Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 11, 2015

L. Xia
Q. Wu
Huawei
D. King
Lancaster University
H. Yokota
KDDI Lab
N. Khan
Verizon

November 11, 2014

Requirements and Use Cases for Virtual Network Functions
draft-xia-vnfpool-use-cases-02

Abstract

Network function appliances such as subscriber termination, firewalls, tunnel switching, intrusion detection, and routing are currently provided using dedicated network function hardware. As network function is migrated from dedicated hardware platforms into a virtualized environment, a set of use cases with application specific resilience requirements begin to emerge.

These use cases and requirements cover a broad range of capabilities and objectives, which will require detailed investigation and documentation in order to identify relevant architecture, protocol and procedure solutions to ensure reliance of user services using virtualized functions.

This document provides an analysis of the key reliability requirements for applications and functions that may be hosted within a virtualized environment. These NFV engineering requirements are based on a variety of uses cases and goals , which include reliability scalability, performance, operation and automation.

Note that this document is not intended to provide or recommend protocol solutions.

This Internet-Draft will expire on May 11, 2011.

Table of Contents

1. Introduction

Network virtualization technologies are finding increasing support among network and Data Center (DC) operators. This is due to demonstrable capital cost reduction and operational energy savings, simplification of service management, potential for increased network and service resiliency, network automation, and service and traffic elasticity.

Within traditional DC networks, varied middleware boxes including FW (Fire Wall), NAT (Network Address Translation), LB (Load Balancers), WoC (Wan Optimization Controller), etc., are being used to provide network functions, traffic control and optimization.
Each function is an essential part of the entire operator and DC network, and overall service chain (required traffic path for users) Combined these functions and capabilities.

Currently, a significant amount of network functions are being migrated into virtualized entities, in essence the middleware capability is implemented in software on commodity hardware using well defined industry standard servers. Thus allowing the creation, modification, deletion, scaling, and migration of single or groups of network functions, across few or many servers.

These virtual network functions (VNF) may be location independent, i.e., they may exist across distributed or centralized DC hardware. This architecture will pose new issues and great challenges to the automated provisioning across the DC network, while maintaining high availability, fault-tolerant, load balancing, and plethora of other requirements some of which are technology and policy based.

Today, architecture and protocol mechanisms exist for the management and operation of server hardware supporting applications, these hardware resources are known as server node pools, which may be accessed by other servers and clients. These server node pools have a well-established set of requirements related to management, availability, scalability and performance.

[I-D.zong-vnfpool-problem-statement] provides an overview of the problems related to the reliability of a VNF set, and also introduces briefly a VNF pooling architecture. This document provides an analysis of the key reliability requirements for applications and functions that may be hosted within a virtualized environment. These Network Functions Virtualization (NFV) engineering requirements are based on a variety of uses cases and goals , which include reliability scalability, performance, operation and automation.

This document is not intended to provide or recommend solutions. The

intention of this document is to present an agreed set of objectives and use cases providing network function using virtualized instances, identification of key requirements across use cases.

1.1. Network Function Virtualization (NFV) Effort

NFV, an initiative started within the European Telecommunications Standards Institute (ETSI), aims to transform the way that network operators architect networks by evolving standard IT virtualization technology to consolidate many network equipment types to industry standard high volume servers, switches and storage.

The objectives for NFV being specified within the ETSI organization include:

- Rapid service innovation through software-based deployment and operationalization of network functions and end-to-end services;
- Improved operational efficiencies resulting from common automation and operating procedures;
- Reduced power usage achieved by migrating workloads and powering down unused hardware;
- Standardized and open interfaces between network functions and their management entities so that such decoupled network elements can be provided by different players;
- Greater flexibility in assigning Virtual Network Functions (VNF) to hardware;
- Improved capital efficiencies compared with dedicated hardware implementations.

1.2. Virtual Network Functions (VNF) Resilience Requirements

Deployment of NFV-based services will require the transition of resilient capabilities from physical network nodes, which are typically highly available, entities running Virtual Network Functions (VNFs) on abstracted pool of hardware resources.

