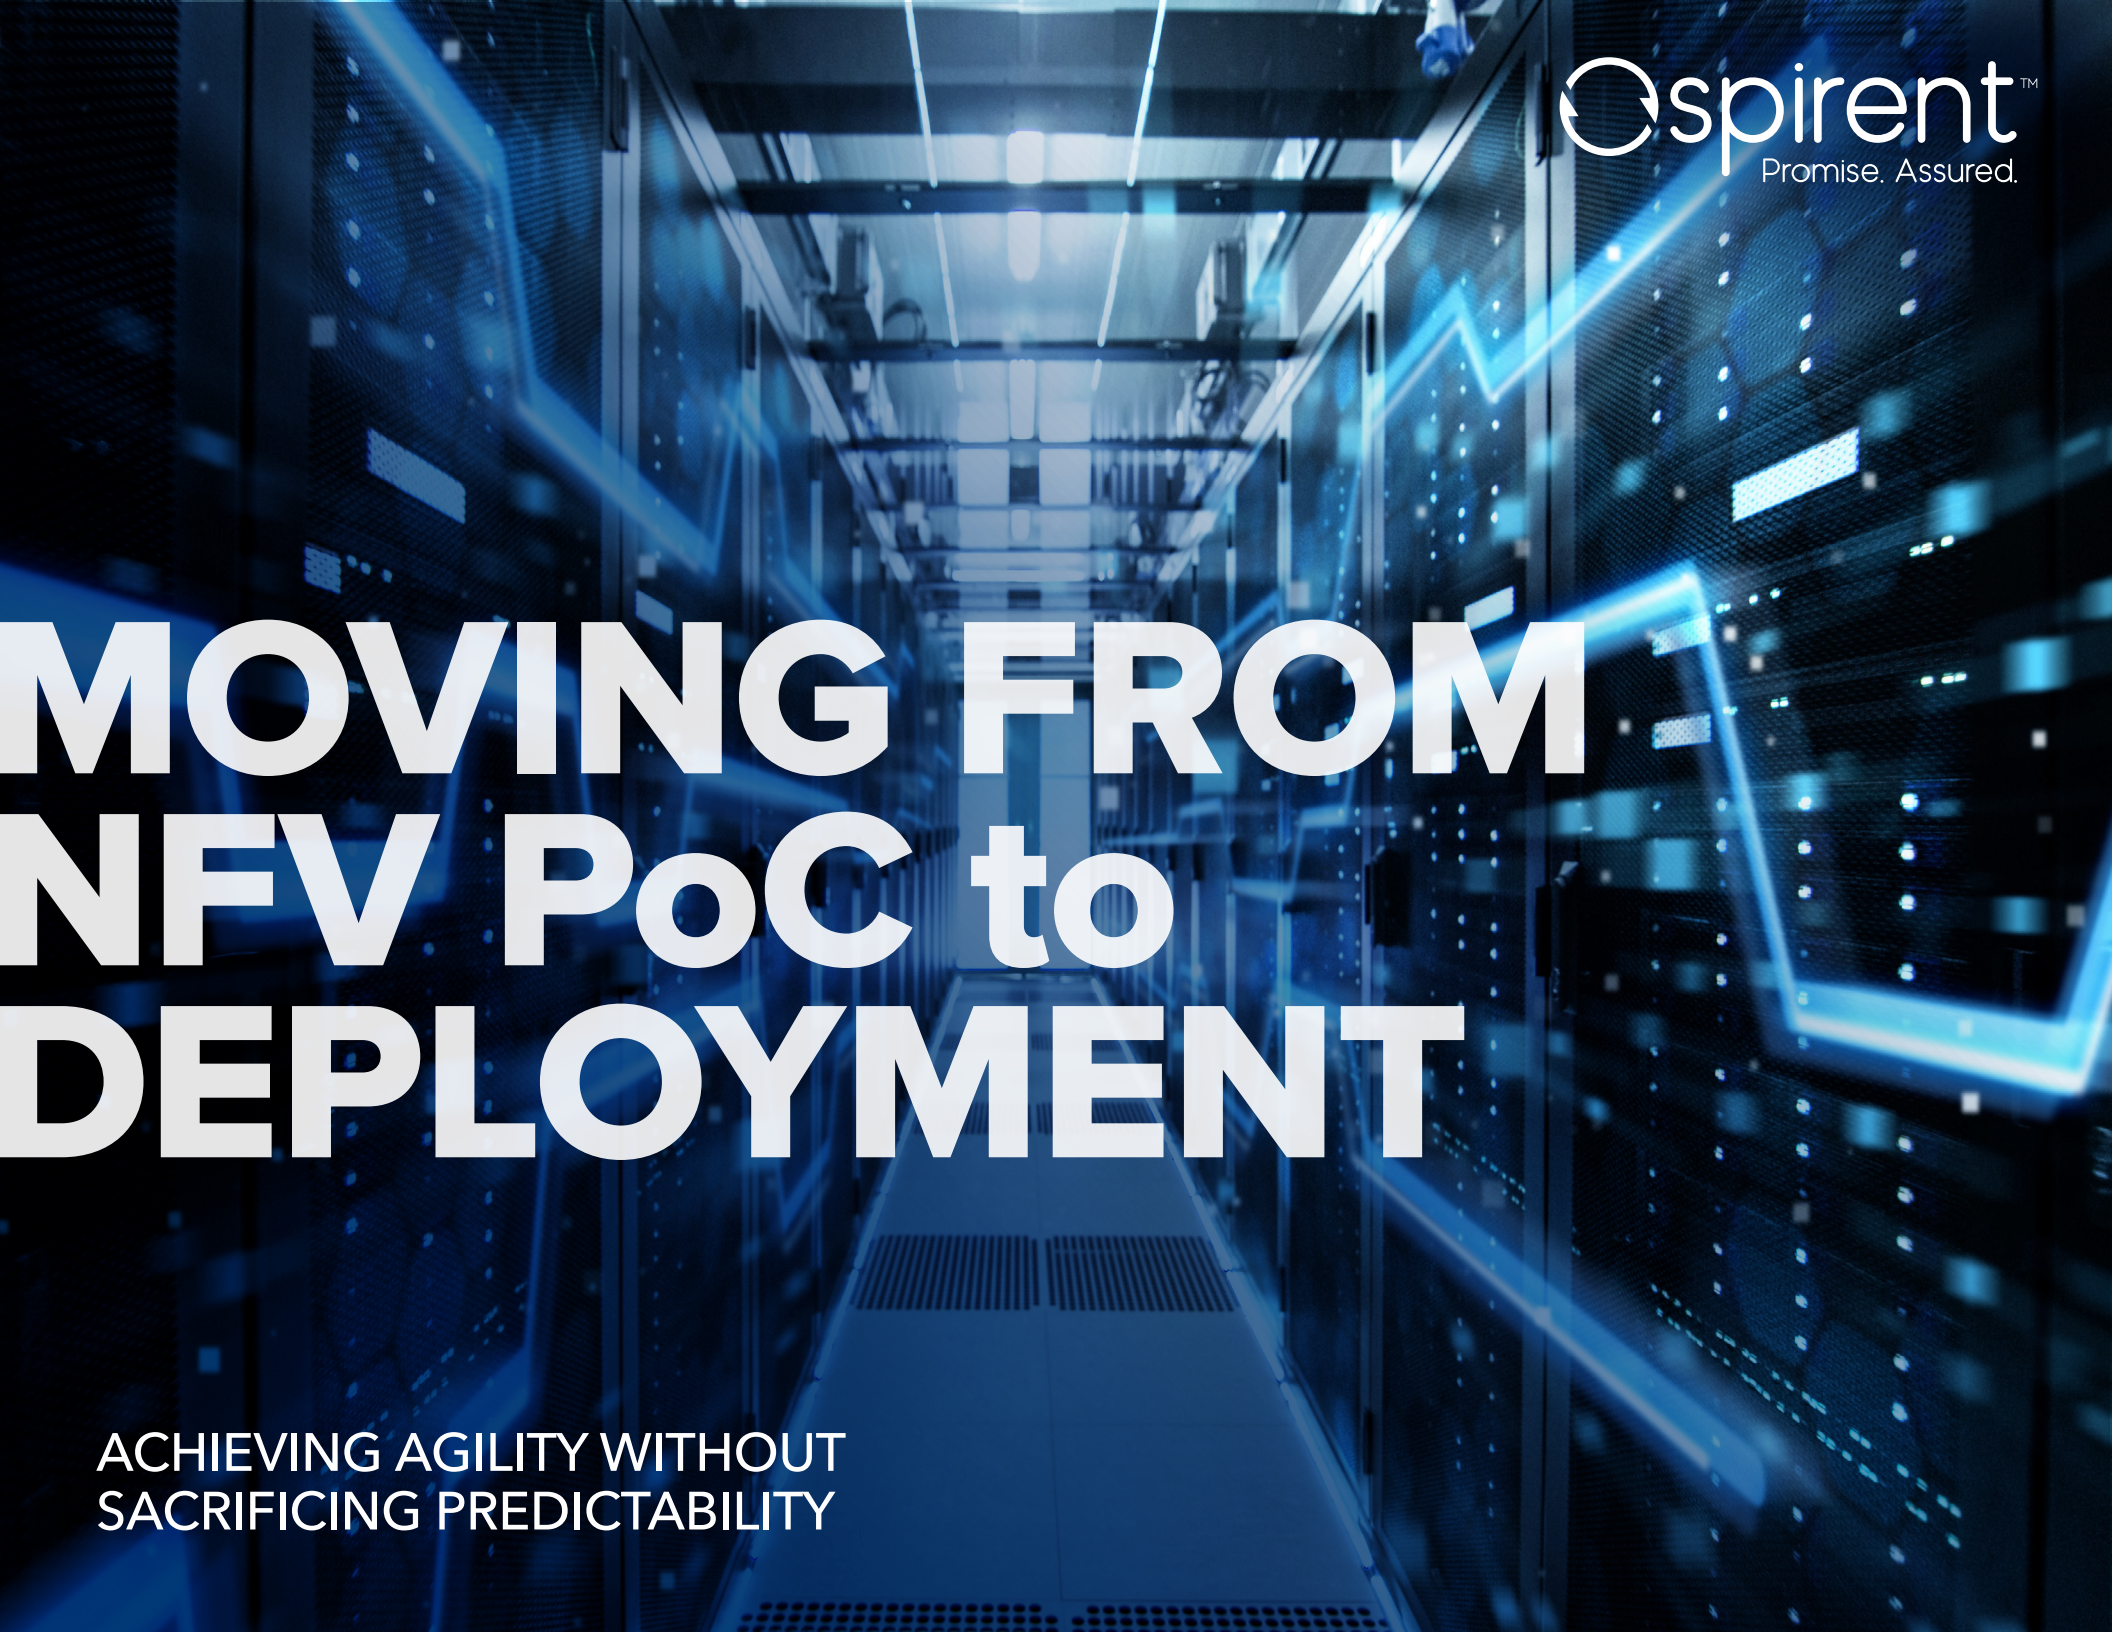
Spirent
Promise. Assured.
MOVING FROM NFV PoC to DEPLOYMENT
ACHIEVING AGILITY WITHOUT SACRIFICING PREDICTABILITY

