
**White Paper**

# Achieving Carrier-Grade OpenStack for NFV

Prepared by

Roz Roseboro
Senior Analyst, Heavy Reading
www.heavyreading.com

on behalf of

WIND RIVER

www.hp.com/go/nfv www.windriver.com

**January 2015**

# Introduction

Network functions virtualization (NFV) offers communications service providers the promise of more rapid service creation, easier management and lower delivery cost. With sustained competitive pressures, including from newer Web-scale players, service providers are responding to this message and moving with urgency to transform their networks and operations.

Many of the components of the platform to deliver NFV have come from community-led developments such as OpenStack and Apache. These have been driven mostly from the enterprise world. Service providers, however, have significantly more stringent requirements than enterprises, meaning work needs to be done in order to make these components truly carrier-grade.

Companies such as HP and Wind River have enhanced and innovated on top of open source solutions, including OpenStack and Linux. They have addressed the key areas of performance, scalability, resiliency, reliability, security and manageability to provide an NFV platform that is carrier-grade.

This white paper is structured as follows:

- **Section II** introduces the concepts of open source as a key component of NFV platforms.
- **Section III** discusses the efforts being undertaken in order to make open source solutions viable in a service provider environment. It addresses how OpenStack, KVM, Linux and vSwitch are being made "carrier-grade."
- **Section IV** shows how HP's integrated NFV platform, built on carrier-grade components, can help service providers confidently transform their networks to support NFV.

# Leveraging Community-Led Development in NFV

Numerous community-led organizations and innovations from the enterprise are increasingly important to service providers. They believe that the open source solutions coming out of OpenStack, OpenDaylight, Apache and others will allow them to avoid vendor lock-in at the infrastructure and platform layers, while also achieving a more robust product deployed on a cloud-based, agile platform. Linux and KVM are tools initially designed for use in enterprise environments, and are also appealing for some service provider use cases. Additional benefits service providers hope to achieve include cost savings, interoperability and greater operational and service agility.

## Service Provider vs. Enterprise Requirements for Software & Platforms

Service providers, however, have significantly higher requirements for their software and platforms than do enterprises. In the same way that a Ford and a Tesla are both cars, but one offers more features at a higher price, carrier-grade solutions provide more availability, serviceability and reliability than do enterprise solutions, in large part due to their need to support SLAs. Service providers expect "Teslas" – and are willing to pay the necessary premium.

Service providers have SLA and regulatory/compliance obligations that lead them to have specific needs in many areas, including the following:

- **Performance –** the ability to process information and traffic: as an example, the network infrastructure must ensure a deterministic interrupt latency of 10µs or less, in order for virtualization to be feasible for the most demanding CPE and access functions.
- **Scalability –** the number of users, transactions, packets that can be supported on a given platform: the platform must automatically and dynamically instantiate new virtual machines (VMs) in response to increases in network traffic or requirements for additional network services ("scale-up"), while also de-instantiating VMs that are no longer needed in order to optimize data center resource allocation and energy consumption.
- **Resiliency –** the ability of a device or network to respond in the event of a failure: in order to respond to failures of physical or virtual elements within the platform, the management software must be able to detect failed controllers, hosts or VMs very quickly and implement hot data synchronization, so that no calls are dropped or data lost when failovers occur.
- **Reliability/availability –** the uptime that a device can achieve: traditional telecom networks based on physical infrastructure enable service providers to guarantee "five-nines" (99.999 percent) uptime for services provided to enterprise customers (and service providers execute service-level agreements based on this reliability). As NFV is progressively deployed in their networks, service providers need to ensure that their new virtualized infrastructure enables the same reliability at the service level.
- **Security –** the ability to keep unauthorized users from impacting device and network operations: telecom networks have security requirements that go beyond typical enterprise installations. For example, in a 4G system, there must be no user traffic that is observable but not encrypted. Similarly, visible user data cannot be stored in the system. In an NFV data center or cloud deployment, operators have to implement efficient multi-tenant isolation and security, so that it's impossible for one subscriber to access the traffic or

data of another subscriber. And the network must fully implement protocols for AAA (authentication, authorization and accounting) security, to prevent unauthorized access, hacking or terrorist attack.

- **Manageability –** usability of the tools offered to configure, provision and operate network devices: the network infrastructure must support hitless software upgrades, hitless patches and integrated backup/recovery systems.

## "Carrier Grade" Is a Must-Have

In order to meet service provider requirements, various industry players are intent on making enterprise solutions "carrier grade." This expression has both positive and negative connotations. On the one hand, "carrier grade" is sometimes perceived to be expensive, cumbersome and unnecessarily complex. On the other, it is thought of as robust, solid and dependable. The real issue is that the cost of failure more than offsets the high initial cost, and "carrier-grade" requirements must be met in order for service providers to deploy solutions with confidence.

