



White Paper

TESTING METHODOLOGIES FOR VALIDATING NFV ENVIRONMENTS

October 2013

SPIRENT

1325 Borregas Avenue
Sunnyvale, CA 94089 USA

Email: sales@spirent.com
Web: <http://www.spirent.com>

AMERICAS 1-800-SPIRENT • +1-818-676-2683 • sales@spirent.com

EUROPE AND THE MIDDLE EAST +44 (0) 1293 767979 • emeainfo@spirent.com

ASIA AND THE PACIFIC +86-10-8518-2539 • salesasia@spirent.com

© 2013 Spirent. All Rights Reserved.

All of the company names and/or brand names and/or product names referred to in this document, in particular, the name “Spirent” and its logo device, are either registered trademarks or trademarks of Spirent plc and its subsidiaries, pending registration in accordance with relevant national laws. All other registered trademarks or trademarks are the property of their respective owners.

The information contained in this document is subject to change without notice and does not represent a commitment on the part of Spirent. The information in this document is believed to be accurate and reliable; however, Spirent assumes no responsibility or liability for any errors or inaccuracies that may appear in the document.

Testing methodologies for validating NFV environments

CONTENTS

Introduction.....	1
Drivers for virtualizing network functions	2
Challenges to the development of NFV	3
The challenges in validating NFV environments	4
Functional and performance testing of network functions.....	4
Validating reliability of network services	4
Ensuring portability of VMs and stability of NFV environments.....	4
Active and passive monitoring of virtual networks	4
Testing methodologies for validating NFV	5
Functional and performance testing of network functions.....	5
Validating reliability of network services.....	7
Ensuring portability of VMs and stability of VM environment	9
Active and Passive monitoring of virtual networks	10
Spirent offerings powered by Intel for NFV	11
Conclusion.....	13
Glossary.....	13
Note 1	15

INTRODUCTION

Robust and scalable networks are an essential element in the critical infrastructure that supports individual organizations and national economies alike. Data from the International Telecoms Union for 2013 indicates how dependent the world has become on networking infrastructure:

- 75 percent of people in the developed world have a mobile broadband subscription.
- 2.3 billion people worldwide are online – an increase of 100 percent over five years.
- More than 100 countries have mobile cellular penetration greater than 100 percent.

The potential for further growth is huge. Wide-scale adoption of mobile devices, the growth of cloud computing and the potential for the Internet of Things are set to drive even more business, government and consumer data over private and public networks. In 2017 it is predicted that the gigabyte equivalent of all movies ever made will cross the global internet every three minutes¹.

The growth in IP traffic is both an opportunity and a challenge for network operators. Revenues from the consumption of data can be expected to grow, but it is widely acknowledged that they will be outpaced by the costs of handling the growth in traffic.

Although the implication for long-term profitability is unsettling, it is not surprising. Even in a single metropolitan area, service providers and operators deploy dozens of routers and appliances, each built using custom ASICs, running millions of lines of code and performing specialized functions. Scaling these elements to match increased demand is inevitably costly.

As a result, network operators are looking for ways to effectively scale their services, while ensuring growth in revenue is not consumed by a corresponding growth in costs. Network Functions Virtualization (NFV) has emerged as a potential solution to the problem and is attracting significant attention from industry experts.

This paper briefly considers various use cases for NFV and the benefits it offers, before looking at the challenges faced in its deployment. It then discusses at greater depth the issues surrounding validation of NFV environments and the role testing plays in ensuring successful deployments.

¹ The Zettabyte Era- Trends and Analysis. Cisco, 29 May 2013.

DRIVERS FOR VIRTUALIZING NETWORK FUNCTIONS

Offering networking services is becoming an expensive proposition for operators. Networks comprise a diverse set of functions running in specialized devices. The capital expenditure for increasing capacity raises the cost per gigabyte, and is increased still further by operational expenditure for ongoing management of services, power consumption, and the cost of employing the increasingly rare but necessary skills. Given the complexity of today's networks, even a simple upgrade to meet surging demand can take months, if not years.