The Network is in the Way

The current market drivers for network operators are embodied in the expectations of the always-on, always-connected generation. Today's subscribers expect unbounded mobility, seamless availability, no-wait response, and unlimited options for applications and content.

Fixed broadband connection rates are expected to reach 42 mpbs in 2018, a 2.6-fold increase from 2013 rates. By 2019 there will be 5.2 billion mobile users globally and 11.4 billion mobile devices[1].

In such an environment, service providers must be nimble to maintain profitability. They need to quickly roll out services to keep up with market demand while driving down OpEx by automating provisioning, monitoring, and maintenance.

The good news is that the cost per subscriber is falling. The bad news is that the annual revenue per user is falling faster. According to analysts, if service providers don't find a way to change the trajectory of one or both of the curves, profits will be driven out of the industry by 2017

The legacy network limits the ability of the service provider to react quickly to market needs with new service offerings. Their networks contain thousands of purpose-built proprietary hardware appliances, such as routers, session border controllers, broadband remote access servers, firewalls, deep packet inspection appliances, WAN accelerators, radio access network nodes, and the like. In most cases, launching a new service requires a significant capital investment in more proprietary hardware appliances and a commensurate investment in operations to provision and maintain the new service.

When the investment in the existing network has not been amortized, it can be difficult to make the business case for committing to an additional investment in new infrastructure. Such analysis paralysis can blunt a service provider's responsiveness to the market, slowing the development and deployment of new services and features.

"The good news is that the cost per subscriber is falling. The bad news is that the annual revenue per user is falling faster."

