WIND™

AN INTEL COMPANY

Carrier Grade Performance and Reliability in Network Virtualization

How It's Measured and Why It Matters

AGENTS OF CHANGE™

EXECUTIVE SUMMARY

Independent test specialist The Tolly Group recently validated the performance and reliability of Wind River® Titanium Server™ Network Functions Virtualization infrastructure (NFVI) software powering the Hewlett Packard Enterprise (HPE) Helion OpenStack Carrier Grade solution, employing a test bed from Spirent Communications, Spirent TestCenter. The test results confirm that Titanium Server delivers the carrier grade performance and reliability that service providers require for network virtualization.

Tolly found that the HPE Helion OpenStack Carrier Grade delivered consistent, high throughput and robust responses to failures. Integral to the HPE solution, Titanium Server's integrated accelerated virtual switch (AVS) achieved:

- Line-rate 10 GbE forwarding performance for a single virtual network function (VNF) at a frame size of 256 bytes
- Line-rate 10 GbE forwarding performance for multiple VNFs at a frame size of 512 bytes
- Consistent performance across 10 test runs
- Consistent low latency of less than 50 microseconds across all packet sizes
- Link aggregation group (LAG) failover in 27 milliseconds
- Compute failure detection and recovery in 1 second

This white paper reviews Tolly's findings and shows how high-performance virtual switching and carrier grade reliability relate to reducing operational expenses and protecting revenues for service providers.

TABLE OF CONTENTS

INTRODUCTION

There is an inherent challenge in the concept of Network Functions Virtualization (NFV): How can components designed for the enterprise deliver telco-grade services? In principle, network virtualization relies on commercial off-the-shelf (COTS) hardware, software-based network functions, and open source technology—none of which was originally designed for carrier grade communication networks.

Compared to enterprise networks, telecom networks are built to guarantee much higher service availability. Traditional physical telecom infrastructure delivers five nines (99.999%) reliability, which equates to about 5.26 minutes of service downtime per year or 25.9 seconds of downtime per month. By contrast, standard enterprise-grade systems typically guarantee only up to three nines (99.9%) availability, which translates into 8.76 hours of service downtime per year or 43.8 minutes per month.

As service providers transition to network virtualization, they are looking for the benefits of agility, flexibility, and cost reduction, but they can't afford to sacrifice network performance and reliability. Ensuring carrier grade service uptime is critical to protecting top-line revenue and minimizing operating costs.

So the question becomes: How can COTS hardware perform at the level of purpose-built hardware? The answer lies in the performance capabilities of the NFVI.

NFVI performance greatly determines the extent to which service providers can minimize operating costs and protect top-line revenue. By benchmarking critical components within the NFVI—such as the virtual infrastructure manager (VIM), virtual switch (vSwitch), virtual compute, virtual storage, and virtual network nodes—service providers will gain confidence in the level of performance and reliability they can achieve for services delivered via VNFs hosted on the platform.

Independent test specialist The Tolly Group evaluated the performance and reliability of HPE Helion OpenStack Carrier Grade powered by Titanium Server NFVI software. As the integral NFVI platform, Titanium Server is the engine that drives HPE's NFV solution. Titanium Server's integrated AVS provides the core networking functions among all virtual machines (VMs) deployed on the HPE server.

Since Titanium Server determines the performance and reliability of HPE Helion OpenStack Carrier Grade, the independent benchmark test proves the carrier grade capabilities of the Wind River software as well as the HPE NFV solution. Tolly observed that as part of HPE's NFV system Titanium Server delivered high throughput, consistent performance, and robust fault detection and recovery.

