blueprints.launchpad.net/nova/+spec/pci-passthrough-sriov*
MANO VIM	*Vi-Vnfm and VNF Descriptor*	*This feature aims to give users and administrators the ability to control the vCPU topology exposed to guests. This enables them to avoid any of the limitations on vCPU toplogies that OS vendors place on their projects.*	*Identified of ETSI GS NFV-PER 001 V0.0.9 section 8.1.1.3*	*OpenStack Blueprint in Flight* ***vCPU Topology*** *https://blueprints.launchpad.net/nova/+spec/virt-driver-vcpu-topology*
MANO VIM & MANO Virtual Machine Descriptor	*Vi-Vnfm and Service VNF And Infrastructure Descriptor – Se –Ma*	*This feature aims to enhance the libvirt driver to be able todo intelligent NUMA node placement for guests. This will increase the effective utilization of compute resources and decrease latency by avoiding cross-node memory*	*Identified of ETSI GS NFV-PER 001 V0.0.9 in section 8.2.2.3 Deterministic memory allocation*	*OpenStack Blueprint* ***NUMA Placement*** *https://blueprints.launchpad.net/nova/+spec/virt-driver-numa-placement*

		accesses by guests.		
VIM & MANO Virtual Machine Descriptor	*Vi-Vnfm and Service VNF And Infrastructure Descriptor – Se –Ma*	*This feature aims to improve the libvirt driver so that it can use large pages for backing the guest RAM allocation. This will improve the performance of guest workloads by improves TLB cache efficiency. It will effectively create the concept of "dedicated RAM" guest instances*	*Identified in Annex A Gap Analysis Table of ETSI GS NFV-PER 001 V0.0.9 under Large Pages Available*	*OpenStack Blueprint* ***VM Large Page Allocation*** *https://blueprints.launchpad.net/nova/+spec/virt-driver-large-pages*
VIM & MANO Virtual Machine Descriptor	*Vi Vnfm and Service VNF And Infrastructure Descriptor – Se –Ma*	*This feature aims to improve the libvirt driver so that it is able to strictly pin guest vCPUS to host pCPUs. This provides the concept of "dedicated CPU" guest instances and improves data plane determinism*	*Reference in section 8.2.2.3 Deterministic allocation of threads in CPUS*	*OpenStack Blueprint* ***CPU Pinning*** https://blueprints.launchpad.net/nova/+spec/virt-driver-cpu-pinning
VIM & MANO Virtual Machine Descriptor	*Vi Vnfm and Service VNF And Infrastructure Descriptor – Se –Ma*	*The NUMA locality of I/O devices is another important characteristic to consider when configuring a high performance, low latency system for NFV workloads. Otherwise excessive latency is introduced into*	*Identified of ETSI GS NFV-PER 001 V0.0.9 section 8.1.1.4*	*OpenStack Blueprint* ***NUMA and I/O Locality*** https://blueprints.launchpad.net/nova/+spec/input-output-based-numa-scheduling

		the system		
VIM & MANO Virtual Machine Descriptor	*Vi Vnfm and Service VNF And Infrastructure Descriptor – Se –Ma*	*Supporting NUMA and VCPU topology configuration for virtual guests.NUMA and VCPU topology are so closely related that an effective design must consider both of these aspects at the same time.*	*Reference in section 8.2.2.3 Deterministic allocation of threads in CPUS*	*OpenStack Blueprint* ***NUMA and vCPU Support*** *https://blueprints.launchpad.net/nova/+spec/virt-driver-guest-cpu-memory-placement*

Gaps identified focusing on workload performance are captured in the table below:

NFV Entity	Identifier	Semantics	Reference	Gap Analysis
Vi-Vnfm, Nfvo-Vi	Virtualised resource management	The interface defines operations comprising the management of resource reservations, including the creation, querying, update and release of such reservations.	GS NFV-MAN 001 Clause 7.3.3	OpenStack Havana release does not support a resource reservation service and consequently it does not provide any interface for resource reservation. NOTE 1: Blazar [i.3] (formerly named Climate), a resource reservation service for OpenStack is under development.
		The interface defines operations comprising the allocation, query, update, scale, migration, operation and release of virtualised resources, including those resources related to networking	GS NFV-MAN 001 Clause 7.3.3	OpenStack Havana release API for networking resources includes specification and information models for defining connectivity among virtual machines in terms of networks (L2 domain) and subnets (L3 domain). A detailed description of network resources including QoS requirements is not provided. NOTE 2: OpenStack's Neutron (Networking API v2.0) is extensible, and it enables introducing new features in the API, introducing vendor-specific niche functionality, and supporting experimental functionalities.
	Virtualised resource catalogue management	The interface defines operations to retrieve the	GS NFV-MAN 001 Clause 7.3.2	OpenStack Havana release does not offer a common virtualized resource catalogue advertisement service. A consumer needs to individually retrieve

		resource catalogue information available from the producer functional block, including notification of catalogue changes		the list with types of resources from each individual OpenStack service. NOTE 3: OpenStack's Nova service advertises the list of VM sizes/flavours that it offers, Glance service advertises the list of images/OS available, Keystone service advertises the list of available OpenStack services endpoints/APIs, etc. OpenStack Havana release does not provide notifications when the individual resource catalogues change or get updated (e.g., when new virtualized resources are made available for consumption to consumer functional blocks).

2. Actions:

OpenStack is invited to take note of the ETSI NFV referred documents and the initial requirement gaps in section 1. ETSI NFV ISG would welcome the emergence of a dialogue on those requirements, and would appreciate any feedback on OpenStack/NFV community to address these issues, or – even better- feedback regarding how such issues are already in progress to be addressed. In addition, ETSI NFV is interested in feedback on any topics that OpenStack community has identified as relevant in the NFV space, that may be jointly explored between OpenStack/NFV community and ETSI NFV, in order to achieve ETSI NFV's goals.

3. Date of next meetings of the originator:

NFV#8. 18th – 21st Nov 2014, Scottsdale AZ, USA.

References (included in the present document):

- NFV Use Cases (published)
  http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV001v010101p%20-%20Use%20Cases.pdf
- NFV Requirements (published)
  http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV004v010101p%20-%20Virtualisation%20Requirements.pdf
- NFV Architectural Framework (published)
  http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV002v010101p%20-%20Architectural%20Fwk.pdf
- NFV Terminology (published)
  http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV003v010101p%20-%20Terminology.pdf
- NFV Proof of Concept Framework NFV Performance and Portability Best Practices (published)
  http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV-PER001v010101p-%20Perf_and_Portab_Best_Practices.pdf
- NFV Management and Orchestration (draft)
  http://docbox.etsi.org/ISG/NFV/Open/Latest_Drafts/NFV-MAN001v061-%20management%20and%20orchestration.pdf