Thus, it is critical to ensure that end-to-end user services which may require a variety of virtualized functions to be reliable, and in the event failure would support seamless failover when required to negate or minimize impact on user services.

A number of requirements have been discussed and documented within the NFV Industry Steering Group (ISG) working groups, including [ETSI-HA-USECASE] and are highlighted in following sub-sections.

1.2.1. Service Continuity

VNFs provide the capability to execute and operate network functions on varying types of Virtual machines (VMs), and subsequently physical equipment. It should be possible to inherently provides resiliency at the function level, as well as physically.

Network Functions (NFs) are assigned session IDs, Sequence IDs and Authentication IDs. This information may be static, dynamic and temporal so will need to be replicated and maintained as needed for failure scenarios.

Hardware entity such as a storage server or networking node are assigned a unique MAC address, which is often pre-configured (hardware encoded) and static.

In the event of a hardware failure or capacity limits (memory and CPU) hosting VMs and therefore VNFs, it may be necessary to move VNFs to another VM, and/or hardware platform. Therefore, service continuity must be maintained with no or negligible impact to users using with services being provided by the NFs.

1.2.2. Topological Transparency

Redundant systems are typically configured as an active and standby nodes, running a specific NF in the same LAN segment. It is possible that they are assigned duplicate IP addresses, and sometimes the same MAC address as well. In the event of an active node failure the standby node can take over transparently. This should be architecture supported by any eventual solution.

In order to achieve topological transparency and seamless hand-over the dependent nodes should replicate and maintain the necessary information so that in the event of failure the standby node takes over the service without any disruption to the users.

1.2.3. Load Balancing or Scaling

When load-balancing or scaling of sessions, the working session may be moved to a new VNF instance, or indeed a new VM on another hardware platform. Again, service continuity must be maintained.

2. Terminology

The following terms have been defined by the ETSI Industry Steering Group (ISG) responsible for the specification of NFV, and are reused in this document:

Network Function (NF): A functional building block within a network infrastructure, which has well-defined external interfaces and a functional behavior. In practical terms, a Network Function is today often a network node or physical appliance.

NFV Orchestrator: The NFV Orchestrator is in charge of the network wide orchestration and management of NFV Infrastructure (NFVI) and resources. The NFV Orchestrator has control and visibility of all VNFs running inside the NFVI. The NFV Orchestrator provides GUI and external NFV-Interfaces to the outside world to interact with the orchestration software.

Service Continuity: The continuous delivery of service in conformance with service, functional and behavioral specification and SLA requirements, both in the control and data planes, for any initiated transaction or session till its full completion even in the events of intervening exceptions or anomalies, whether scheduled or unscheduled, malicious, intentional or unintentional. From an end-user perspective, service continuity implies continuation of ongoing communication sessions with multiple media traversing different network domains (access, aggregation, and core network) or different user equipment.

Hypervisor: Software running on a server that allows multiple VMs to run on the same physical server. The hypervisor manages and provide network connectivity to Virtual machines [RFC7365].

Network Functions Virtualization (NFV): Moving network function from dedicated hardware platforms onto industry standard high volume servers, switches and storage.

Set-top Box (STB): This device contains audio and video decoders and is intended to connects to a variety of home user devices media servers and televisions.

Virtual Machine (VM): Software abstraction of underlying hardware.

Virtual Application (VA): A Virtual Application is the more general term for a piece of software which can be loaded into a Virtual Machine. A VNSF is just one type of VA amongst many others, which may not relate to any VNF (e.g. SW-tools or NFV-Infra-internal applications).

Virtualized Network Function (VNF): a VNF provides the same functional behavior and interfaces as the equivalent network function, but is deployed as software instances building on top of a virtualization layer.

The VNF Problem statement [I-D.zong-vnfpool-problem-statement]

defines the terms reliability, VNF, VNF Pool, VNF Pool Manager, and VNF Set. This draft also uses these definitions. In addition to the terms described above, this document also uses the following additional terminology:

VNF Pool: a group of VNF instances providing the same network function.

