

Internet Research Task Force (IRTF)
Internet Draft
Intended status: Informational
Expires: September 2015

Z. Qiang
Robert Szabo
Ericsson
March 2, 2015

Elasticity VNF
draft-zu-nfvrg-elasticity-vnf-01.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on September 3, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

This draft is an analysis of Network Function Virtualization (NFV) applications based on the NFV architecture, use cases and requirements. The purpose of this analysis is to identify any NFV characteristics related issues. The analysis is focusing on elastic VNF with predicable performance, reliability and security. Only the issues which are unique to NFV are discussed in this document.

Table of Contents

1. Introduction.....	3
2. Conventions used in this document.....	3
3. Terminology.....	4
4. Network Function Virtualization.....	5
4.1. NFV Requirements.....	5
4.2. NFV Use Cases.....	6
4.2.1. Network Function Virtualization Infrastructure.....	6
4.2.2. Telecom Network Functions Migration.....	7
5. Elasticity in a Distributed Cloud.....	7
5.1. NFV Infrastructure.....	8
5.2. Elastic VNF.....	8
5.3. VNF Forwarding Graphs.....	9
5.4. VNF scaling across multiple NFVI PoPs.....	10
6. Elasticity with Predicable Performance.....	10
6.1. Predicable Performance.....	10
6.2. Hardware virtualization features.....	11
6.3. Network Overlay.....	12
7. Elasticity with Reliability.....	12

8. Elasticity with Security.....	13
9. Security Considerations.....	13
10. IANA Considerations.....	13
11. References.....	13
11.1. Normative References.....	13
11.2. Informative References.....	13
12. Acknowledgments.....	14

1. Introduction

Network Functions Virtualization (NFV) is a network architecture concept that proposes using IT virtualization related technologies, to virtualize entire classes of network node functions into building blocks that may be connected, or chained, together to create communication services. NFV aims to transform the traditional operator architect networks by evolving standard IT virtualization technology to consolidate network equipment types onto industry standard high volume services, switches and storage, which could be located in a variety of NFV Infrastructure Point of Presences (NFVI PoPs) including Data Center (DC), network nodes and in end user premises. It is also indicated that an important part of controlling the NFV environment should be done through automation network management and orchestration.

This draft is an analysis of NFV applications based on the NFV architecture, use cases and requirements. The purpose of this analysis is to identify any NFV characteristics related issues. The analysis is focusing on elastic VNFs with predicable performance, reliability and security. Only the issues which are unique to NFV are discussed in this document. The intention is to identify what is missing, and what is needed to be addressed in terms of protocol / solution specifications which may be the potential work for IETF.

The reader is assumed to be familiar with the terminology as defined in the NFV document [nfv-tem].

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

3. Terminology

This document uses the same terminology as found in the NFV end to end architecture [nfv-tem]:

Network Function Consumer: a Network Function Consumer (NFC) is the consumer of virtual network functions. It can be either an individual user, home user or the enterprise user.

NFV: network function virtualization. NFV technology uses the commodity servers to replace the dedicated hardware boxes for the network functions, for example, home gateway, enterprise access router, carrier grade NAT and etc. So as to improve the reusability, allow more vendors into the market, and reduce time to market. NFV architecture includes a NFV Control and Management Plane (orchestrator) to manage the virtual network functions and the infrastructure resources.

NF: A functional building block within an operator's network infrastructure, which has well-defined external interfaces and a well-defined functional behavior. Note that the totality of all network functions constitutes the entire network and services infrastructure of an operator/service provider. In practical terms, a Network Function is today often a network node or physical appliance.

Network Function Provider: a Network Function Provider (NFP) provides virtual network function software.

Network Service Provider (NSP): a company or organization that provides a network service on a commercial basis to third parties. A network service is a composition of network functions and defined by its functional and behavior specification. The NSP operates the NFV Control Plane.

NFV Infrastructure (NFVI): NFV Infrastructure indicates the computing, storage and network resources to implement the virtual network function. High performance acceleration platform is also part of it.

VNF: virtual network function, an implementation of an executable software program that constitutes the whole or a part of an NF that can be deployed on a virtualization infrastructure.

VM: virtual machines, a program and configuration of part of a host computer server. Note that the Virtual Machine inherits the

properties of its host computer server e.g. location, network interfaces.

