ETSI NFV Announcement on Document Availability

During the first six months of its second phase, the ETSI NFV ISG has been actively working on the development of normative specifications for the reference points identified by the NFV Architecture Framework, addressing the interoperability goals that constitute its key objective, and on continuing the exploration of NFV technical aspects in the essential areas identified during the inception of this second phase. This has been performed in a framework of continuous and tight collaboration with those external bodies (SDOs and open-source projects) most directly concerned with NFV technologies.

One important milestone in this development, aimed to facilitate open collaboration with external bodies and the industry and academia at large, has been the decision to make all NFV draft documents publicly available. The results of the NFV ISG activities can be analyzed and commented on by all those interested in them, and incorporated into third party activities as soon as they become available (with all due caveats regarding their draft nature). The NFV ISG also expects to increase the quality of its results by having a tighter feedback loop with the wider community and more direct experimentation capabilities.

The ETSI NFV ISG second phase was launched with a very ambitious work program, in which the vast majority of activities were committed to conclude along the first year of the two-year phase lifetime. In what follows we provide a brief report of the work already completed (or very close to completion) and the status of those documents that had committed to achieve significant results during the first six months of the second phase term. This is intended to be the first of a succession of announcements to be made public with a six-month periodicity.

Documents Completed or Close to Completion

IFA001 - Acceleration Technologies; Part 1: Report on Acceleration Technologies and Use Cases

The report on NFV acceleration outlines a common acceleration abstraction layer, which allows deployment of various accelerators within the NFVI and facilitates interoperability between VNFs and accelerators. This document also categorizes the accelerators on the aspects of accelerator type, location, etc. with twelve supporting use cases to illustrate the usage of acceleration techniques in NFV environment.

Work on this report is nearing completion, and it is not foreseen that further significant additions will be added in this release. Maybe some editorial changes are needed for the final release, but the technical changes are likely to be minor.

REL002 - Reliability; Report on Scalable Architectures for Reliability Management

The goal of this document was to provide a technical report providing a feasibility study of the use of the scalable architecture techniques currently adopted in cloud datacenters to enable reliability management in an NFV environment. In particular, two techniques have been described.

The first one, Migration Avoidance, enables dynamic scaling of VNFs when existing VNFs are unable to cope with unexpected bursts of incoming telecommunications traffic. The second, termed Lightweight Rollback Recovery, enables the recovery of failed VNFs without degrading existing traffic flows, by instantiating a backup VNF to assume the role of the failed VNF immediately after the failure.

Results from lab tests on these techniques are provided in detail – they demonstrate the efficacy of these techniques for reliability management. The development of specific architectures to support these techniques in actual telecommunications network conditions is a topic for further study.

SEC002 - Security; Cataloguing Security Features in Management Software

This document is a survey of the security features in the open-source management software relevant to NFV, in particular OpenStack as the first case study. The document addresses the OpenStack modules that provide security services (such as authentication, authorization, confidentiality protection, integrity protection, and logging) together with the full graphs of their respective dependencies down to the ones that implement cryptographic protocols and algorithms. It also identifies a set of recommendations on the use of and enhancements to OpenStack as pertinent to NFV.

SEC004 - Privacy and Regulation; Report on Lawful Interception Implications

The present report provides a problem statement on implementing LI in NFV and identifies the necessary capabilities to be provided in NFV to meet the requirements outlined for telecommunications capabilities in general in ETSI TS 101 331. The document identifies the challenges of providing LI in an NFV, and it is intended to give guidance to the NFV community and to the wider LI community on the provision of LI in an NFV.

It considers the generic global requirements for Lawful Interception, including their legal basis and general CSP obligations, together with a series of recommendations for the NFV implementation of LI requirements, and the challenges of applying legacy LI models to NFV against a consideration of specific NFV problem sets.

Status of Selected Documents

EVE005 - Ecosystem; Report on SDN Usage in NFV Architectural Framework

This document provides a technical report on the use of SDN in an NFV architectural framework, including guidance with a number of design patterns and recommendations for potential requirements and further work in the ETSI NFV ISG.

EVE005 leverages existing work from ETSI ISG NFV, especially SWA, MANO and INF documents from Phase 1. It identifies use cases and defines typical design patterns on the usage of SDN within an NFV architecture framework, including position of SDN resources, SDN controller and SDN applications and the different combinations and associated reference points. This includes SDN Controller as a VNF, SDN Controller as a realization of the Infrastructure network controller, or SDN controller in the tenant domain. Network domains to be covered include datacenter SDN, datacenter-WAN interworking, access network and WAN. ETSI NFV PoC teams have been invited to study this topic and their feedback is included as an annex. A comparison of open-source network controllers has also been performed to identify the scope of an SDN controller.

This technical report supports discussions with other SDO and open-source projects such as IETF, OPNFV, ONF, OpenStack, OpenDaylight and others as appropriate. The report will also make recommendations as to whether normative work should be initiated as a follow-up activity.

At present the content is nearly completed. The team is collecting some feedback from external entities and from the open area, before finalizing the set of recommendations and the document.