However, encouraged by advances in computing and the success of virtualization techniques in IT environments, network operators are pushing for the adoption of similar techniques in their own environments as a means of addressing this scalability problem.

Network Functions Virtualization (NFV) is one such option. It leverages standard IT virtualization technology to consolidate many network equipment types onto industry-standard high-volume servers, switches and storage located in data centers, network nodes or in end-user premises.

Use cases for NFV

NFV is applicable to many network functions, including edge routing, BNG, DPI, IPSec gateways, mobility gateways, and security functions. As a result of this inherent flexibility, operators have identified a number of attractive potential use cases. These include:

- Virtualization of content distribution networks to deliver on-demand scaling of content-delivery services.
- Implementations of IP-nodes that support BNG, PE routing, and CG-NAT capabilities on standard high-volume servers.
- Virtualization of residential gateways and set-top boxes.
- Virtualization of the mobile core network to enable operators to cope with ever-increasing traffic demand.
- Virtualization of security-related capabilities, including firewalls, IDS, IPS and WAN acceleration, that need dedicated appliances on customer premises.
- Software-based DPI implementations that provide traffic analysis and reporting of line rates.

² Network Functions Virtualization – Introductory White Paper. Published October 22-24 2012 at SDN and OpenFlow World Congress, Darmstadt, Germany. http://www.tid.es/es/Documents/NFV_White_PaperV2.pdf

Benefits of NFV

In its original whitepaper² introducing NFV, members of the European Telecommunications Standards Institute (ETSI) listed a number of the benefits offered, summarized as follows:

- Reduced capital expenditure through consolidation of equipment.
- Reduced operational expenditure from lower power consumption.
- Consolidated cycle of innovation for shorter time to market.
- Availability of multi-version and multi-tenancy network appliances for greater sharing of resources.
- Targeted introduction of services based on geography or customer sets.
- Greater openness to encourage pure software firms, small players and academia to join the ecosystem.
- Lowered risk when introducing new services and new revenue streams.

CHALLENGES TO THE DEPLOYMENT OF NFV

However, replacing thousands of specialized routers and appliances networks with NFV-based servers is still a significant challenge for operators. The process will almost inevitably be a gradual one, taking into account the budgetary imperative to allow existing equipment to be fully depreciated before replacement.

Operators will also have to deal with new pricing and revenue-sharing models in an expanded ecosystem. The operations staff and technicians who are well versed in traditional network deployments, CLIs and element management systems will need to be trained to work with cloud-management systems.

Most significantly, virtualized functions, virtual environments and end-to-end networks need to be thoroughly validated prior to deployment. Failures in the field resulting from lack of thorough testing lead to lost revenues and angry customers, and will ultimately slow down the adoption of NFV.

Consequently, testing challenges associated with validating NFV environments are not to be underestimated. Below we describe the most common challenges and the means of addressing them.

THE CHALLENGES IN VALIDATING NFV ENVIRONMENTS

Functional and performance testing of network functions

For end-users, the primary concerns when it comes to network services are performance of their applications and the quality of experience. They need protocols to work in accordance with specifications, and state machines to be consistent, but are not interested in whether the BNG, routing, CDN or mobility functionalities are implemented in standard servers or purpose-built appliances.

For operators there are additional concerns regarding the control-plane and data-plane scale, and whether, for example the number of PPPoE sessions, throughput and forwarding rates, number of MPLS tunnels and routes supported are broadly similar between physical and virtual environments. Testing must ensure that the performance of virtual environments is equivalent to that of the corresponding physical environment.

Validating reliability of network services

Operators and users accustomed to 99.999 percent availability of physical network services will have the same expectations for virtual environments. It is important to ensure that node, link and service failures are detected within milliseconds and that corrective action is taken promptly without degradation of services.