VNF Pool Manager: an entity that manages a VNF pool, and interacts with the service control entity to provide the network function.

VNF Set: a group of VNF instances that can be used to build network services.

3. Virtual Network Function (VNF) Pool Architecture

Shifting towards virtual network function presents a number of challenges and requirements, this document focuses on those related to network function availability and reliability. In large DC environments, a virtual server may need to deal with traffic from millions of hosts. This represents a significant scaling challenge for Virtual network function deployment and operation.

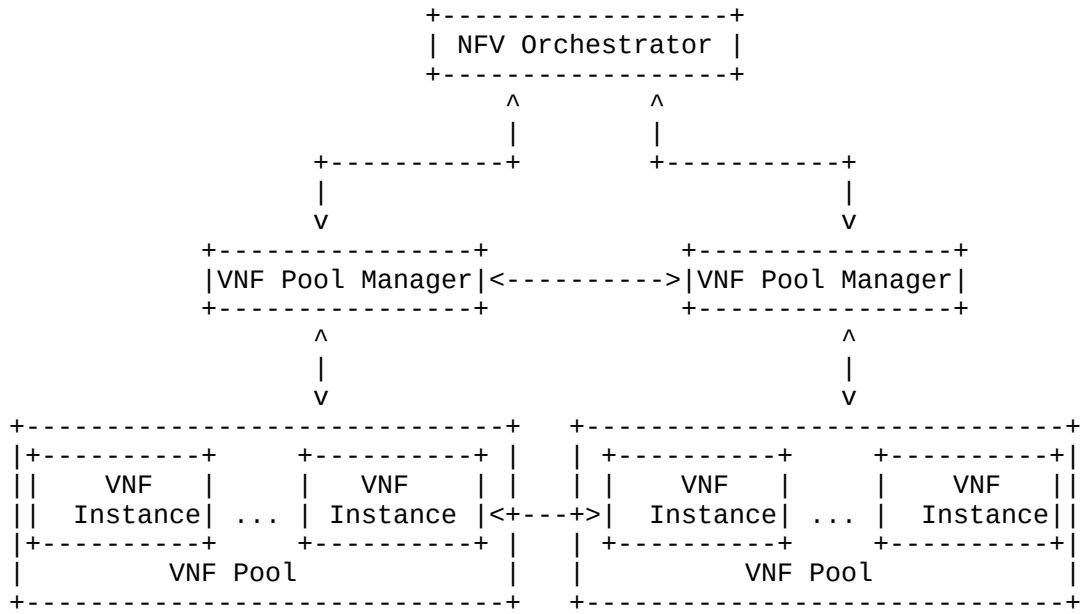

Figure 1: Typical VNF Pool Network Architecture

As shown in Figure 1, the overall architecture of VNF Pool-based network includes:

- o VNF Instances
- o VNF Pool
- o VNF Pool Manager

Rserpool [RFC5351] has the similar architecture to provide high-availability and load balancing, However Rserpool are only used to manage physical servers and can not deal with VNF instance when it was designed.

3.1. VNF Instance Resilience Objectives

In order to manage VNF-based nodes and provide fault tolerant and load sharing across nodes, the VNF instances may be initiated and established as logical element. A set of VNFs providing the same service type, is known as a VNF Pool, or groups of network functions (FW, LB, DPI) running on multiple VNFs, is known as a VNF Set.

Considering the reliability requirements of a VNF-based node architecture it should support several key points detailed below:

- o Application resource monitoring and health checking;
- o Automatic detection of application failure;
- o Failover to another VNF instance;
- o Transparency to other VNF instances;
- o Isolation and reporting of failures;
- o Replication of state for active/standby network functions.

3.2. Resilience of Network Connectivity

The other category of reliability requirements concerns the network connectivity between any two VNFs, across a VNF set, or between VNF Pool Manager.

The connectivity between the VNF Pool Manager and the VNF instance is used to provide registry service to the VNF Set. A set of VNF Pool managers might be configured to provide reliable registration.