"High availability" is a term that's been used in the telecom industry for many years. It basically refers to systems which include enough excess capacity in the design to accommodate a performance decline or a subsystem failure. This is a key feature of carrier-grade networks, but it's not sufficient. Reviewing the list of carrier-grade requirements above, it's clear there are many performance- and security-related constraints, none of which are addressed purely by redundancy in either hardware or software.

It's important not to confuse "high availability" with "carrier grade." The latter is much more demanding, but it's an essential feature of today's telecom networks. Enterprise customers (and to a certain extent consumers) have been conditioned to expect extreme reliability in their network services. Service providers need to continue to meet those expectations as they transition to NFV; otherwise, they run the risk of losing their high-value customers (seeing increased subscriber churn), while their top-line revenue is impacted by significant penalties resulting from SLA violations.

## Open Source Solutions Are Increasingly Preferred for NFV Platforms

Service providers are currently testing NFV, with the underlying platform decision being among the most scrutinized. OpenStack and KVM are emerging as preferred components of an NFV platform. They are seen as cost-effective and scalable alternatives for resource management and virtualization respectively. Because they are based on open source code, service providers expect to avoid vendor lock-in and benefit from rapid innovation coming out of the rich ecosystems that exist around each of them. In addition, products that come from open source efforts tend to cost less than proprietary products from a single vendor.

# Creating a Carrier-Grade NFV Platform

Because OpenStack and KVM were initially developed for enterprise environments, work is being done to enhance them to provide the level of functionality service providers require. Linux, too, is being addressed, as it is the OS that underpins most open source development. As with any solution, service providers demand availability, reliability and high (and predictable) performance and manageability. These are the aspects that must be addressed for a solution to be considered "carrier-grade."

## Making OpenStack Carrier-Grade

OpenStack is responsible for ensuring the optimal placement of VMs. In a telco environment today, VMs are usually IT workloads. Once NFV is implemented, though, virtual network functions (VNFs) will be run as VMs. As such, it is imperative that OpenStack provide automatic failure detection and recovery to ensure the functions are available. Failure restarts needs to be handled much more quickly in telecom environments as well, which means the process needs to be automated.

Service providers, unlike enterprises, have SLAs and regulatory/compliance mandates that compel them to ensure service availability with the requisite quality. Latency and performance requirements dictate that VNFs receive guaranteed resources – unlike in an enterprise environment where over-subscription is acceptable. Non-uniform memory architecture awareness (NUMA) optimization will provide the assurance that VNFs can leverage the server processing resources that it requires.

Service providers are accustomed to robust management tools for their operations, so providing hardened orchestration, management and APIs are critical. OpenStack thus needs to supply APIs to support functions such as graceful shutdown and health check. Unlike enterprises, service providers have BSS/OSS platforms that must interact with the virtual resources. OpenStack's Ceilometer and Monasca projects aim to provide the data aggregation and feeds needed to support service provider billing functions.

Network security requirements present major challenges for telecom infrastructure. Carrier-grade security can't be implemented as a collection of bolt-on enhancements to enterprise-class software; rather, it must be designed in from the start as a set of coordinated, fully-embedded features, including: full protection for the program store and hypervisor; AAA security for the configuration and control point; rate limiting, overload and denial-of-service (DoS) protection to secure critical network and inter-VM connectivity; encryption and localization of tenant data; secure, isolated VM networks; secure password management and the prevention of OpenStack component spoofing.

## Making KVM Carrier-Grade

KVM is a Linux-based, open source hypervisor that is responsible for spinning up and running VMs. As noted above, VMs in a telecom environment will increasingly support VNFs. Because service providers expect the same levels of performance in a virtualized environment as in the physical one, enhancements to the hypervisor element are required. Many VNFs – especially those supporting voice and video services, such as wireless baseband and media caching – require deterministic behavior, driving the need for real-time virtualization, low-latency interrupt handling, preemptible KVM etc. Patches, scripts and tuning need to be added to KVM in order to address the high performance and low-latency requirements of these functions.

The standard implementation of KVM doesn't provide the response time needed to minimize downtime during orchestration operations for power management, software

upgrades or reliability spare reconfiguration. In order to respond to failures of physical or virtual elements within the platform, the management software must be able to detect failed controllers, hosts or VMs very quickly and implement hot data synchronization, so that no calls are dropped or data lost when failovers occur. The system must automatically act to recover failed components and to restore sparing capability if that has been degraded. To do this, the platform must provide a full range of carrier-grade availability APIs (hot sync, VM monitoring, etc.), compatible with the needs of the OSS and orchestration systems and VNFs deployed by the service provider. The software design must ensure there is no single point of failure that can bring down a network component, nor any "silent" VM failures that can go undetected.

