



Executive Summary

Communication service providers (CSP) demand infrastructures that offer 99.999% availability to their customers. This translates into a downtime of no longer than five minutes per year. There are plenty of cases where network downtime has caused major negative impact on economic and social issues. The cost of long downtimes can be quantified by service level agreement penalty clauses, as well as the inherent opportunity cost in terms of higher customer churn rate and a poor image in the industry.

CSPs are currently migrating from their traditional purpose-built networks to virtualized network function (VNF) infrastructures. This trend brings new opportunities and challenges to both CSPs and their vendors. Lower total cost of ownership (TCO) levels and service agility are naturally attractive, but they do not change service providers' fundamental requirements. Virtual network function or not, they demand carrier-grade, highly available (6 9s) solutions to ensure that mission-critical applications are serviced at 99.999%.

This paper focuses on topics covering carrier-grade network function virtualization infrastructure (NFVI) solutions and their advantages over enterprise-grade versions, as the latter has also targeted the CSP market. The definition of a carrier-grade solution is discussed. The discussions also include technical but more importantly the economic advantages of a carrier-grade NFVI.

ACG Research performed the economic analysis for performance, high availability, and manageability advantages of HPE's Helion OpenStack Carrier-Grade (HPE HCG) solution versus a competing enterprise-grade product over five years. ACG found a 64% TCO based return on investment and total TCO savings of 19% in favor of HPE HCG based on its superior performance. In high availability, a \$1.2B in revenue differential over five years was found for a service provider that decides to deploy HPE HCG instead of an enterprise-grade solution. The carrier-grade manageability of HPE HCG lowers overall operations expense with a marginal savings of \$55K per year for every hour saved in a day with lower labor intervention in the network operations.

Introduction

A communication service provider's (CSP) business can be construed as mission-critical and, therefore, it demands infrastructures that can support this point. Accordingly, the infrastructure's stability becomes a critical factor to ensure uninterrupted delivery of services to customers. This remains true for virtualized network infrastructures, which have to support the classic 5 9s availability for customers' service levels. This 5 9s service availability translates into a 6 9s availability for the network infrastructure and its hardware/software components. In essence, CSPs want the same level of availability as they have had with their purpose-built networks.

Evolution to network virtualization has its known benefits. However, in the area of high availability, it has proven to be challenging. Even with directives from standards bodies, e.g., ETSI¹, many of the current solutions in the market do not satisfy the 6 9s requirements. Enterprise-grade 3 9s solutions are being offered for CSPs' infrastructure market. The 6 9s requirement in the virtualized network infrastructure is necessary for all of its software and hardware: NFVI software (OS, Hypervisor, vSwitches, SDN controllers, OpenStack, middleware), orchestration, VNFs, and servers, switches (leaf/spine), and the underlying physical transport. Achieving 6 9s availability is not trivial. There cannot be a weak link in this infrastructure as network downtime carries unacceptable cost levels.

Redundancy of hardware is usually the norm for protecting it from failure, generally with a 1+1 or N+1 protection schemes. Software failures, however, has been the cause of downtime more often than hardware failures. For software to be carrier-grade, it has to have been designed and coded from the ground up. The discussions in this paper cover the technical and, more importantly, the economics of a carrier-grade versus an enterprise-grade software in a Virtualized Network Infrastructure (VNI). The focus will be on the NFVI software.

Economics of Downtime

Network down time has been called the silent killer that deprives a service provider from its revenue. The revenue loss comes from diminished reputation, security breaches, and work disruptions. Reports from service providers for cost of downtime per server have been anywhere from \$210,000 to \$660,000 to per hour in SLA penalties!

In worst-case scenarios, crippled networks can render emergency services and public safety communications systems useless. The nationwide outage of Canada's Rogers Wireless in October left millions of customers without access to 911, with city officials advising people to use neighbors' landlines or phones at nearby stores in case of an emergency².

Network errors leading to down times directly impact the service providers' ROI, operation expenses (opex), and result in opportunity costs for the service provider. For every 1% loss of annual revenue, the top four North America service providers face \$790M in lost revenues. Considering that the average cost of downtime is estimated to be between 3 to 4% annually, the math to calculate the total loss in revenue is easy and very painful. For Internet-based businesses, the cost levels are also painful. In 2013,

¹ http://www.etsi.org/deliver/etsi_gs/NFV-REL/001_099/001/01.01.01_60/gs_nfv-rel001v010101p.pdf.