When one VNF instance cannot obtain a register response from the assigned VNF Pool Manager, it should be capable of fail-over to

another VNF Pool Manager. Connectivity may also be monitored by the VNF Pool Manager to the VNF Instance periodically as well.

The connectivity between Pool Managers is used to maintain synchronization of data between VNFs located in different VNF Pools or VNF Sets. This allows every Pool Manager to acquire and maintain overall information of all VNFs and provide protection for each other.

For all types of network connectivity discussed previously, the key reliability requirements stay consistent and include:

- Automatic detection of link failure;
- Failover to another usable link;
- Automated routing recovery.

3.3. Service Continuity

It is critical to ensure end-to-end service continuity over both physical and virtual infrastructure. A number of requests exist to maintain user services in the event of network or VNF instance failure, these include:

- Storage and transfer of state information within the VNFs;
- VNF capacity (memory and CPU) limitations per instance to avoid overbooking, and failure of end-to-end services;
- Automated recovery of end-to-end services after failure situations;

4. General Resilience Requirements For VNF Use Cases

4.1. Resilience for Stateful Service

In the service continuity use case provided by the European Telecommunications Standards Institute (ETSI) Network Function Virtualization (NFV) Industry Specification Group (ISG) [NFV-REL-REQ] , which describes virtual middlebox appliances providing layer-3 to layer-7 services may require maintaining stateful information, e.g., stateful vFW. In case of hardware failure or processing overload of VNF, in addition to the replacement of VNF, it is necessary to move its key status information to new VNF for service continuity. See Figure 2 (Resilience for Stateful Service) for clarification.

In case of multiple vFws on one VM and not enough resources are available at the time of failure, two strategies can be taken: one is

to move as many vFws as possible to a new place according to the available resources, and the other is to suspend one or more running VNFs in the new place and move all vFws on the failed hardware to it.

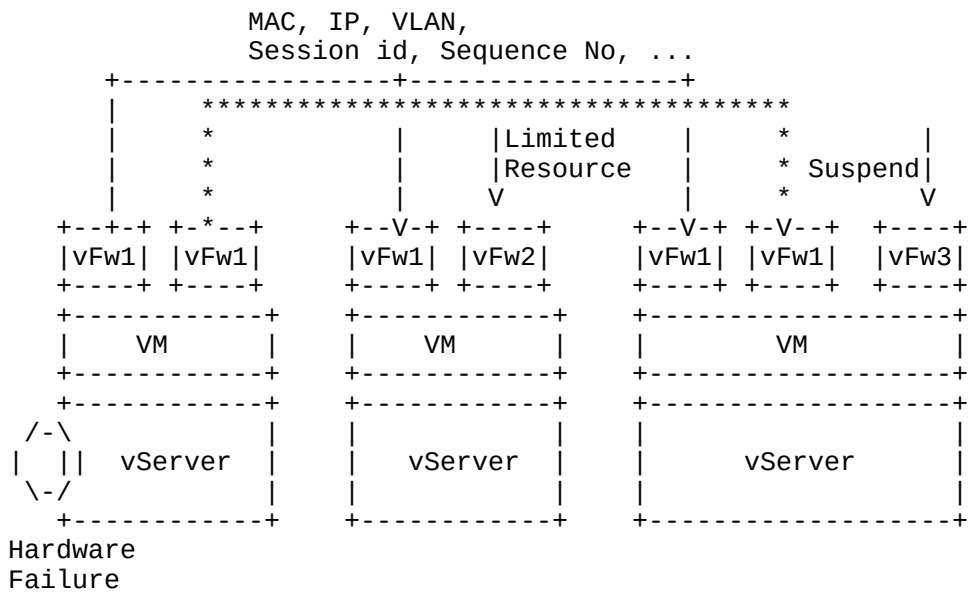

Figure 2: Resilience for Stateful Service

In both scenarios, the following requirements need to be satisfied:

- Supporting status information maintaining;
- Supporting status information moving;
- Supporting VNF moving from one VM to another VM;
- Supporting partial VNFs moving;
- Seamless switching user traffic to alternative VMs and VNFs.