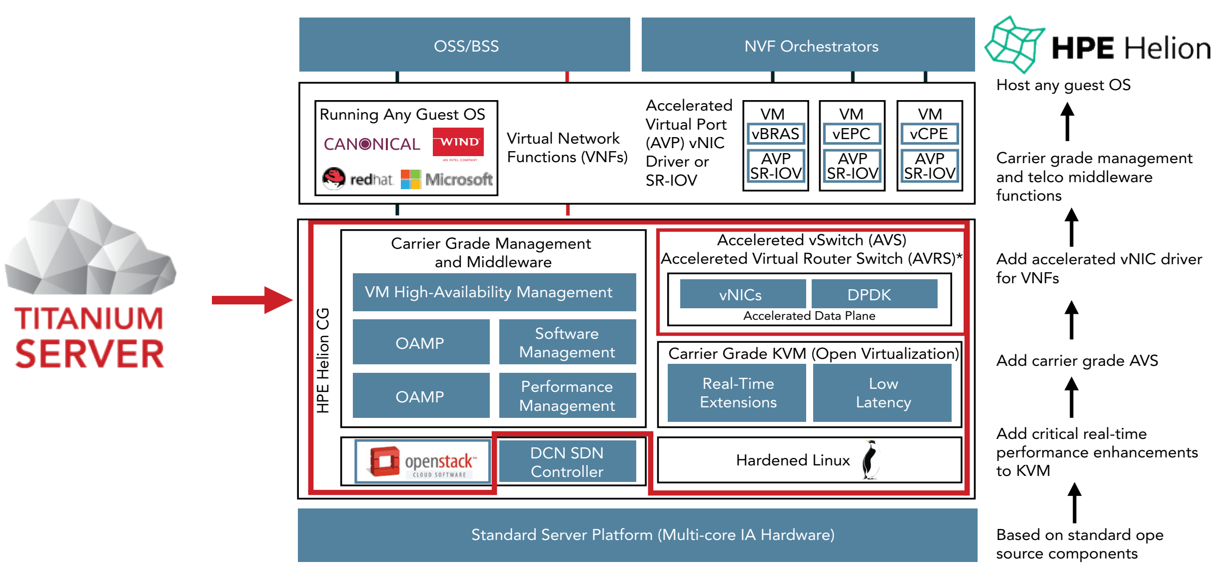

Figure 1. HPE Helion OpenStack Carrier Grade powered by Titanium Server

TEST RESULTS FOR TITANIUM SERVER PERFORMANCE

The Tolly Group measured Layer 2 10 GbE forwarding performance in three virtual switching scenarios: AVS port forwarding, single VNF forwarding, and VNF-to-VNF forwarding. The test setup comprised two physical 10 GbE ports from the Spirent TestCenter into the compute system. The VNFs were Wind River Linux 6 guest VMs, each with a Layer 2 DPDK forwarding application. Each VM used three cores: one for each of the two virtual 10 GbE ports on the VMs and a third management core for the Linux core of the VM, which was shared across all VMs. Titanium Server's AVS was allocated two cores.

10 GbE Forwarding Performance

The forwarding performance results across the three virtual switching scenarios were as follows:

- In the simple AVS port-forwarding scenario, Titanium Server achieved 99.9% or 100% of the theoretical maximum throughput at a frame size of 128 bytes and above.
- In the single VNF forwarding scenario, Titanium Server achieved 99.9% or 100% of the theoretical maximum throughput at a frame size of 256 bytes and above.
- In the VNF-to-VNF forwarding, Titanium Server achieved 99.9% or 100% of the maximum throughput starting at a frame size of 512 bytes and above.

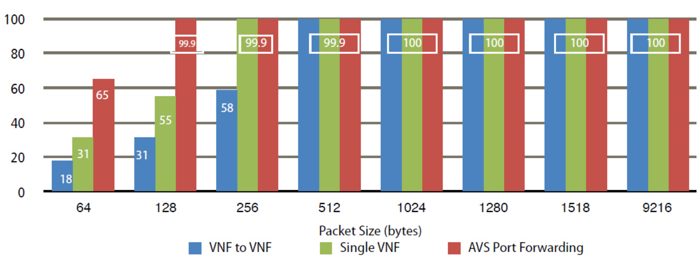

Figure 2. Layer 2 10 GbE forwarding performance

Consistency

Tolly engineers ran the single VNF forwarding test 10 different times and compared the performance results. Across 10 test runs, the highest variance in results was 0.46% at a frame size of 256 bytes. At higher frame sizes of 512 bytes and above, there was no variation across the test runs.