² <http://www.cbc.ca/news/canada/rogers-cellphone-outage-highlights-911-vulnerability-1.1959393>.

the Amazon.com website went down for 30 to 40 minutes, costing the company between \$3 and \$4 million dollars³.

The current migration of service providers' networks to a virtualized network infrastructure-based (VNI) operation can create additional challenges in reducing network downtimes. The risks stem from support for new VNFs, a higher level of complexity in debugging and virtual machines moving dynamically across servers. In contrast, purpose-built infrastructures have been in operation for many years with thousands of deployments. These infrastructures have been delivering 6 9s high availability levels. Obviously, there is a lack of adequate real-world sample points for VNI-based infrastructures.

Service providers insist that their network infrastructures adhere to 6 9s high availability requirements, regardless of the technology. The software in VNI, including the NFVI, must be carrier-grade. A discussion of what separates a carrier-grade versus an enterprise-grade VNI, focusing on NFVI, is given in the following section.

Carrier-Grade Software in NFVI

It is useful to examine the classic definitions of high availability and reliability. These terms have different meanings:

1. Reliability: A system is called reliable when it can operate even in presence of a fault. Reliability is usually achieved through redundancy, e.g., 1+1.
2. High availability: It measures the percentage of time when the system is operating according to its specifications. A system that is 99.9999% available has a downtime of no longer than 31 seconds per year. It is important to note that the seconds that make up this time frame do not need to be sequential. Therefore, a system that has 99.9999% availability could see an average of downtime of 0.6 seconds per week or 2.6 seconds per month. For VNIs, this fact is problematic since fault-detection latency and mitigation must have latency levels lower than the aforementioned time frames in downtime.

Availability of a system can be obtained by:

Availability = $MTBF / (MTBF + MTTR)$, with MTBF signifying "Mean Time in Between Failures" and MTTR as "Mean Time To Repair". Minimizing the latter greatly increases the availability of a system.

Existing enterprise-grade NFVI, which only offer a 3 9s availability level, can experience an average of 8.76 hours per year. Deployment of an enterprise-grade NFVI into a service provider's network can cost as high as of \$5.7M per year per server. Clearly, this is not the right solution for this market segment.

For NFVI software, which controls many critical functions (operating system, Hypervisor, vSwitches, orchestration engine (OpenStack), and middleware), carrier-grade quality is a necessity. The discussion, so far, has implied that high availability and carrier grade are interchangeable concepts. For VNI and by

³ <http://www.fool.com/investing/general/2013/08/20/heres-how-much-yesterdays-outage-cost-amazon.aspx>.

extension NFVI, carrier-grade systems must have the following characteristics to survive the requirements of service providers and address the limitations of enterprise-grade NFVI:

1. Fault-tolerant techniques to achieve 6 9s high availability to maximize service uptime. Classic techniques include active/standby, active/active, and advanced self-healing. For active/standby techniques, failover time frame of less than 1 second is necessary to ensure the necessary 6 9s HA level. In addition, if a virtual machine (VM) is used to back up another VM, it must be instantiated on a different physical server. An enterprise-grade solution cannot guarantee the fast failover time frame requirement; it uses load balancing across many VMs and servers to achieve higher availability levels. This band-aid technique, by nature, can lead to a, potentially, substantial increase in TCO; nevertheless, it cannot reach the same level of a 6 9s high availability solution.
2. Performance and scalability to instantiate hundreds of virtual machines benefiting from high-performance vSwitches. Additionally, the ability to increase performance for inter VM communication, 10 to 40 to 100Gbps, is necessary. High-performance solutions must have near line rate throughput and high rates for bare-metal compute.
3. Management software must support in-service introduction of software patches, hitless upgrades, low latency fault-detection, and mitigation and be capable of advanced scheduling of VM instantiation. The software must have built-in security to reduce any risk in violation of the network operation and potential disruptions in service delivery.