## Making Linux Carrier-Grade

Linux is the operating system at the heart of most open source development – the vast majority of which are being undertaken to support enterprise applications. Carrier-grade Linux integrates technologies and capabilities to address needs not found in standard Linux. Providing this type of Linux requires registration with the Linux foundation. As specified by the Linux Foundation, to call a Linux distribution "carrier-grade," it must meet its specifications in the following areas: standards, hardware, serviceability, performance, availability, clusters and security. Enhancements made to meet these specifications must be contributed back into the community.

## Making vSwitch Carrier-Grade

Because switching performance is such an important driver of opex reductions, two approaches have been developed that boost performance while compromising on functionality: PCI Pass-through and Single-Root I/O Virtualization (SR-IOV). As we'll see, though, the functions that are dropped by these approaches turn out to be critical for carrier-grade telecom networks. Fortunately, there is now an alternative solution that provides best-in-class performance as well as these key functions, so the compromises turn out to be unnecessary.

PCI Pass-through is the simplest approach to switching for NFV infrastructure. It allows a physical PCI network interface card (NIC) on the host server to be assigned directly to a guest VM. The guest OS drivers can use the device hardware directly without relying on any driver capabilities from the host OS. SR-IOV, which is implemented in some but not all NICs, provides a mechanism by which a single Ethernet port can appear to be multiple separate physical devices. This enables a single NIC to be shared between multiple VMs.

For service providers that are deploying NFV in their live networks, neither PCI Pass-through nor SR-IOV enable them to provide the carrier-grade reliability they need – namely, six-nines (99.9999 percent) service uptime. To meet their carrier-grade requirements, service providers need to deploy a telco-grade accelerated vSwitch (AVS) that not only delivers high switching bandwidth but also provides the carrier-grade features that are absent from the other two solutions that we've discussed:

- ACL and QoS protection, providing protection against DoS attacks and enabling intelligent discards in overload situations
- Full live VM migration with less than 150ms service impact, instead of the limited "cold migration" option
- Hitless software patching and upgrades
- Link protection with failover in 50ms or less
- Fully isolated NIC interfaces

# HP's Fully Integrated Carrier-Grade NFV Platform

HP and Wind River are two of the companies enhancing the tools created for the enterprise to make them suitable for service provider deployments. In somewhat general terms, open source provides 80 percent of what is needed – suppliers such as HP and Wind River provide the remaining 20 percent.

HP has integrated numerous elements from Wind River's portfolio to deliver an NFV platform based on open source, but hardened to achieve carrier-grade performance. Wind River has benefitted from the riding the innovation paths of non-telecom industries to provide components designed for highly advanced environments. HP has built on its leading position in servers in developing a platform tailored to a telecom environment.

For service providers, HP's integrated offering gives them a single point of accountability in building their next-generation Open NFV platform. It takes responsibility for integrating components from multiple solution providers, eliminating the need for the service provider to manage the process – which is particularly valuable given the specialized IT skills such integrations require.

To help alleviate concerns about using solutions based on open source code, HP is offering HP's OpenStack Indemnification Program. This program is designed to protect qualified customers using HP Helion OpenStack code from third-party patent, copyright and trade-secret infringement claims directed to OpenStack code alone or in combination with Linux. Because it is actively involved in all of the major open source initiatives, HP can ensure that any work it does on behalf of its service provider customers is fed back into the community – resulting in ever more robust solutions that provide the carrier-grade functions and features that service providers demand.

HP, with its long history in IT security, bootstrapped the Security Vulnerability Team and created the OpenStack Security Group. HP continues to lead this group, has authored an overhaul of the OpenStack Security Notes that have been issued, and continues to lead in finding, reporting and fixing many security issues, including releasing a Python-based tool for performing security analysis of OpenStack code.

HP is now focused on solving the challenges faced by enterprises and carriers by building integration with other HP security products and inventing new technologies (such as ephemeral PKI) to solve some of the fundamental challenges with OpenStack security.

# Conclusion

Service providers are keen to begin their transformation toward NFV. The decision regarding the underlying platform used to deliver NFV are critical. Community-led efforts such as OpenStack are seen as important building blocks of the platform, but service providers have much stronger requirements than enterprises, which have led much of the development.

Vendors such as HP and Wind River have and continue to innovate on top of these open source efforts in order to provide a platform that can meet service providers' exacting performance, resiliency, reliability, security and manageability needs. An integrated carrier-grade NFV platform – such as the one from HP – is needed to give service providers confidence that they will be able to achieve the full benefits of NFV.