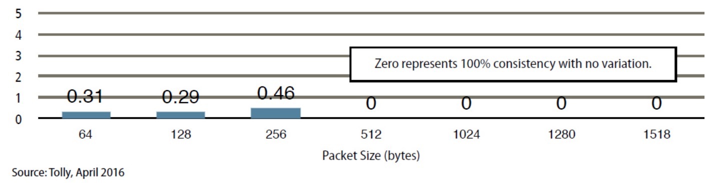

Figure 3. Layer 2 10 GbE forwarding performance consistency

Latency

Latency performance was evaluated based on results from the single VNF forwarding test. The average latency across all packet sizes was consistently less than 50 microseconds, as 99% of all results were below this time measurement. The traffic profile for the latency test was as follows: 80% of line rate for 256-byte and above frame sizes, 44% of line rate for 128-byte frame size, and 25% of line rate for 64-byte frame size.

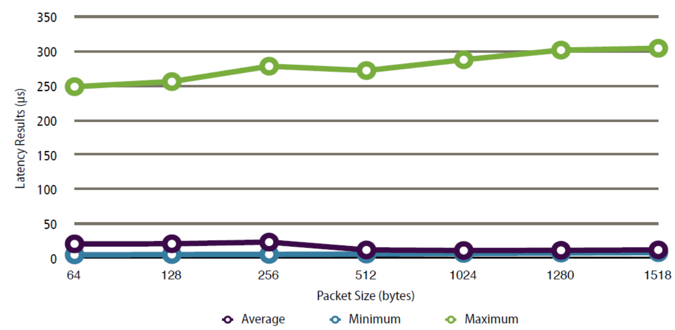

Packet Size (bytes)	Latency Distribution by Various Ranges (Percentage of the Total Packets)							
	<50µs	50-75µs	75-100µs	100-150µs	150-200µs	200-300µs	300-400µs	>400µs
64	99.42%	0.09%	0.1%	0.21%	0.18%	0.01%	0%	0%
128	99.43%	0.08%	0.1%	0.19%	0.18%	0.01%	0%	0%
256	99.37%	0.1%	0.1%	0.2%	0.2%	0.01%	0%	0%
512	99.41%	0.1%	0.1%	0.2%	0.16%	0.01%	0%	0%
1024	99.31%	0.1%	0.1%	0.21%	0.19%	0.06%	0%	0%
1280	99.29%	0.1%	0.1%	0.21%	0.19%	0.07%	0%	0%
1518	99.26%	0.1%	0.1%	0.22%	0.18%	0.1%	0%	0%

Source: Tolly, April 2016

Figure 4. 10 GbE Layer 2 latency

INTERPRETING THE PERFORMANCE BENCHMARK RESULTS

The independent test results show Titanium Server achieves predictable high performance across different packet sizes and exhibits consistently low latency, which are essential requirements for carrier grade communications systems. Titanium Server achieved an impressive line-rate 10 GbE at frame sizes of 512 bytes and

WIND
AN INTEL COMPANY

above across all three test scenarios. Equally important is the fact that Titanium Server achieved the strong performance results by employing only two processor cores of the server platform, which liberates processing capacity for VNFs delivering customer services.

The test bed's three performance-forwarding scenarios represent varying degrees of sophistication. The most basic scenario is AVS port forwarding, where traffic simply flows from the TestCenter to the AVS and back, which establishes a baseline for AVS performance in a test environment. In this test, Titanium Server hit line-rate forwarding performance at the relatively small packet size of 128 bytes. While this is an excellent test result, it does not necessarily reflect Titanium Server's performance expectations in more realistic NFV scenarios.

The key performance results are from the single VNF and the VNF-to-VNF forwarding scenarios, because they indicate how Titanium Server will behave in real-world NFV systems where there are multiple VMs running VNFs. In the simple, single VNF-forwarding test, where traffic flows from the TestCenter to the AVS to the VNF and back again, it is impressive that Titanium Server clocked line-rate 10 GbE forwarding with 256-byte packet sizes using only two processor cores. High performance at a 256-byte frame size is important because it is one of the most commonly used frame sizes for data traffic in communications networks, which indicates Titanium Server's capability to meet service provider requirements.