NFV Control and Management Plane (NFVCMMP): a NFV Control and Management Plane is operated by a NSP and orchestrates the NFV NFV Overview

4. Network Function Virtualization

4.1. NFV Requirements

There are many virtualization requirements described by NFV in [nfv-req]. The followings are highlights of a few NFV requirements which are related to this document:

- Portability: VNF portability is a reasonable generic virtualization requirement. It allows VNF mobility across different but standard multi-vendor environment. However, moving a VNF within the NFV framework with the Service Level Specification (SLA) requirements including performance, reliability and security could be a challenge.
- Performance: Virtualization adds additional processing overhead and increases the latency. For latency-sensitive VNFs, it is a big concern for NFV on how to achieve predictable low-latency performance.
- Elasticity: NFV elasticity requirement allows the VNF to be scaled within NFVI. Within the NFV framework, it is important to support VNF scaling with the SLA requirements including performance, reliability and security.
- Resiliency: NFV resiliency is a must requirement for NFV network, including both the control plane and data plane. Necessary mechanisms must be provided to improve the service availability and fault management.
- Security: The traditional telecom network functions are developed in dedicated hardware located in an isolated network. Security is provided by underlay network. When moving VNF into a DC network with shared Infrastructure, security becomes a big concern.
- Service Continuity: At VNF failure over, migration, mobility, and upgrading, service downtime may not be avoided. In NFV, service continuity must be supported which means the provided service must be restored at the VNF instance updated / replaced / recovered. This procedure includes the restoration of any ongoing data sessions. And it shall be transparent to the user of NFV service.

4.2. NFV Use Cases

Multiple use cases are described by NFV in [nfv-uc]. The followings are a highlight of the NFV use cases.

4.2.1. Network Function Virtualization Infrastructure

Network Function Virtualization Infrastructure as a Service (NFVIaaS), Virtual Network Function as a Service (VNFaaS) and Virtual Network Platform as a Service (VNPaaS) are the NFV use cases which describe how the telecom operators would like to build up their telecom cloud infrastructure using virtualization.

Network Function Virtualization Infrastructure (NFVI) is the totality of all hardware and software components which build up the environment in which VNFs are deployed. The NFVI can span across several locations. The network providing connectivity between these locations is regarded to be part of the NFVI.

NFVIaaS is a generic IaaS plus NaaS requirement which allows the telecom operator to build up a VNF cloud on top of their own DCs Infrastructure and any external DCs Infrastructure. This will allow a telecom operator to migrate some of its network functions into a 3rd party DC when it is needed. Furthermore, a larger telecom operator may have multiple DCs in different geography locations. The operator may want to setup multiple virtual data center (vDC), where each vDC may cross several of its physical DCs geography locations. Each vDC is defined for providing one specific function, e.g. Telco Cloud.

VNFaaS is more focusing on enterprise network which may have its own cloud infrastructure with some specific services / applications running. VNFaaS allows the enterprise to merge and/or extend its specific services / applications into a 3rd party commercial DC provided by a telecom operator. With this VNFaaS, the enterprise does not need to manage and control the NFVI or the VNF. However, NFV Performance & portability considerations will apply to deployments that strive to meet high performance and low latency considerations.

With VNPaaS, the mobile network traffic, including WiFi traffic, is routed based on the APN to a specific packet data service server over the mobile packet core network. Applications running at the packet data service server may be provided by the enterprise. And it is possible to have an interface to route the traffic into an enterprise network. But the infrastructure hosting the application is fully under controlled by the operator. However, the enterprise

has full admin control of the application and needs to apply all configurations on its own, potentially via a vDC like management interface with support of the hosting operator.

All the above use cases need solutions for the operator to share the infrastructure resources with 3rd parties. Therefore cross domain orchestration with access control is needed. Besides, the infrastructure resource management needs to provide a mechanism to isolate the traffic, not only based on the traffic type, but also from different operators and enterprises.

4.2.2. Telecom Network Functions Migration

Virtualization of telecom network functions, including Mobile Core Network functions, IMS functions, Mobile base station functions, Content Delivery Networks (CDN) functions, Home Environment functions, and Fixed Access Network functions, are described in the NFV use case document [nfv-uc]. In additional, VNF forwarding Graphs is another use case which describes how the user data packets are forwarded by traversing more than one operator service chain functions, such as DPI, Firewall, Content Filtering, before reaching the service server.