IFA002 - Acceleration Technologies; Part 2: VNF Interfaces Specification

IFA002 focuses on making hardware and software accelerators available to VNFCs in an implementation independent way and giving means to control acceleration within a VNF. Acceleration encompasses network traffic optimizations between VNFCs of a single VNF, network features offloads (e.g. IPSec), compute offloads (e.g. compression, cryptographic operations) or storage access acceleration.

An acceleration model has been defined around the concept of Extensible Para-virtualized Device (EPD) which is derived from an extension of the Virtio device model specified by the OASIS group. An EPD and its associated device driver are located in Virtualization Containers and communicate with a hypervisor domain backend that helps adapting to hardware or software implementation of the accelerated function. The EPD may receive software plugin and resources from the hypervisor domain to allow the most direct access to acceleration while preserving portability.

At present, the defined requirements for the Common Acceleration Virtualization interface and EPD drivers have been defined. A number of abstract interfaces have been identified, together with requirements for some of them.

IFA003 - Acceleration Technologies; Part 3: vSwitch Benchmarking and Acceleration Specification

IFA003 specifies performance-benchmarking metrics for virtual switching, with the goal that the metrics will adequately quantify performance gains achieved through virtual switch acceleration conforming to the associated requirements specified herein. The acceleration-related requirements will be applicable to common virtual switching functions across usage models such as packet delivery into VNFs, network overlay and tunnel termination, stateful Network Address Translators (NAT), service chaining, load balancing and, in general, match-action based policies/flows applied to traffic going to/from the VMs. The document will also provide deployment scenarios with applicability to multiple vendor implementations, and recommendations for follow-on proof of concept activities.

The document currently contains the definition of the vSwitch requirements (already formulated in the early drafts), the set of metrics, and a format defined to develop and document deployment scenarios. The plans is to complete the set of requirements (and any additional metrics) as well as define the deployment scenarios that combine these requirements together.

IFA005 - Management and Orchestration; Or-Vi Reference Point – Interface and Information Model Specification

This document is dedicated to the specification of the interfaces and related information elements exposed by the VIM towards the NFVO, and exposed by the NFVO towards the VIM over the Or-Vi reference point. It addresses requirements describing, at a high level, which interfaces are to be supported and what capabilities are to be supported over those interfaces. The detailed interface definitions describe the request and response messages used to invoke operations on the VIM, and list the information elements exchanged in those messages. Where necessary (typically if the information element has multiple attributes) the information elements are also defined.

The interfaces currently considered in the document are those exposed by the VIM, and related to the management of VNF software images, virtual resources, forwarding paths, resource performance, and resource faults, as well as the notifications of resource changes. Although interfaces exposed by the NFVO towards the VIM are in the scope of the document, no such interfaces have yet been identified.

Work on IFA005 is nearing completion, and it is not foreseen that further significant additions (new interfaces, or new operations on existing interfaces) will be added in this release. However, it is possible that activity in other ongoing documents (in particular, IFA010) could lead to some additions being required.

IFA006 - Management and Orchestration; Vi-Vnfm Reference Point – Interface and Information Model Specification

This document seeks the specification of the interfaces and related information elements exposed by the VIM towards the VNFM, and exposed by the VNFM towards the VIM over the Vi-Vnfm reference point. It considers requirements describing, at a high level, which interfaces are to be supported and what capabilities are to be supported over those interfaces. The detailed interface definitions describe the request and response messages used to invoke operations on the VIM, and list the information elements exchanged in those messages. Where necessary (typically if the information element has multiple attributes) the information elements are also defined.

The interfaces currently defined by this document are those exposed by the VIM, and related to the management of VNF software images, virtual resources, resource performance, and resource faults, as well as the notifications of resource changes. Although interfaces exposed by the VNFM towards the VIM are in the scope of the document, no such interfaces have yet been identified.

Work on IFA006 GS is nearing completion, and it is not foreseen that further significant additions (new interfaces or new operations on existing interfaces) will be added in this release. However, it is possible that activity in other ongoing documents (in particular, IFA010 and IFA007) could lead to some additions being required. In addition, the section on “Security Considerations” still needs to be addressed, and may also depend on the feedback from the SEC WG.

IFA009 - Management and Orchestration; Report on Architectural Options

The NFV architectural framework will appear in multiple different deployments. A deployment may use or combine functional blocks in a specific way to optimize for requirements of this deployment. The IFA009 report on architectural options outlines some of the major options that may be present. Since these options have had a direct impact on the normative specifications we produce, we thought it is crucial to document those options.

The objective of the IFA009 work is to give the reader of the ETSI NFV IFA specifications some background information why certain conditional requirements for interfaces or the architecture exist. One of the present examples is the way in which resource management for VNFs is invoked. Another example is how a VNF Manager can be used standalone or be combined with a VNF or Element Manager.

While the work on IFA is progressing IFA009 may be extended to further clarify existing options or outline new options that become relevant to document.

IFA010 - Management and Orchestration; Functional Requirements Specification

In order to guide the development of the specification of the interfaces exposed between the NFV-MANO functional blocks, IFA010 specifies functional requirements for three main Management and Orchestration blocks: NFVO, VNFM and VIM, and also specifies general guidelines and requirements for NFV management and orchestration interface design. The functional requirements are documented in IFA010 by identifying what functional capabilities of management and orchestration should be supported by those function blocks.