In the event that virtual machines are migrated between servers, it is important to ensure that any loss of packets or services is within acceptable limits set by the relevant SLAs.

Ensuring portability of VMs and stability of NFV environments

The ability to load and run virtual functions in a variety of hypervisor and server environments must also be tested. Unlike physical environments, instantiating or deleting VMs can affect the performance of existing VMs as well as services on the server. In accordance with established policies, new VMs should be assigned the appropriate number of compute cores and storage without degrading existing services. It is also critically important to test the virtual environment, including the orchestrator and cloud-management system.

Active and passive monitoring of virtual networks

In addition to pre-deployment and turn-up testing, it is also important to monitor services and network functions on either an ongoing, passive basis or an as-needed, active basis. Monitoring virtual environments is more complex than their physical equivalents because operators need to tap into either an entire service chain or just a subset of that service chain. For active monitoring, a connection between the monitoring end-points must also be created on an on-demand basis, again without degrading the performance of other functions that are not being monitored in that environment.

TESTING METHODOLOGIES FOR VALIDATING NFV

Functional and performance testing of network functions

1) Using physical test appliances to validate NFV

The methodologies to validate NFV functions in virtual environments are similar to those employed to validate physical DUT devices. Any of the hundreds of access, routing, MPLS and mobile backhaul functions and protocols typically tested on physical DUTs will also need to be tested in NFV environments. In the example shown below in Figure 1, a physical test system is used to validate virtual BNG and PE functions implemented on a standard server. Emulated PPPoE clients, BGP routers and MPLS tunnels are instantiated in the test systems, and connected to peer sessions inside the NFV servers using control plane protocols defined by IETF RFCs. The associated user data-plane streams

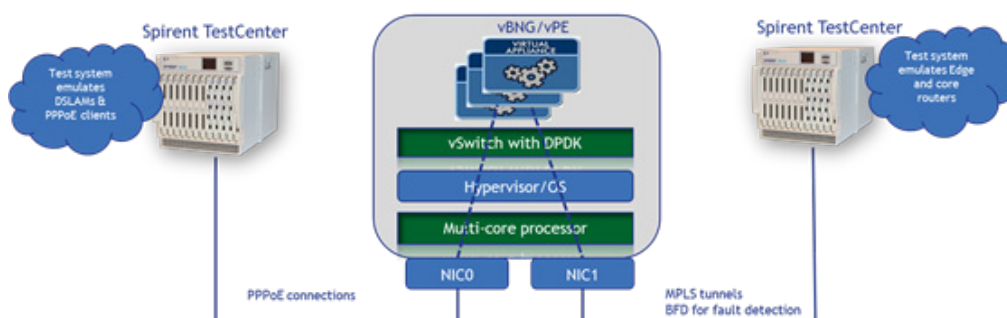


Figure 1: Validating virtual BNG functions

are also instantiated for originating and terminating traffic.

The network functions in the NFV servers are verified by measuring various data plane and control plane metrics as shown below:

Data-plane metrics:

- Latency on each of tens of thousands of data streams.
- Throughput and forwarding rate.
- Packet-delay variation and short-term average latency.
- Dropped frames and errored frames.

Control-plane metrics:

- States and state transitions of hundreds of control plane protocols (in Figure 1, these are PPPoE, RSVP and OSPF).
- Control-plane frames sent and received on each session, for example OSPF hellos.
- Control-plane error notifications.
- Validation of control-plane protocols at high scale:
 - Scaling up on one protocol (for example, tens of thousands of BGP sessions, OSPF routers or LDP sessions) and validating protocol state machines and data plane.
 - Scaling up on multiple protocols at the same time (for example, BGP sessions, OSPF/IS-IS routers and RSVP tunnels) and validating protocol state machines and data plane.
 - Scaling up on routes (for example, BGP routes and OSPF LSAs) and MPLS tunnels.