Migrate the telecom functions includes moving the control plane, data plane and service network into a cloud based network and using cloud based protocol to control the data plane. Service continuity, network security, service availability, resiliency in both control plane and data plane must be ensured at this migration.

5. Elasticity in a Distributed Cloud

Today the usage of personal devices, e.g. smartphones, for internet service traffic, telecom specific service access, and accessing the corporate network, is increased significantly. At the same time, telecom operators are under pressure to accommodate the increased service traffic in a fine-grained manner. Services provided by telecom network must be done in an environment of increased security, compliance, and auditing requirements, along with traffic load may be changed dramatically overtime. Providing self-service provisioning in telecom cloud requires elastic scaling of the VNF based on the dynamic service traffic load and resource management e.g. computing, storage, and networking.

The existing telecom network functions may not be cloud technologies ready yet. Most of the NFV functions are stateful and running on either specific hardware or a big VM. It is not designed to tolerate

any system failure in many VMs. The network functions are very difficult in term of configuration, scale updating, etc.

Re-engineering may be needed for virtualization enabling, e.g. software adaption for software and hardware decoupling. For cloud technologies readiness, telecom network functions need to be re-designed to run on small VMs with multiple instances which can provide higher application availability. Such VMs may be stateless in operations or may need to support state migration (e.g., OpenNF <http://opennf.cs.wisc.edu/>). With cloud ready network functions, applications' dynamic scaling can be achieved by adding more VMs into the service.

Virtualization provides the elasticity ability to scale up / down, scale out / in with guaranteed computational resources, security isolation and API access for provisioning it all, without any of the overhead of managing physical servers. However, there are still many optimizations which can be used to avoid the increasingly overhead.

5.1. NFV Infrastructure

Virtualized Network Function (VNF) is an implementation of a network function that can be deployed on Network Function Virtualization Infrastructure (NFVI).

For a large telecom operator, multiple NFVI Point of Presences (NFVI PoPs) may be created according to multiple physical data centers. As NFVI PoPs may be located in different geography locations, networking characteristics should be taken into account when selecting an NFVI PoP to host a VNF.

5.2. Elastic VNF

In many cases, a VNF may not be designed for scaling up/down. As scaling up/down may require a restart of the VNF which the state data may be lost. In that case either stateless operation is needed, or the support of state information migration procedure is required, which will increase the complexities of the VNF implementation.

Normally a VNF may be capable for scaling in/out only. Such VNF is designed running on top of a small VM and grouped as a pool of one VNF function.

VNF capacity may be limited if it only can be scaled within one NFVI PoP, e.g., within one DC in a geography location. As an NFVI which may be crossing multiple NFVI PoPs (or data center)s, it is possible

to scale an elastic VNF crossing different network zones if it is needed. At cross DC scaling, the result is that the new VNF instance may be placed at a remote cloud location. It is a must requirement to provide the same level of SLA including performance, reliability and security.

5.3. VNF Forwarding Graphs

In NFV network, a VNF Forwarding Graph (VNF FG) (an application) may consist of multiple VNFs, where each VNF may consist of multiple VNF instances. Normally the VNFs are working as such that the services provided by the VNFs may need to process the user data packets with several selected VNF instances before delivering it to its destination

For instance, when mobile users setup a PDN connection for IMS services, there are multiple network entities involved along the PDN connection, including eNB, Serving GW, PDN GW, P-CSCF, S-CSCF, etc. Another example is service function chaining, where a service chain is referring to one or more service processing functions in a specific order which are chained to provide a composite service.

In telecom cloud, a service session may traverse multiple stateful and stateless VNF functions of a VNF set. And with an NFVI consisting of multiple NFVI PoPs, it may be crossing multiple DCs. In such cloud, an incoming data packet may be processed by multi-VNF instance before delivering to the final destination. Therefore the east-west traffic (i.e. data traffic between VNFs within the DC) is much heavier comparing to the north-south traffic (i.e. data traffic in/out from the DC).

When placing VNF Forwarding Graphs, it is better spread the VNF components across many NFVI PoPs, which may give a better availability. However, multiple NFVI PoPs may also increases the network latency, which can be considerably big compared to latencies within a single NFVI PoP. Therefore, the whole VNF Forwarding Graph should be taken into account instead of a single VNF component during orchestration. Furthermore, during VNF scaling, dependencies (interconnection) with other service instances of the VNF Forwarding Graph shall also be considered.