4.1.1 State Synchronization

As identified in section 4.1 (Resilience for Stateful Service) their is a requirement for for state synchronization. A failure of a vFW would result in the loss of active connections transiting the node. Any connection-orientated or secure sessions, including enterprise and financial transactions, may be critical, and losing them would result in the loss of data.

If required it should be possible to ensure that the VNF Pool infrastructure should minimise or negate session data traffic if a vFW failures. Prior to the failure the vFW might advertise and synch the connection information transitioning its node. The connection state synchronization to other vFWs acting as stand-by nodes would provide fast fail-over and minimal connection interruption to users.

This synchronization mechanism should be supported by the (NFV) infrastructure level, that is, ideally each application does not need to code the redundancy procedures (reserve a VM resource, instantiate one or more backup server(s), copy the state, keep them in sync, etc). Also, such a state can be embedded in each vNF or stored in an external virtual storage, which should be supported by the NFV infrastructure.

4.2. Auto Scale of Virtual Network Function Instances

Adjusting resource to achieve dynamic scaling of VMs described in the ETSI [NFV-INF-UC] use case and [NFV-REL-REQ]. As shown in Figure 3, if more service requests come to a VNF than one physical node can accommodate, processing overload occurs. In this case, the movement of the VNF instance to another physical node with the same resource constraints will create a similar overload situation. A more desirable approach is to replicate VNF instance to one or more new VNF instances and at the same time distribute the incoming requests to those VNF instances.

In a scenario where a particular VNF requires increased resource allocation to improve overall application performance, the network function might be distributed across multiple VMs. To guarantee performance improvement, the hypervisor dynamically adjusts (scaling up or scaling down) resources to each VNF in line with the current or predicted performance needs.

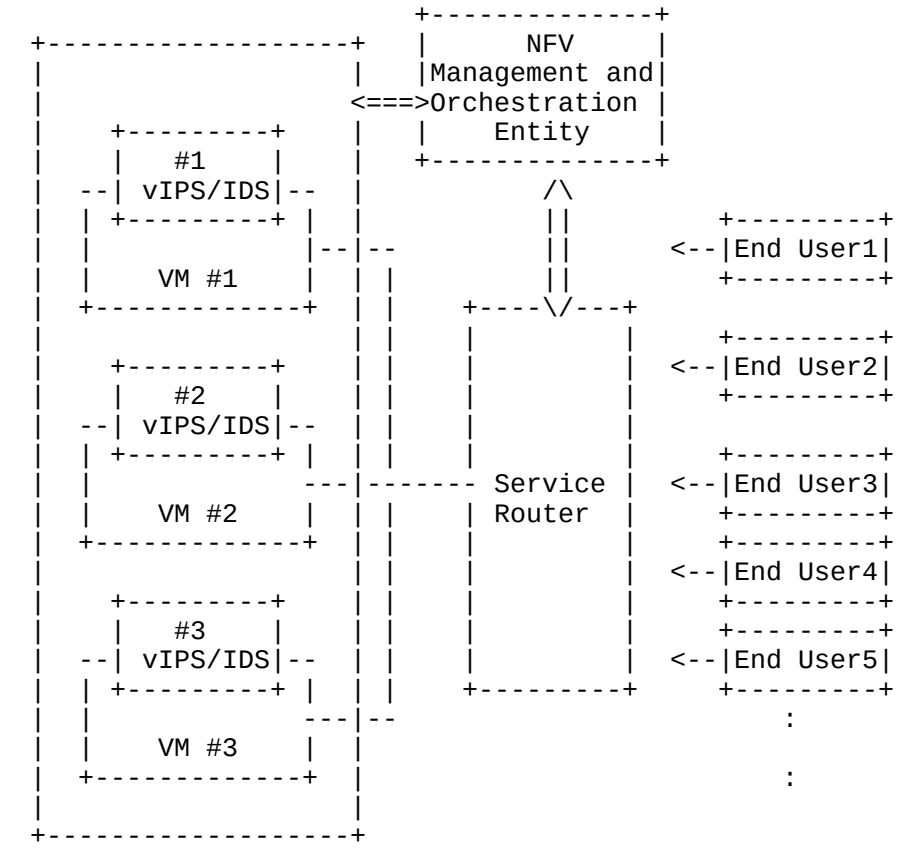

Figure 3: Auto Scaling of Virtual network Function Instances

In this case, the following requirements need to be satisfied:

- Monitoring/fault detection/diagnosis/recovery - appropriate mechanism for monitoring/fault detection/diagnosis/recovery of all components and their states after virtualization, e.g. VNF, hardware, hypervisor;

- Resource scaling - elastic service aware resource allocation to network functions.

4.3. Reliable Network Connectivity between Network Nodes

In the reliable network connectivity between VNFs use case provided by ETSI [NFV-INF-UC], the management and orchestration entities must be informed of changes in network connectivity resources between VNFs. For example, Some network

connectivity resources may be temporarily put in power savings mode when resources are not in use. This change is not desirable since it may have great impact on reachability and topology. Another example, some network connectivity resource may be temporarily in a fault state and comes back into an active state, however some other network connectivity resource becomes permanent in a fault state and is not available for use.

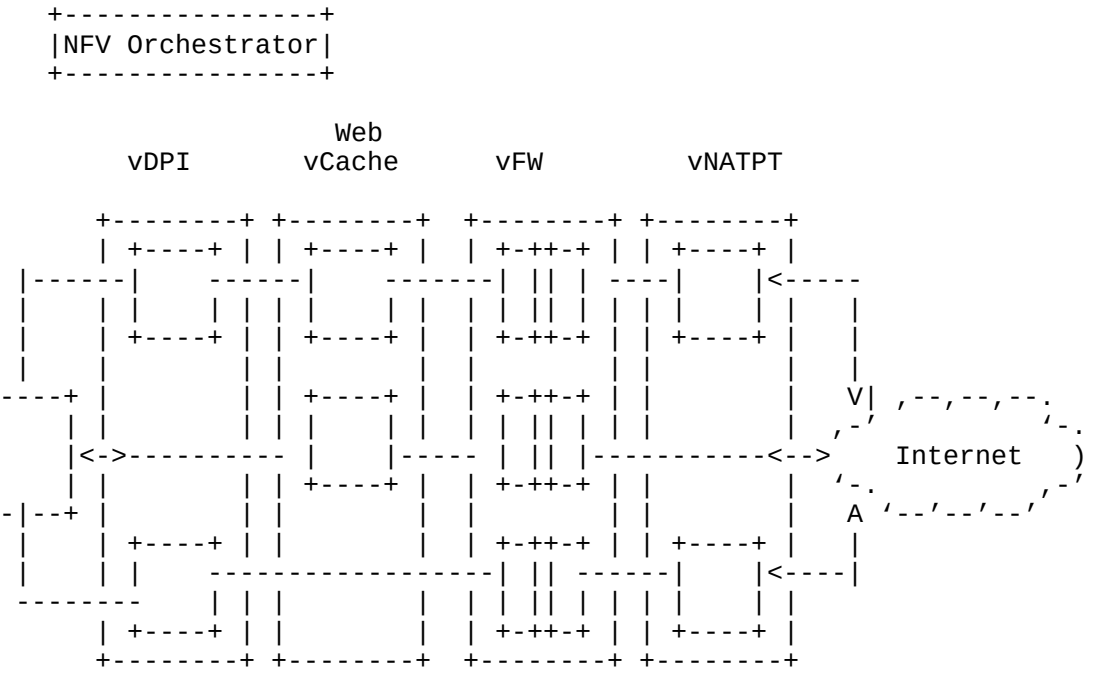

Figure 4: Reliable Network connectivity

In this case, the following requirements need to be satisfied:

- Quick detection of link failures;
- Adding or removing VNF instances;
- Adding or removing network links between VNFs.