All of these points are necessary to enable a robust (5 9s) service delivery to customers. Briefly, the first point was discussed previously as it is one of the top two requirements of service providers (the other is stringent network security). Point number 2 is necessary to scale up/down and in/out for rapid introduction of resources and low latency in data delivery to manage today's video-based network traffic. Point number 3 is highly critical in ensuring a smooth running operation. Specifically, automation of fault-detection and mitigation contribute to rapid restoration of the system to a healthy state with minimal or no disruption to the network. Advance resource scheduling, such as VM instantiation or deletion, can be used to ensure optimizing network resources. Introduction of software upgrades/patches must be done with ease. For example, the use of techniques such as network slicing (a logical instantiation of a network) can be used to ensure continuity in the current network operation while new patches or a new software image are being introduced in the instantiated (parallel) network.

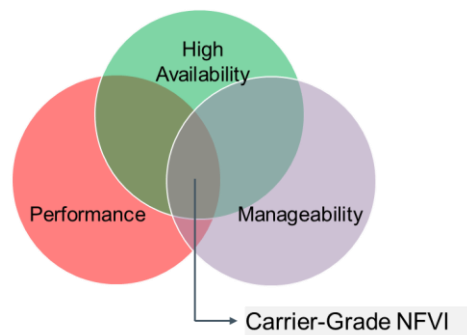


Figure 1. High-Level Requirements for a Carrier-Grade NFVI

Hewlett Packard Enterprise Helion OpenStack Carrier-Grade (HPE HCG) Solution

HPE HCG is a carrier-grade NFVI solution that encompasses the integration of several critical modules beyond OpenStack. It offers a holistic approach to a carrier-grade NFV solution that includes hardened versions of OpenStack, KVM and Linux, VIM, a federated SDN controller, and an accelerated vSwitch. The combination of the modules has value-added carrier-grade features.

Mapping open source software into a carrier-grade quality requires strict guidelines, specifically: high availability, scalability, manageability, and service response characteristics. HPE HCG solution has integrated these requirements in its Linux OS and Hypervisor. For the former, low latency is of utmost importance to ensure time-critical VNF applications, such as packet core’s vEPG, behave in a consistent manner according to 3GPP specifications. Therefore, real-time kernel extensions are necessary to achieve the latency thresholds required by any time-critical function.

Performance of open-source-based products can be boosted and differentiated by introducing innovative techniques. For example, open vSwitch is one such module. It can be enhanced to ensure carrier-grade performance via increased throughput and efficiency. This has been achieved in HPE HCG by its Accelerated vSwitch (AVS), which enables high throughput for data paths.

HPE HCG management layer offers ease of network operation and contributes to its high availability. It offers automation, for example, reassignment of resources at predefined time intervals to shut down servers (to save power consumption) or reassign them for other tasks during off-peak hours to achieve higher efficiency levels. It also offers tools for enhanced security, high performance, and fault mitigation. Specifically, its fault mitigation capabilities contribute to increased high availability, for example, a) via auto-VM-recovery on a VM failure, b) ensuring redundancy for the critical control node and auto-failover. The latter two are not currently available in standard OpenStack implementation. Figure 2 depicts the HPE HCG high-level architecture.