These are a representative sample of a comprehensive set of control-plane and data-plane statistics, states and error conditions that are measured for a thorough validation of NFV functions.

Performance Note: Spirent's test systems are internally powered by x86-based Intel multi-core processors. Spirent's software takes advantage of Intel architecture by running different software threads on different cores. This capability enables Spirent to scale up on multiple dimensions simultaneously. This includes, for example, achieving multi-protocol scale on BGP and LDP by running BGP on one core and LDP on another.

2) Using virtual test appliances to validate NFV

Physical test appliances are recommended for the testing virtual environments that require the highest levels of data-plane performance (line-rate) or microsecond-level timing accuracy. However, virtual test systems offer equivalent capabilities for almost all other scenarios. Test VMs, such as Spirent TestCenter Virtual or Avalanche Virtual, are ideal for organizations that need the flexibility of testing network functions on multiple, geographically disparate servers. They are also the right choice for organizations that need to test individual network functions within a service chain, or organizations that are on a limited budget.

As shown in Figure 2, the Spirent Avalanche Virtual test appliances emulate realistic user behavior and originate and terminate stateful HTTP, FTP and video traffic. The test appliances are at the end of a service chain that comprises a virtual load balancer, virtual firewall and virtual CE router.

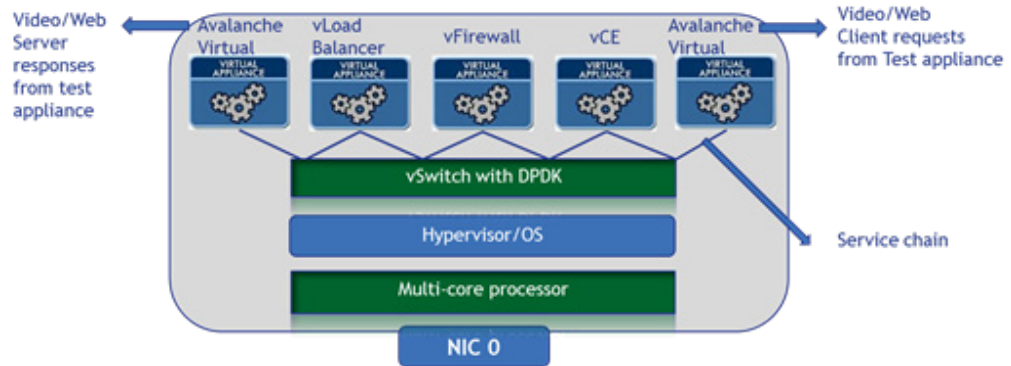


Figure 2: Validating a service chain using virtual test appliances

The following are some of the key metrics measured by the virtual test appliances:

- Sustained data transfer rate.
- Connections establishment rate, and transactions per second.
- Total number of connections.
- Round trip time and goodput.
- Denial of service handling.
- Packet loss.
- Service chain scale.
- Packet leakage across service chains.
- Time between VM instantiation and first available packet.

Performance Note: By using Intel DPDK drivers in the user space, Spirent demonstrated a five- to ten-fold improvement in the data-plane performance of Spirent Avalanche Virtual appliances.

Validating reliability of network services

In traditional networks, a number of tests are performed to establish the reliability of network services. These include validation of control plane and data plane at high scale and testing failure detection and recovery scenarios using mechanisms such as high-frequency BFD and high-frequency Ethernet OAM. These tests are equally appropriate for establishing the reliability of network services when offered on a virtual environment.

Figure 1 showed the example of a virtual BNG that was tested using a Spirent TestCenter chassis by establishing PPPoE connections and MPLS tunnels. In addition to functional and performance testing, control plane reliability is tested by establishing high frequency BFD sessions at 3.33 ms heartbeat intervals between the vBNG and the test system and declaring failures if three consecutive BFD packets are not received at 3.33 ms intervals.