4.4. Existing Operating Virtual Network Function Instance Replacement

In the Replacement of existing operating VNF instance use case provided by ETSI [NFV-INF-UC] use case, the Management and Orchestration entity may be configured to support virtualized network function replacement. For example, the Network Service Provider has a virtual firewall that is operating. When the operating vFW

overloads or fails, the Management and Orchestration entity determines that this vFW instance needs to be replaced by another vFW instance.

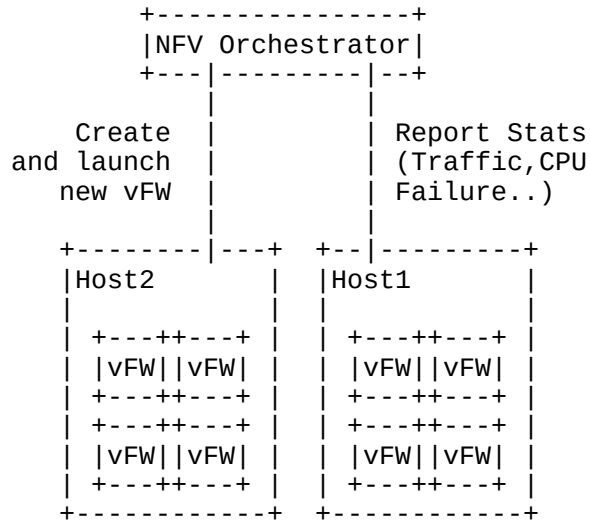

Figure 5: Existing vFW replacement

In this case, the following requirements need to be satisfied:

- Verifying if capacity is available for a new instance of the VNF at some location;
- Instantiating the new instance of a VNF at the location;
- Transferring the traffic input and output connections from the old instance to the new instance. This may require transfer of state between the instances, and reconfiguration of redundancy mechanisms;
- Pausing or deleting the old VNF instance.

4.5. Combining Different VNF Functions (a VNF Set)

A VNF Set is used to assemble a collection of network functions together to support a type of user or end-to-end service. Connectivity between the VNF sets is known as a VNF Forwarding Graph (a graph of logical links connecting VNFs together for steering traffic between network function). To support the reliability of an end-to-end service, except for satisfying the aforementioned basic use case requirements, a VNF Set presents further requirements of reliability as followed:

- o As a whole, any failures (i.e., VNF failures, link failures, performance degradation, etc) of a VNF Set can be detected and recovered in time;

- o Keeping the VNF order and relation unchanged when the VNF Set is updated;

- o The integrated VNF Set performance is not denigrated after it is updated;

4.6. VNF Resilience Classes

Different end-to-end services(e.g., Web, Video, financial backend, etc) have different classes of resilience requirement for the VNFs.

The use of class-based resiliency to achieve service resiliency SLAs, without "building to peak" is critical for operators.

VNF resilience classes can be specified by some attributes and metrics as followed:

- o Does the VNF need status synchronization;

- o Fault Detection and Restoration Time Objective (e.g., real-time, near-real time, non-realtime) and metrics;

- o Service availability metrics;

- o Service Quality metrics;

- o Service reliability;

- o Service Latency metrics for components.

[More description is needed.]

4.7. Multi-tier Network Service

Many network services require multiple network functions to be performed sequentially on data packets. A traditional model for multi-tier service is shown as below, where for each network function, all instances connect to the corresponding entrance point (e.g. LB) responsible for sending/receiving data packets to/from selected instance(s), and steering the data packets between different network functions.