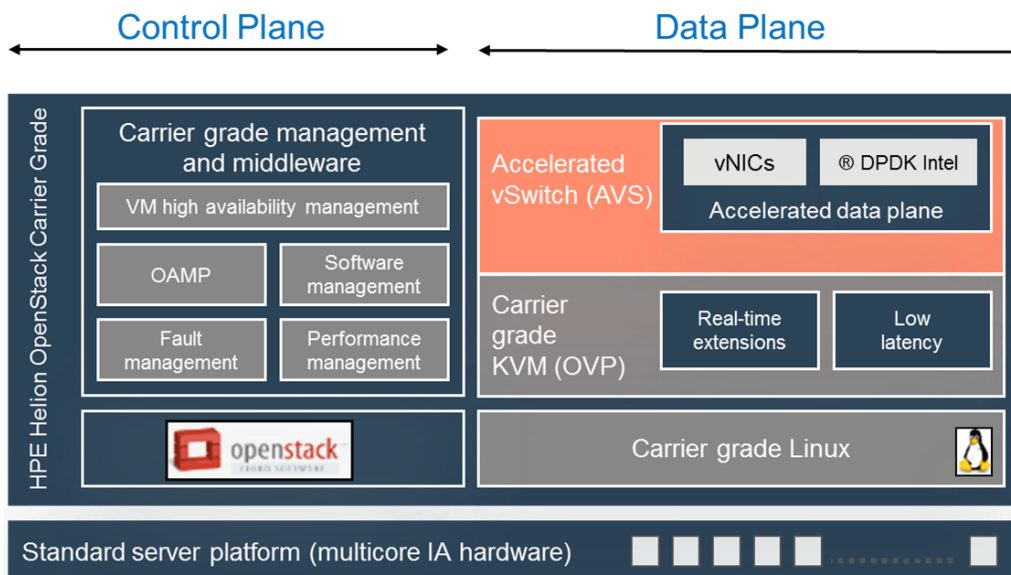


Figure 2. HPE HCG High-Level Architecture

Table 1 highlights how HPE HCG addresses the carrier-grade requirements for an NFVI stack. It also shows additional attributes that are unique to it.

Subject	Features/Description	Benefits
High Availability	<ul style="list-style-type: none"> • Redundant architecture. • Automatic, sub-second failure detection and recovery. • Advanced self-healing techniques. • Ensures VMs of same server groups are created on different compute nodes and/or different availability zones. 	<ul style="list-style-type: none"> • Achieves high availability levels demanded by CSPs. • No single point of failure for continuous network operation. • Minimal disruptions in services to ensure minimal loss, if any, of revenue levels due to SLA clauses.
Performance	Improved vSwitch offers <ul style="list-style-type: none"> • Near line rate network throughput. • Bare metal compute performance. 	Achieves performance levels required by computation intensive or high bandwidth applications and jitter/delay sensitive applications.
Manageability	<ul style="list-style-type: none"> • Performance, fault, high availability, and software management features for tight control of network operation. • Advanced resource scheduling. • Load balancing across processor nodes, selects the processor with lighter loads. • Enhanced security. • Ease of node, software and feature upgrades to update the network operations. 	<ul style="list-style-type: none"> • Ensures high availability for the network. • High efficiency in resource utilization leading to lower capital expenses (capex). • Reduces operational expenditure via reduction in potential security violations. • Optimizes network resources for target applications/customers. • Maximizes automation to reduce MTTR.
End-to-End Integrated, Operational Solution	<ul style="list-style-type: none"> • Preintegrated with other HPE's products such as the NFV Director and HPE's servers. • Pretested/Performance-tested with HPE's partners' VNFs • Use of HPE professional services organization to integrate 3rd party vendors' modules. 	<ul style="list-style-type: none"> • Reduces time, cost and risk in adding new VNFs • Save on network roll out and consequently achieve faster time to market and revenue.
Openness & Flexibility	Based on open standards.	No vendor lock-in for VNF integration.

Table 1. HPE HCG Features and Benefits

ACG Research assessed HPE HCG overall approach and capabilities to its competitors offering enterprise-grade solutions to CSPs. The following radar chart depicts our findings:




Figure 3. High-Level Comparison of NFVI Solutions

Business Analysis of a Carrier-Grade NFVI

To examine HPE HCG economic advantages, ACG Research conducted an analysis of HPE HCG versus a competitor that offers an enterprise-grade solution to CSPs. Typical scenarios were used to assess the economic differences. ACG performed the analysis for performance and high availability. Analysis on its performance was centered on its accelerated vSwitch (AVS). The analysis on high availability focused on comparing potential revenue loss. Although the management layer of HPE HCG is an important piece in its carrier-grade solution, its features basically enable and contribute to the high availability capability of the solution. Its overall capability to ease network operational tasks contributes to a reduction in opex.

Analysis of HPE HCG's Performance

Table 2 summarizes the variables that were assumed for analysis and quantification of HPE HCG's AVS performance advantages. Performance numbers were supplied by HPE and gleaned from publicly available data. Note that the packet size used in this analysis was 256 bytes, which is not realistic in network operations; however, it does offer a good approximation of the relative capex and opex levels.

Table 2 shows the basic assumptions for this case.