The VNF-to-VNF forwarding test is more complex and the results are more indicative of performance in real-world NFV scenarios. Referred to as the "bunny ears" test, traffic comes in to Titanium Server and its AVS sends the traffic on to a VNF, then it goes back to the AVS and is sent to another VNF, then back to the AVS and the TestCenter. Titanium Server performed extremely well with line-rate 10 GbE at a 512-byte frame size. And with Titanium Server's AVS using just two cores, the result is even more impressive.

"Not only was the performance line rate for most of the frame sizes under consideration, the performance consistency across multiple runs was one of the best observed in the public tests for NFVI benchmarking."

—Gurpreet Singh, Product Manager, NFV Solutions at Spirent Communication

In any telecommunications system, performance consistency is imperative for delivering carrier grade services. The very low variance observed across 10 test runs proves that the performance of Titanium Server is consistent and predictable. This validation is particularly important given the elastic nature of virtualized environments compared to hardware-based systems.

Latency directly impacts performance and can degrade many services if it is too high, which is why it is important to understand the latency characteristics of the NFVI in a virtualized environment. As the test results show, Titanium Server performed with very low latency—less than 50 microseconds on average—and the latency results were consistent throughout the test.

SYSTEM UNDER TEST: HPE HELION OPENSTACK CARRIER GRADE, VERSION 2.0 PATCH 3

The hardware compute server in the NFV platform comprised one HPE DL360 Gen9 server with dual Intel® Xeon® E5-2640 v3 @2.6 GHz eight-core processors, 128 GB RAM, dual HPE Ethernet 10 Gb 2-port 560SFP+ Adapter network cards. The software was HPE Helion OpenStack Carrier Grade, version 2.0 patch 3, powered by Wind River Titanium Server software.

The test tool was Spirent TestCenter N4U with test module MX-10G and TestCenter software version 4.59. One Spirent TestCenter physical ports used 40 IP addresses while the other physical port used 50 IP addresses. The test relied on an RFC 2544 throughput test wizard and the frame loss threshold was set at 0.001%.

TEST RESULTS FOR TITANIUM SERVER FAULT DETECTION

To validate the reliability of Titanium Server within the HPE Helion OpenStack Carrier Grade environment, Tolly engineers set up five failure scenarios using the single VNF forwarding test and measured how long it took for the system to recover as well as the impact, if any, on the traffic at the point of failure. Tolly found that all failure situations were detected and recovered.

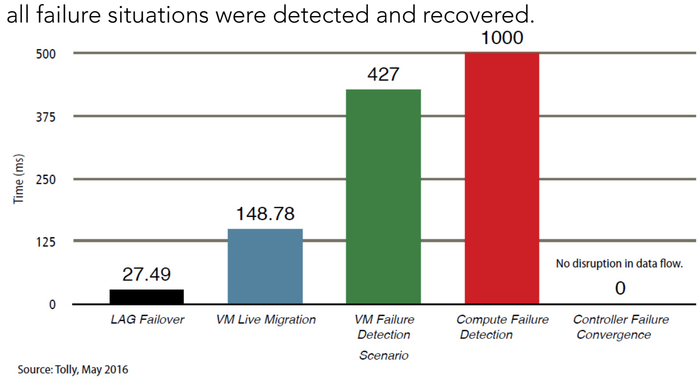

Figure 5. Carrier Grade fault detection scenarios

- **Link Aggregation Group (LAG) failover:** A LAG comprising two 10 GbE interfaces was configured between Titanium Server's AVS and a VNF. When engineers removed one active link, the traffic convergence time was 27.49 milliseconds.
- **VM live migration:** Engineers manually triggered the VM live migration in the GUI-based management console of the HPE Helion OpenStack Carrier Grade controller, and the traffic convergence time was 148.78 milliseconds.
- **VM failure detection:** Engineers manually inserted a critical fault that caused the VM to fail. The fault was detected in 427 milliseconds.
- **Compute failure:** Engineers used HPE's Integrated Lights Out (iLO) management processor to turn off the physical compute server. The system detected the failure in 1 second and restarted the VNF on a backup compute source.
- **Controller failure:** The system had one active controller and one backup controller. Engineers turned off the active controller and the backup controller activated immediately, resulting in a "hitless" failover and no traffic loss.