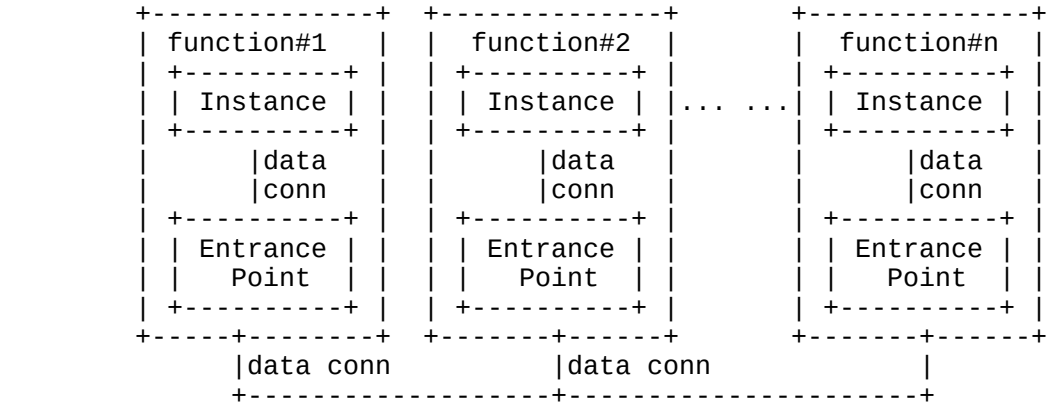

Figure 7: Multi-tier Service

Such model works well when all instances of the same network function are topologically close to each other. However, VNF instances are highly distributed in DC networks, Network Operator networks and even customer premised. When VNF instances are topologically far from each other, there could be many network links/nodes between them for transferring the data packets. For two different VNF instances, it is possible that they are on the same physical server, but the entrance points are many links/nodes away. To improve network efficiency, it is desirable to establish direct data connections between VNF instances, as shown below:

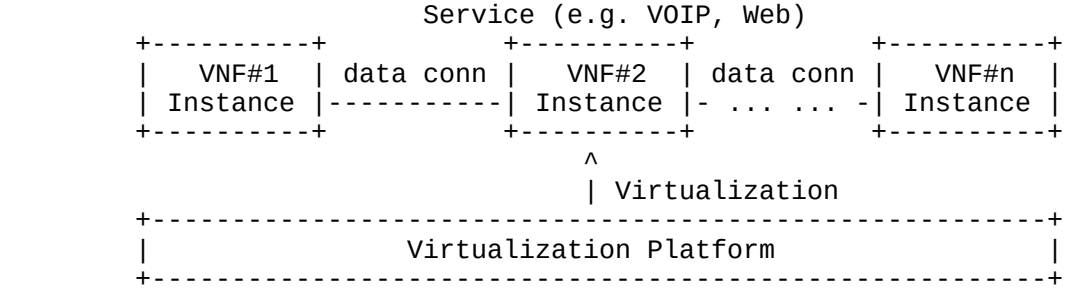

Figure 8: VNF Instances Direct Connection’

In this case, the following requirements need to be satisfied:

o End to end failure detection of VNFs or links for multi-tier service;

- o Keep running service not be influenced during VNF instance transition or failure in the model of VNF instances direct connection.

5. IANA Considerations

This document has no actions for IANA.

6. Security Considerations

TBD.

7. References

7.1. Normative References

7.2. Informative References

[NFV-INF-UC] "Network Functions Virtualisation Infrastructure Architecture Part 2: Use Cases", ISG INF Use Case, June 2013.

[ETSI-HA-USECASE] "Network Function Virtualisation; Use Cases;", ISG NFV Use Case, June 2013.

[NFV-REL-REQ] "Network Function Virtualisation Resiliency Requirements", ISG REL Requirements, June 2013.

[I-D.zong-vnfpool-problem-statement] Zong, N., "Problem Statement for Reliable Virtualized Network Function (VNF) Pool", July 2014.

[RFC7365] Lasserre, M., et al. "Framework for DC Network Virtualization", RFC7365, October 2014.

[RFC5351] Lei, P., Ong, L., Tuexen, M., and T. Dreibholz, "An Overview of Reliable Server Pooling Protocols", May 2008.

Authors' Addresses

Liang Xia(Frank)
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: frank.xialiang@huawei.com

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: bill.wu@huawei.com

Daniel King
Lancaster University
UK

Email: d.king@lancaster.ac.uk

Hidetoshi Yokota
KDDI Lab
Japan

Email: yokota@kddilabs.jp

Naseem Khan
Verizon
USA

Email: naseem.a.khan@verizon.com