In virtual environments, reliability and efficiency are also attained by migrating VMs from server to server, on an as-needed basis. Alternatively, migration can be triggered by events such as time-of-day, traffic levels or failure conditions. During, and immediately after, the migration of a VM, it is important to ensure that services are not degraded beyond acceptable levels as set in the relevant SLAs.

Figure 3 shows an example of a VM migration between two servers.

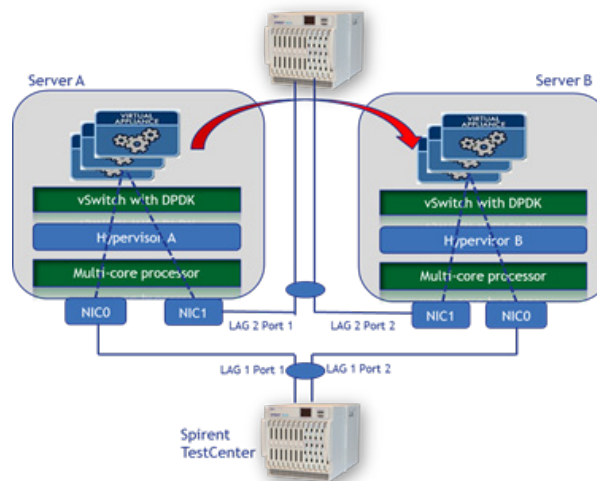


Figure 3: Testing VM migration

The following method establishes the VM migration time and any associated packet loss:

- Create link aggregation groups (LAG) between each test chassis and the NICs on the two servers. In the example shown in Figure 3, two test chassis are shown in the test setup, for clarity. In reality only two test modules are needed on a single chassis for performing the test described below.
- Initially Port 1 of both LAG 1 and LAG 2 is active, while Port 2 of LAG 1 and LAG 2 is on standby. Create a data stream originating on LAG 1 and terminating on LAG 2: this is Stream ID 1. Create a data stream originating on LAG 2 and terminating on LAG 1: this is Stream ID 2.
- Initiate the migration of the VM from Server A to Server B. At the same time, change the state of Port 1 of both LAG 1 and LAG 2 to standby, and change the state of Port 2 to active. (The exact method of synchronizing the VM migration with the change of state for the LAG port is outside the scope of this whitepaper.)
- The number of packets lost in the forward direction is calculated as the transmit time of Stream ID 1 on LAG 1 less the receive time of Stream ID 1 on LAG 2.
- The number of packets lost in the reverse direction is calculated as the transmit time of Stream ID 2 on LAG 2 less the receive time of Stream ID 2 on LAG 1.
- VM migration time is the greater of: the time of arrival of the first packet on Port 2 of LAG 2 less the time of arrival of last packet on Port 1 of LAG 2, and the time of arrival of the first packet on Port 2 of LAG 1 less the time of arrival of last packet on Port 1 of LAG 1.

Ensuring portability of VMs and stability of VM environment

The interoperability between a guest OS and host OS/hypervisor is essential if vendor lock-ins are to be avoided. For this reason, the ETSI NFV Forum has recommended that a vendor's virtual network functions and its associated guest OS should be deployable on standard NFV infrastructures that contain host OS/hypervisors from another vendor.

As a new network function is provisioned on a server through the deployment of a VM, it is important to validate the following:

- The network function performs as expected.
- The NFV infrastructure continues to provide the necessary compute, storage and network resources to existing VMs without any performance degradation, even as new VMs are instantiated.

Figure 4 shows an NFV service chain containing a WAN accelerator, traffic classifier, DPI engine and test appliances that emulate clients and servers.