INTERPRETING THE FAULT DETECTION RESULTS

Tolly's evaluation of fault detection and recovery across a variety of scenarios clearly shows that Titanium Server delivers carrier grade service uptime. The results are impressive because the system not only detected all the test failures quickly but also took rapid remedial actions.

According to Spirent, it is important to evaluate NFVI resiliency based solely on NFVI components rather than including VNFs. Although VNFs influence end-to-end service reliability and availability, because the instantiation time for VNFs varies among different vendors they don't reveal much about the fault detection and tolerance of the NFVI. This is why the test focused on fault detection time and avoided measurements dependant on the VNF's fault-handling ability.

In the case of a link failure, Titanium Server's AVS minimizes service downtime by automatically switching to another link. The AVS detected a link fault and switched from one link to another in 27 milliseconds—very fast compared to standard off-the-shelf operating systems. Typically, competing systems perform LAG failover in 50 milliseconds.

Titanium Server delivers best-in-class performance when performing VM live migration. In less than 150 milliseconds, the system can transfer live VMs from one server to another, thereby ensuring maximum uptime for the services running on those VMs. By contrast, enterprise-grade systems can take from 200 milliseconds up to 2,000 milliseconds to migrate VMs.

If a VM crashes, Titanium Server can detect the failure in just 427 milliseconds and take immediate action to resolve the fault. Competing systems can take minutes to detect a VM failure, and most do not have the capability to recover the VM automatically. Similarly, if an entire compute server goes down Titanium Server can recover in 1 second or less by restarting the affected VMs on another compute node. In other systems, compute fault detection typically takes minutes rather than 1 second, and the only action instigated is a fault notification rather than full recovery.

And in the case of a controller node failure, the test results highlight the importance of Titanium Server's redundant controller design. If a controller node fails, the second controller instantly takes over and no traffic is lost.

	Enterprise IT Platform Capability	Carrier Grade Cloud Requirements	Titanium Server
Detection of failed VM	> 1 minute	< 1 s	500 ms
Detection of failed compute node	> 1 minute	~ 1 s	1 s
Recovery from control node failure	No support	< 25 s	< 25 s
vSwitch performance	1–2 Gbps	Line rate with minimum core utilization	Line rate with single core (512 B/frame)
Network link failure detection	Depends on Linux distribution	50ms	50ms
Live migration for DPDK- based VMs	No support	Full support	Full support: 200 ms

Figure 6. Meeting telco reliability requirements

CARRIER GRADE PERFORMANCE AND RELIABILITY MEAN GOOD BUSINESS FOR SERVICE PROVIDERS

Titanium Server's high performance and carrier grade reliability translate into real business value for service providers because these capabilities are significant factors in reducing operational expenditure (OPEX) and protecting revenue.

Service Provider Operating Costs Correlate to vSwitch Performance in Network Virtualization Implementations

Higher vSwitch performance within the NFVI maximizes VM density, which enables service providers to serve more customers per server blade, in turn lowering operating costs. Since the vSwitch runs on the same server platform as the VNFs, fewer processor cores are dedicated to switching traffic, meaning more system capacity is available for VMs running revenue-generating service workloads. Thus a high-performing vSwitch contributes to reduced OPEX by enabling service providers to serve more subscribers per NFVI platform.

A recent customer use case for a virtualized media gateway demonstrates the power of Titanium Server's performance and what that means for OPEX. The system configuration was a typical 28-core platform running one VM per core and the VNF was a virtualized media gateway. The customer required a bandwidth of 3.5 Gbps per core or 6.8 million packets per second per core. Using Open vSwitch (OVS), the service provider would have needed 23 cores to switch the traffic consumed by one core running a single VM, resulting in a VM density of one. But with Titanium Server, the service provider could use just 10 cores for switching traffic consumed by 17 cores, each running a VM. In this case, Titanium Server produced a massive 17x improvement in VM density, resulting in a significant reduction in OPEX.

In Addition to High Performance, Carrier Grade Reliability Is a Business Imperative for Service Providers

Any amount of service downtime can result in onerous financial penalties imposed by national regulators or incurred by missing service level agreement (SLA) commitments. Service providers need to protect top-line revenue streams in the transition to network virtualization by delivering carrier grade service quality and avoiding SLA penalties.