IFA010 has developed a good amount of functional requirements for VIM, NFVO and VNFM related to virtualized resource management, aligned with those interface requirements specified by other documents (IFA05, IFA06...) The current results also include some basic functional requirements for NFVO, VNFM and VIM on other management aspects. IFA010 has also developed basic guideline and requirements for NFV management and orchestration interface design, and accommodated some different deployment options.

IFA010 is still a work in progress, and continues to develop new and missing functional requirements for NFVO, VNFM and VIM, solve inconsistencies among requirements, address valuable feedback from external bodies, and fix technical errors in the existing functional requirements.

IFA011 - Management and Orchestration; VNF Packaging Specification

This document specifies the packaging for VNFs to be delivered to service providers. Standard packaging enables delivery of multiple vendor VNFs to service providers ensuring consistency in orchestration and management. A VNF package includes the required files and meta-data descriptors (e.g. VNFD) required to verify and successfully instantiate a VNF. The specification for VNF packages includes the definition of the structure of the package; how versioning, licensing and certificates are included; considerations on security and integrity, and the integration of VNF deployment templates.

Work on IFA011 has been slow mainly due to conflicting priorities with other documents. As work in these other documents progress, IFA011 is expected to pick up momentum. Today the effort is largely around the VNF deployment template and there are promising contributions in this area that will help bring work back on track.

REL004 - Assurance; Report on Active Monitoring and Failure Detection

This technical report proposes a framework for active monitoring and fault isolation for NFV environments. The document discusses the pertinent uses cases of fault isolation, periodic performance monitoring and capacity planning and the need to augment traditional active monitoring techniques with passive monitoring and NFVI analytics. The key focus is on describing methods that take into account the increased variability that multi-tenancy and multi-vendor scenarios introduce in NFV.

The report proposes two management entities as part of the active monitoring system – Test Controller and Test Result Analysis Module. It also defines the recommendations for implementation of these two entities and proposes the use of virtual test agents for increased network visibility and maintaining the monitoring point of presence in NFV based networks. Additionally, the document also takes concepts from E2E service monitoring in cloud environment and discusses its applicability to NFV. At the same time it presents the need to characterize SP centric services in terms of NFVI resource usage and to understand the performance impact of multiple services in a multi-

tenant environment. Key aspects for QoE measurement are discussed and example methodologies for QoE measurement of E2E services are described as well.

SEC009 – Report on Use Cases and Technical Approaches for Multi-Layer Host Administration

SEC009 addresses one of the enduring issues within complex administration domains: the provision of multi-layer administration within a single host.

Several different use cases have been identified, and currently exist in various stages of detail. These include operator-service related use cases (e.g. multi-tenant hosting and IaaS as a service), security sensitive or security network monitoring functions, and compliance-related use cases such as Retained Data, Lawful Interception and customer data privacy. The document describes recommendations, maps the use cases to the recommendations, and then describes some measures that could be used to meet these recommendations. The final section provides a description of three approaches to meeting the recommendations, addressing advantages and disadvantages.

The document is nearing maturity and is expected to meet its planned publication date. Future work contemplates the improvement of use cases, the analysis of the different approaches with further descriptions and diagrams.

TST001 - Pre-deployment Testing; Report on Validation of NFV Environments and Services

This document is committed to provide recommendations for lab validation of VNFs, their interaction with the NFV functional blocks and the NFV blocks themselves, including guidelines for user and control plane performance validation along with reliability and availability features.

The TST001 report identifies the peculiarities of testing virtual network functions in an NFV environment with respect to their physical counterparts, discussing the impact of virtualization on testing methods and assuring that the System Under Test (SUT) and Test Environments are properly identified for the cases where either the NFV Infrastructure, the VNFs, the Network Service (NS) or the own NFV MANO stack are under test. Additionally, the report provides step-by-step methodologies for common VNF and NS tests (e.g. VNF instantiation testing, data plane benchmarking, speed of activation of a NS, auto-scaling validation, etc.), which augment others commonly used in traditional physical environments.

Work in TST001 is in an advanced stage where most tests cases can be considered close to their final definition. However, some activity is expected to split some test cases (creating more atomic ones when needed) and align the test descriptions to the formal methodologies under elaboration in other activities in the TST WG, prominently TST002.

TST002 - Testing Methodology; Report on Interoperability Testing Methodology

The goal of TST002 is to study how interoperability test methodology can be applied to NFV by analyzing the functional blocks and interfaces defined within the NFV architecture and the NFV capabilities enabled by the current release.

The TST002 report provides methodology guidelines for interoperability testing for NFV, including a review of basic concepts for interoperability testing and their fit in an NFV environment, or a methodology for the development of interoperability test specifications that is illustrated with examples related to realistic NFV operations. While these sections are in a pretty mature state and can be considered stable, there is also active work in the areas of the definition of a generic SUT architecture and the collection of NFV interoperability features, which need to be aligned to the contents of the next interim specs from normative documents.