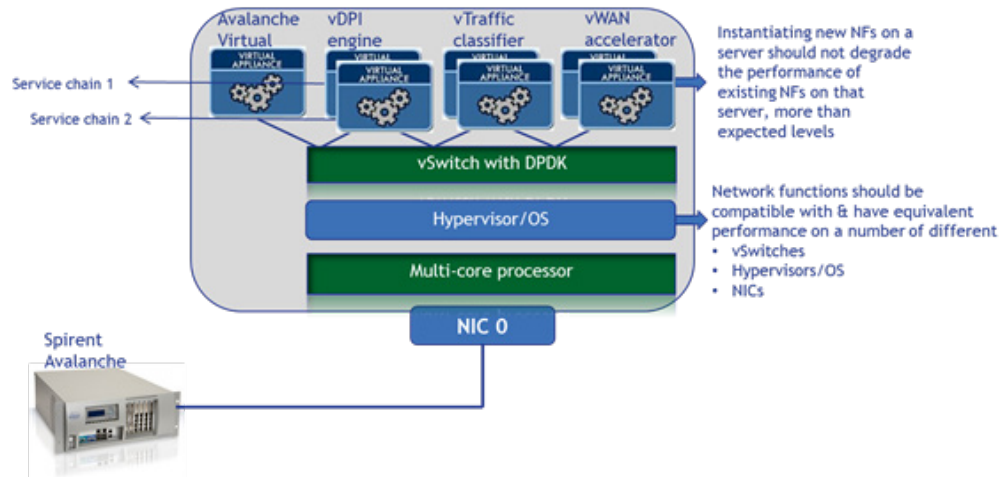


Figure 4: Portability of VMs and VM environment stability

The test appliances (in this example, Spirent Avalanche), should validate the following:

- The performance of the existing Service Chain 1; that is, the combination of WAN accelerator, traffic classifier and DPI engine.
- The instantiation of a new Service Chain 2 leads to the allocation of appropriate compute, storage and network resources, in accordance with the relevant SLAs and without degrading the performance of the existing Service Chain 1.
- Network functions are compatible with a range of different OS/ hypervisors, cloud- management systems and NICs. One way to validate this level of portability is to use a test NFV infrastructure built with multiple hypervisors (such as KVM, Hyper-V and ESX), multiple cloud management systems (such as OpenStack and CloudStack) and multiple NICs. Such a test setup is used to ensure that the virtual network functions under test are compatible with the infrastructure.

Active and passive monitoring of virtual networks

To ensure that end customers enjoy services that adhere to established SLAs, visibility of performance across the entire footprint of service-provider and datacenter networks is necessary. NFV servers are expected to contain dozens of service chains and hundreds of network functions. Proactive and reactive performance monitoring, identification of failures, isolation of any faulty network functions and corrective action are as important in virtual environments as they are in physical environments. But they are also more challenging.

The figure below depicts a service chain that comprises a load balancer, a firewall and a CE:

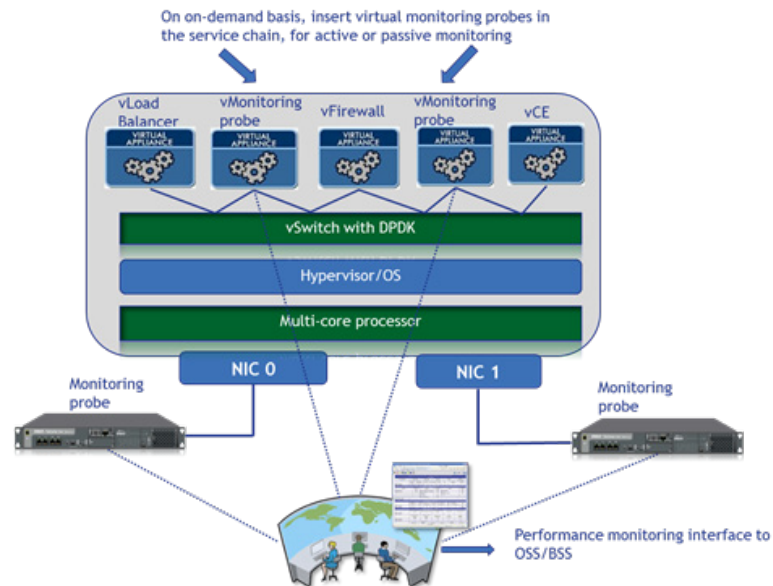







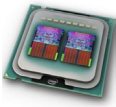




Figure 5: Performance monitoring and network visibility of VM environments

Virtual and physical monitoring probes (in this example, Spirent's STC Live 7500) continuously feed monitoring reports to a centralized operations center. The virtualized monitoring function must be 'pluggable' at any point in the service chain on an as-needed basis, and must be able to monitor, actively or passively either: a single function; a set of functions; or the entire service chain. Test access points must be provided at different stages in the service chain. If secure tunnels are used, the virtual monitoring functions must be able to proactively originate traffic and receive responses.

SPIRENT OFFERINGS POWERED BY INTEL FOR NFV

Spirent has a number of product offerings to help network operators, NEMs, and independent software vendors test and validate their NFV infrastructure and network functions. The offerings and their applications are described as follows:

SPIRENT OFFERINGS POWERED BY INTEL FOR NFV	
Product/Capability	NFV Application
Spirent TestCenter Chassis 	Functional and performance testing of vBRAS, vPE, vCE, vSTB devices and routing, access and MPLS protocols. Line rate L2-L7 traffic generation and analysis capable and timing accuracy in nanoseconds.
Spirent TestCenter Virtual 	Compatible with many OS/ hypervisors and runs on standard servers. Identical to the chassis solution in look and feel and control-plane/protocol testing.
Spirent TestCenter Appliance 	Multi-10 Gbps capacity, security and performance testing of vFirewalls, vLoad balancers, vWAN Accelerators, vIPS/IDS, vDPI Boxes and virtual web application infrastructures
Spirent Avalanche Virtual 	Compatible with many Hypervisor/OSs and runs on standard servers. Identical to the appliance solution in look and capable of high performance using DPDK drivers.
Spirent Landslide 	Testing of virtualized EPC functions. Emulates millions of data subscribers and tests mobility functions.
Spirent TestCenter Live 	Integrated device capable of performance-monitoring and diagnostics for the turn-up, monitoring and troubleshooting of Ethernet services.
Spirent iTest 	iTest is an integrated test-authoring and execution solution for rapidly developing and automating test cases.
Intel multi-core processor 	Many of Spirent's products are powered by Intel multi-core processors. Intel enables Spirent to simultaneously scale up on multiple dimensions by assigning separate cores to different functions/ protocols.
Intel DPDK 	Intel DPDK drivers vastly improve the data-plane performance of Spirent offerings such as TestCenter and Avalanche.
Intel VT-d 	By using the VT-d hardware assistance built into Intel's chipsets, Spirent products achieve higher level of performance and availability.

CONCLUSION

Network Functions Virtualization (NFV) will enable service providers to deploy many virtualized network functions on industry-standard high-performance servers. In addition to savings to both operational and capital expenditure, operators benefit from shortened time to market, the entry of more independent software vendors into the ecosystem, and easier development and deployment of new revenue-generating features.

Thorough testing of NFV is essential to minimize network outages post deployment and is critical to its fast adoption. NFV functions and infrastructure must be tested in the lab at the time of turn-up of services, as well as being constantly monitored for performance.

Spirent offers a broad portfolio of physical and virtual solutions for functional and performance testing of virtualized L2-L3 network functions, application security functions and evolved packet core functions.

For more information, please visit www.spirent.com

Glossary

ASIC – Application Specific Integrated Circuit

BFD – Bidirectional Forwarding Detection

BGP – Border Gateway Protocol

BNG – Broadband Network Gateway

BRAS – Broadband Remote Access Server

BSS – Business Support System

CDN – Content Delivery Network

CE Router – Customer Edge Router

CG-NAT/NAT – Carrier Grade Network Address Translation

CLI – Command-line Interface. Also known as command-line user interface, console user interface, and character user interface

DPDK – Data Plane Development Kit

DPI – Deep Packet Inspection

DSLAM – Digital Subscriber Line Access Multiplexer

DUT – Device Under Test

Ecosystem – Equipment and software vendors, integration services vendors, network operators, service providers, academia, etc.