While difficult to calculate, service downtime can have a severe financial impact on service providers. According to Heavy Reading, network outages and service degradations cost mobile operators about $20 billion per year, and operators estimate they spend an average of 1.7% of annual revenue to handle such incidents. Then there are the unquantifiable costs, such as lost confidence among customers and customer churn that are damaging to service providers over the long run.

CARRIER GRADE PERFORMANCE EXTENDS ACROSS TITANIUM CLOUD ECOSYSTEM

The independent evaluation by The Tolly Group is an important milestone for the partnership between HPE and Wind River, with Titanium Server positioned as an integral component in the Helion OpenStack Carrier Grade product. But the carrier grade performance and reliability of Titanium Server is not limited to HPE's hardware platform. Wind River has established the Wind River Titanium Cloud™ partner ecosystem to facilitate the availability of interoperable hardware and software products optimized for rollout with Titanium Server.

Partners include nearly 40 vendors of hardware infrastructure, VNFs, service orchestrators, and operational and business support systems (OSS/BSS). The same carrier grade capabilities that Tolly independently observed in the HPE NFV server are replicable across the entire Titanium Cloud ecosystem. Any hardware platform partner that relies on Titanium Server will be able to achieve the same carrier grade results. Examples include Adlink, Dell, Huawei, and Kontron.

The Wind River ecosystem extends the applicability of Titanium Server across a range of NFV use cases and has already attracted service providers such as China Unicom. The Chinese operator is currently evaluating Titanium Server as part of an NFV strategy initially focused on three use cases: virtual customer premises equipment (vCPE), virtual evolved packet core (vEPC) and virtual IP multimedia subsystem (vIMS).

HOW DOES TITANIUM SERVER ACHIEVE CARRIER GRADE NETWORK VIRTUALIZATION?

As the only commercial cloud that achieves carrier grade reliability, Titanium Server enables carrier grade quality in network virtualization. Unlike off-the-shelf enterprise-grade software, Titanium Server is designed to deliver the high performance service providers require without sacrificing service uptime. Titanium Server's AVS delivers 40 times better performance than OVS, which is what is typically used in enterprise data centers and originally designed for IT applications.

Titanium Server's architecture is designed with open source projects, including OpenStack, KVM, Linux, and Ceph, which have all been hardened for higher performance. Titanium Server supports enhanced performance awareness (EPA) in OpenStack, which enables the addition of extensions that specify optimal ways of performing CPU modelling, core pinning, and optimal VNF placement.

But high performance via EPA is not enough. Service providers require performance and service uptime. To achieve carrier grade reliability, Wind River has also added OpenStack extensions to perform extra monitoring, fault management, alarms, and fault logging as well.

Through optimization and rigorous testing, Titanium Server delivers both the performance and the reliability that service providers demand for carrier grade NFV deployments.

CONCLUSION

Performance and reliability are business imperatives for service providers under any circumstances. In the transition to network virtualization, however, service providers especially require high performance to reduce operating costs and guaranteed service uptime to protect revenues. They need to know that the services they deliver via VNFs will not suffer as a result of being virtualized.

The Tolly Group's independent evaluation of HPE Helion OpenStack Carrier Grade powered by Titanium Server proves the carrier grade credentials Wind River has long claimed. The results show that as part of HPE's NFV solution Titanium Server delivers consistent, high throughput, and robust responses to failures. The test results also suggest that the Wind River NFVI solution resolves an inherent challenge of NFV by proving that network virtualization can indeed deliver carrier grade communications services.