Scenario Assumptions: Five Years	
Number of CSP's enterprise customers	10,000
Average peak bandwidth per customer	25Mbps
CAGR in customer acquisition	1%
CAGR in bandwidth	5%
Server	HPE's DL380s with 2 CPUs and 2x40GE interfaces
Number of VMs per processor	1
Number of cores per processor	12
Standard vSwitch performance assumption	4.64 millions of packets per second (MPPS)
Standard vSwitch performance (measured by HPE)	12MPPS

Table 2. Scenario Assumptions

TCO Results and Cash Flow Analysis

TCO based ROI analysis was conducted over five years, and the study found a TCO based ROI level of **64%** in favor of HPE HCG, based on its superior performance levels. Figures 4, 5 and 6 show the TCO advantages of HPE HCG’s performance in capex and opex. The cumulative TCO savings are also shown.

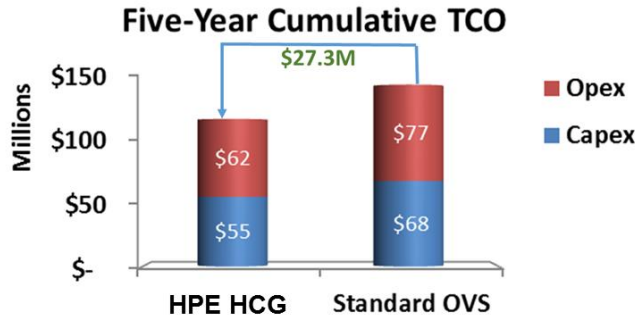




Figure 4. TCO Comparison



Figures 5 and 6. Cash Flow Results

HPE HCG’s High Availability and Manageability

Table 3 summarizes the variables that were assumed for analysis and quantification of HPE HCG’s high availability advantages. In this case, the main objective is to focus on minimization of revenue loss based on network downtime and SLA imposed penalties.

Table 2 shows the basic assumptions for this case.

Scenario Assumptions: Five Years	
Number of Enterprise Customers served by the CSP	10,000
Average annual revenue per enterprise customer	\$50,000
CAGR in customer acquisition	1%
CAGR in bandwidth	5%
SLA revenue penalty per one hour of network downtime	5%
Churn rate per year per hour of network downtime (after first year)	1%
HPE HCG service availability	99.999%
Enterprise-grade service availability	99.9%

Table 3. Scenario Assumptions

The analysis showed a \$1.2B in revenue differential over five years (Figure 7).




Figure 7. Revenue Differentials

A simple methodology in quantifying the advantages of network manageability is in opex savings as a result of lower labor costs. This is attributed to network automation, enhanced security and fault-mitigation. In this case, the computation in incremental annual savings is straightforward. Assuming an average labor rate of \$150 (loaded) for every hour saved per day in network operations, a savings of approximately \$55K per year is gained. Naturally, the labor rate varies per geographical regions and the overall savings depend on the network size and its associated network operations center.

Conclusion

Service providers demand solutions that ensure service continuity to their customers, adhere to their own respective SLA clauses and bring predictive key performance indicators to their business operations. To satisfy these stringent demands, vendors must offer solutions that are carrier-grade to minimize network downtime and lower operational costs.

With the advent of NFV, network infrastructures solutions are more software-centric. Therefore, the software itself must be resilient and also be able to control the reliability of the underlying off-the-shelf hardware. One of the most critical software components is the NFVI, which is currently based on standard OpenStack.

HPE HCG offers technical and economic benefits of a carrier-grade NFVI solution and has an advantage over an enterprise-grade one in three main areas: performance, high availability and manageability. ACG Research found a 64% TCO based ROI and total TCO savings of 19% in favor of HPE HCG. In high availability, a \$1.2B in revenue differential over five years was found for a service provider that deploys HPE HCG instead of an enterprise-grade solution. HPE HCG's carrier-grade manageability lowers overall opex with a marginal savings of \$55K per year for every hour saved in a day with lower labor intervention in the network operations.

About ACG Research: ACG Research is an analyst and consulting company that focuses in the networking and telecom space. We offer comprehensive, high-quality, end-to-end business consulting and syndicated research services. Copyright © 2016 ACG Research. www.acgcc.com.