EPC – Evolved Packet Core

FTP – File Transfer Protocol

IETF – Internet Engineering Task Force

IDS – Intrusion Detection System

Industry-Standard High-Volume Server – a server built on standardized IT components (such as x86 architecture) and sold in the millions

IP – Internet Protocol

IPS – Intrusion Prevention System

IPSec – IP Security

IS-IS – Intermediate System to Intermediate System (a routing protocol designed to move information within a computer network)

LDP – Label Distribution Protocol

LSA – Link-state Advertisement (a communication mechanism of the OSPF routing protocol for IP)

M2M – Machine-to-Machine communications

MPLS – Multiprotocol Label Switching

NEM – Network Equipment Manufacturer

NIC – Network Interface Controller

OAM – Operations Administration and Maintenance

OS – Operating System

OSFP – Open Shortest Path First

OSS – Operations Support System

RFC – Request for Comments (a publication of the Internet Engineering Task Force and the Internet Society)

RSVP – Resource Reservation Protocol

PE Router – Provider Edge Router

PPPoE – Point-to-Point Protocol over Ethernet

SLA – Service Level Agreement

SBC – Session Border Controller

STB – Set-top Box

STCv – Spirent TestCenter Virtual

VM – Virtual Machine

vSwitch – any Ethernet switch implemented in software alongside or inside a hypervisor. There are proprietary and open implementations of vswitch.

WAN – Wide Area Network

NOTE 1:

The methods of calculation of packets lost and VM migration time were expressed as a mixture of mathematical formulae and written text. Ideally, the paper should stick to one or another format. The written version has now been included in the body of the paper. The mathematical formula can be substituted instead and would properly be expressed as follows:

- The number of packets lost in the forward direction is given by:

$$N_{TX}[SID\ 1, LAG\ 1] - N_{RX}[SID\ 1, LAG\ 2]$$

where $N_{TX}[SID\ 1, LAG\ 1]$ is the number of transmitted packages from Stream ID 1 on LAG 1, and $N_{RX}[SID\ 1, LAG\ 2]$ is the number of received packages from Stream ID 1 on LAG 2.

- The number of packets lost in the reverse direction is given by:

$$N_{TX}[SID\ 2, LAG\ 2] - N_{RX}[SID\ 2, LAG\ 1]$$

where $N_{TX}[SID\ 2, LAG\ 2]$ is the number of transmitted packages from Stream ID 2 on LAG 2, and $N_{RX}[SID\ 2, LAG\ 1]$ is the number of received packages from Stream ID 2 on LAG 1.

- VM migration time is given by:

$$\text{Max}(T_i[\text{port } 2, LAG\ 2] - T_l[\text{port } 1, LAG\ 2], T_i[\text{port } 2, LAG\ 1] - T_l[\text{port } 1, LAG\ 1])$$

where $T_i[\text{port } x, LAG\ y]$ is the time of arrival of the first package on port x of LAG y, and $T_l[\text{port } x, LAG\ y]$ is the time of arrival of the last package on port x of LAG y.



At Spirent Communications we work behind the scenes to help the world communicate and collaborate faster, better and more often. The world's leading communications companies rely on Spirent to help design, develop and deliver world-class network devices and services.

Spirent's lab test solutions are used to evaluate performance of the latest technologies. As new communication services and applications are introduced in the market, Spirent provides tools for service management and field test to improve troubleshooting and quality. Spirent also enables enterprises, institutions and government agencies to secure and manage their networks.

To learn more how Spirent can help with your testing requirements, please visit: <http://www.spirent.com/>