When scaling, VNFs are not scaled only in relation to compute and storage PoPs. VNF instances may need to be grouped together according to the VNF FG and subjected to auto-scaling techniques to the entire group. The scaling policies, e.g., ratio between the

different VNFs, need to be applied on the VNF FG in aggregate to control the scaling process.

5.4. VNF scaling across multiple NFVI PoPs

Since in general, a VNF is part of a VNF Forwarding Graph (or a service function chain), meaning the data traffic may traverse multiple stateful and stateless VNF functions in sequence. When some VNF instances of a given service function chain are placed / scaled out in a distant cloud execution, the service traffic may have to traverse multiple VNF instances which are located in multiple physical locations. In the worst case, the data traffic may ping-pong between multiple physical locations.

Therefore it is important to take the whole service function chain's performance into consideration when placing and scaling one of its VNF instance. Network and cloud resources need mutual considerations [unify1].

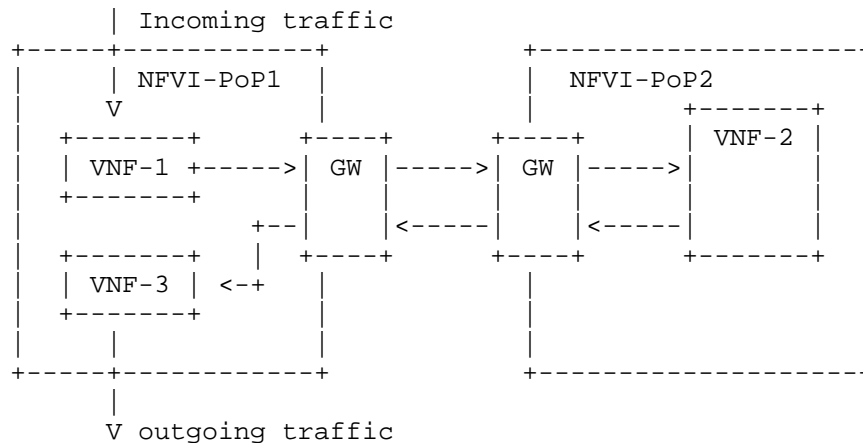


Figure 1 a data traffic flow traversing distributed VNF cloud

6. Elasticity with Predicable Performance

6.1. Predicable Performance

High performance with low-latency VNF is expected in the NFV framework. The NFVI metrics are related to any kind of metrics generated by the NFVI, including not only CPU load on a VM, CPU load

on a host, but also interrupt rate handled by the hypervisor or network latency/packet loss.

Virtualization adds additional overhead which impacts the performance. This additional extra distortion shall be avoided or, at least, minimized. It is a big concern for NFV on how to achieve predictable and low-latency performance, not only at placing a VNF into the DC, but also at VNF scaling.

Operator may wish to run standard test and use the result to provide KPIs of the VNF. A significant part of a VNF vendor's performance guarantees will depend on the choice of the virtualization technology.

Network latency may not be at the same level if the physical connections between the servers are various. Furthermore, geography location of the physical servers also increases the network latency. When placing or moving a VNF, the location of other VNFs of the same VNF set shall be considered to avoid network latency issue. For instance, VNF-a, VNF-b, and VNF-c are grouped as one VNF set. When moving VNF-b into a new location, the network connection between the new location and the existing location may be a concern. As the traffic may traverse from VNF-a to VNF-b, then VNF-c, moving the VNF-b into a new location may create Ping-Pong type of traffic, which the network latency may be doubled. The best choice would be to move the whole VNF set into the new location instead of only one VNF.

6.2. Hardware virtualization features

Virtualization layer adds minimal overhead and delivers a predictable performance between a minimum and maximum threshold for latency and jitter which are far more important. Light weight virtualization, e.g. container or bare metal, may be considered for performance sensitive VNF applications. In addition, hardware virtualization features (e.g. SR-IOV) are important to be supported in order to provide some performance improvement. Many VNF requires direct access to the device hardware so that they can offload functionality with throughput rates of millions of packets a second. Another alternative, which may be more attractive for latency-sensitive applications, is using non-hypervisor virtualization, including bare metal and Linux container.

Optimization to drive high-throughput network workloads associated with such functions as traffic filtering, NATing and firewalling. Avoiding performance bottleneck, the virtualization layer shall have a suitably-architected I/O stack.