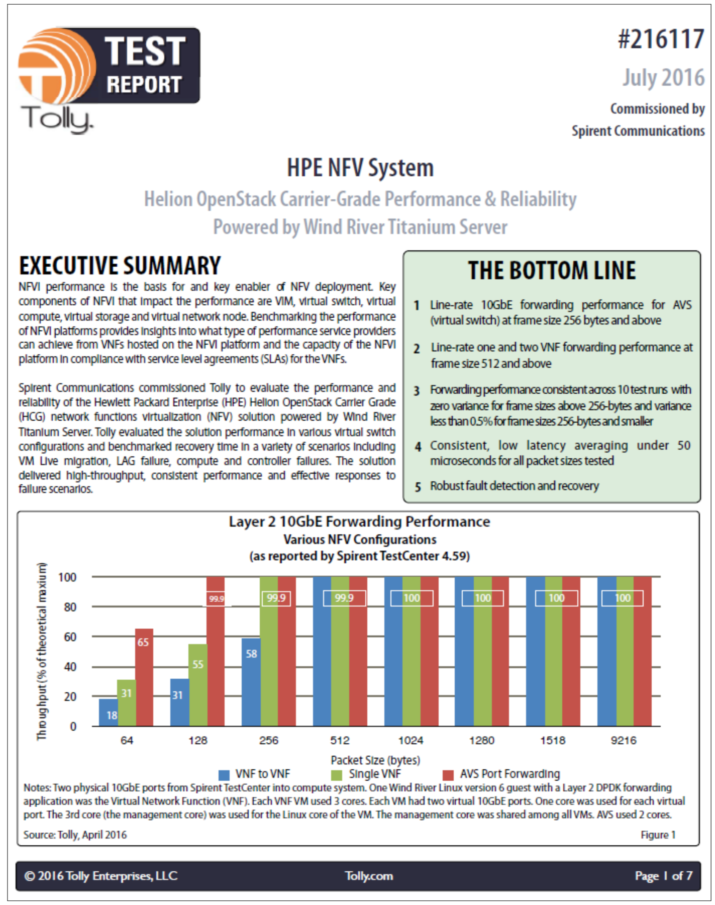

TEST REPORT

Tolly.

#216117

July 2016

Commissioned by

Spirent Communications

HPE NFV System

Helion OpenStack Carrier-Grade Performance & Reliability

Powered by Wind River Titanium Server

EXECUTIVE SUMMARY

NFVI performance is the basis for and key enabler of NFV deployment. Key components of NFVI that impact the performance are VIM, virtual switch, virtual compute, virtual storage and virtual network node. Benchmarking the performance of NFVI platforms provides insights into what type of performance service providers can achieve from VNFs hosted on the NFVI platform and the capacity of the NFVI platform in compliance with service level agreements (SLAs) for the VNFs.

Spirent Communications commissioned Tolly to evaluate the performance and reliability of the Hewlett Packard Enterprise (HPE) Helion OpenStack Carrier Grade (HCG) network functions virtualization (NFV) solution powered by Wind River Titanium Server. Tolly evaluated the solution performance in various virtual switch configurations and benchmarked recovery time in a variety of scenarios including VM Live migration, LAG failure, compute and controller failures. The solution delivered high-throughput, consistent performance and effective responses to failure scenarios.

THE BOTTOM LINE

1 Line-rate 10GbE forwarding performance for AVS (virtual switch) at frame size 256 bytes and above

2 Line-rate one and two VNF forwarding performance at frame size 512 and above

3 Forwarding performance consistent across 10 test runs with zero variance for frame sizes above 256-bytes and variance less than 0.5% for frame sizes 256-bytes and smaller

4 Consistent, low latency averaging under 50 microseconds for all packet sizes tested

5 Robust fault detection and recovery

Layer 2 10GbE Forwarding Performance
Various NFV Configurations
(as reported by Spirent TestCenter 4.59)

Notes: Two physical 10GbE ports from Spirent TestCenter into compute system. One Wind River Linux version 6 guest with a Layer 2 DPDK forwarding application was the Virtual Network Function (VNF). Each VNF VM used 3 cores. Each VM had two virtual 10GbE ports. One core was used for each virtual port. The 3rd core (the management core) was used for the Linux core of the VM. The management core was shared among all VMs. AVS used 2 cores.

Source: Tolly, April 2016 Figure 1

© 2016 Tolly Enterprises, LLC Tolly.com Page 1 of 7

Figure 7. Tolly report #216117, April 2016

Wind River is a global leader in delivering software for the Internet of Things. The company's technology is found in more than 2 billion devices, backed by world-class professional services and customer support. Wind River delivers the software and expertise that enable the innovation and deployment of safe, secure, and reliable intelligent systems.

©2016 Wind River Systems, Inc. The Wind River logo is a trademark of Wind River Systems, Inc., and Wind River and VxWorks are registered trademarks of Wind River Systems, Inc. Rev. 11/2016