6.3. Network Overlay

Network overlay adds additional overhead when forwarding the data packets. Reference [vxlan-p] is a VXLAN performance testing report which indicates the overlay performance is a concern. Avoiding overlay connections may be one option which is more attractive for latency-sensitive applications.

Furthermore, additional network latency may be added when traversing the cross-DC overlay connections. To avoid any additional network latency, all the functions of a VNF set may be placed in the same low-latency network zone, e.g. same host or same DC. However, when the capacity limitation the network zone is reached, scaling-out one VNF into another network zone may be needed. In this case, as the service session has to traverse the same path, the Ping-Pong traffic between the network zones cannot be avoided. Depends on the network overlay technologies used for the cross network zone connection, the overhead network latency can be various. In another words, the network performance may become unpredictable.

7. Elasticity with Reliability

NFV resiliency is a must requirement for NFV network, including both the control plane and data plane. Necessary mechanisms must be provided to improve the service availability and fault management.

With virtualization, the use of VNFs can pose additional challenges on the reliability of the provided services. For a VNF instance, it typically would not have built-in reliability mechanisms on its host (i.e., a general purpose server). Instead, there are more factors of risk such as software failure at various levels including hypervisors and virtual machines, hardware failure, and instance migration that may make a VNF instance unreliable. Even for cloud ready NFV applications, a HA may still be needed as the storage, load balancer may be failure. Service restoration solution is still needed.

One alternative to improve the VNF resiliency is to take snapshot of the VM periodically. At VNF failure, the network can restore the VM at same or different host using the stored snapshot. However, there is a downtime of the provided service due to the snapshot recovering. And the downtime is much longer than the expected value which could be tolerated by NFV. NFV has a completely different level of reliability requirements, e.g. recovering time, comparing to enterprise cloud applications.

To improve the network function resiliency, some kind high availability (HA) solutions may be needed for NFV network, which has the potential to minimize the service downtime at failure. However, in most of the telecom use cases, there are application level restoration procedures available which makes the high availability solution less important.

The VNF reliability can be achieved by eliminating any single points of failure by creating a redundancy of resources, normally, including enough excess capacity in the design to compensate for the performance decline and even failure of individual resources; that is, a group of VNF instances providing the same function works as a network function cluster or pool, which provides protection (e.g. failover) for the applications and therefore an increased availability.

8. Elasticity with Security

TDB

9. Security Considerations

This is a discussion paper which provides inputs for NFV related discussions and in itself does not introduce any new security concerns.

10. IANA Considerations

No actions are required from IANA for this informational document.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2234] Crocker, D. and Overell, P.(Editors), "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, Internet Mail Consortium and Demon Internet Ltd., November 1997.

11.2. Informative References

[nfv-arch] Network Functions Virtualization Infrastructure Architecture Overview; GS NFV INF 001.

[nfv-rel] Network Function Virtualization (NFV) Resiliency Requirements; ETSI GS NFV-REL 001.

- [nfv-uc] Network Function Virtualization (NFV) Use Cases; ETSI GS NFV 001
- [nfv-req] Network Function Virtualization (NFV) Virtualization Requirements; ETSI GS NFV 004
- [nfv-sec] Network Function Virtualization (NFV) NFV Security Problem Statement; ETSI NFV-SEC 001
- [nfv-tem] Network Function Virtualization (NFV) Terminology for Main Concepts in NFV; ETSI GS NFV 003
- [vxlan-p] Problem Statement for VxLAN Performance Test, draft-liu-nvo3-ps-vxlan-performance, (working in progress)
- [unify1] Szabo, R., Csaszar, A., Pentikousis, K., Kind, M., and D. Daino, "Unifying Carrier and Cloud Networks: Problem Statement and Challenges", draft-unify-nfvrg-challenges-00 (work in progress), October 2014.

12. Acknowledgments

Many people have contributed to the development of this document and many more will probably do so before we are done with it. While we cannot thank all contributors, some have played an especially prominent role. The following have provided essential input: Suresh Krishnan.

Authors' Addresses

Zu Qiang
Ericsson
8400, boul. Decarie
Ville Mont-Royal, QC,
Canada

Email: Zu.Qiang@Ericsson.com

Robert Szabo
Ericsson Research, Hungary
Irinnyi Jozsef u. 4-20
Budapest 1117
Hungary

Email: robert.szabo@ericsson.com
URI: <http://www.ericsson.com/>