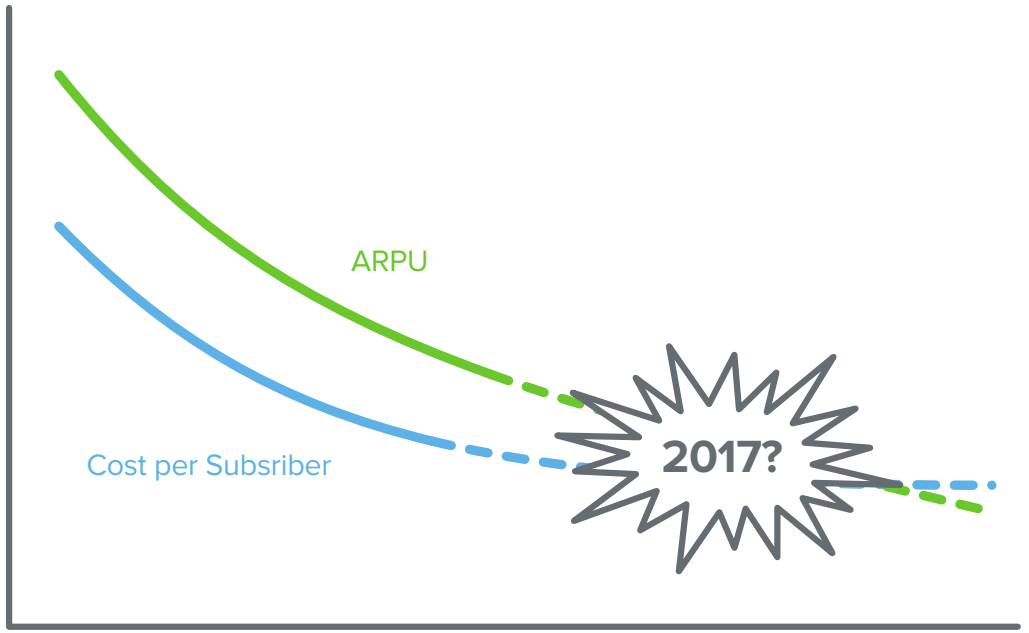

1 Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2014-2019. February 3, 2015.

The Promise of NFV and SDN

To move at the speed of the market, service providers need standardized appliances that can serve any required network function (NF) by remotely and/or automatically loading a software image onto the appliance. And to maintain profitability, the appliance should be able to run multiple instances of the required network function in virtual machines (VMs).

This concept is known as network functions virtualization (NFV). A mature technology in widespread use in the IT world, virtualization is the innovation with the potential to change the profitability curve by taking networking beyond the legacy black-box infrastructure built on inflexible, proprietary, purpose-built hardware.

By contrast, virtualization lives in the highly-adaptable and nimble world of software running on powerful CPUs with multi-processor cores. In the world of networking, the goal of virtualization is to consolidate a range of legacy network elements onto industry-standard, high-volume servers, switches, and storage devices located in data centers, network nodes, and end-user premises. These NFs can be instantiated in various locations in the network as required without the need to install new equipment, thus providing the elastic scale required to change the profitability curve.

NFV offers the flexibility service providers need to create the agile development and deployment environment that is necessary to meet market expectations. In addition, the hardware required to support the virtualized network is a comparatively inexpensive, standardized, white-box platform that opens the door to increased interoperability, taking the possibility of multi-vendor solutions to a new level.

In late 2012, twelve operators joined to publish the seminal white paper on the need for NFV and formed the ETSI NFV Industry Specification Group (ISG). In a little over two years, the membership has expanded to more than 37 operators and 230 individual companies illustrating the challenges the industry faces with traditional networks.

Benefits of NFV/SDN

- Improved agility (network and services on demand)
- Self-provisioning of network policies (automation)
- Reduced time to deliver
- Innovative products that exploit service chaining
- Elastic, demand-based scaling
- Global network visibility
- Optimized traffic steering

NFV starts with commodity hardware—a standard, high-volume platform (switch, server, or storage)—and a virtualization layer (hypervisor or container). This combination of platform and virtualization layer is referred to as the NFV Infrastructure (NFVI) and it is managed by a virtual infrastructure manager (VIM) such as Openstack. Virtual network functions (VNFs) run on top of this NFVI and are managed by a VNF Manager, which is responsible for VNF lifecycle management. Applications then access the VNFs transparently. The NFV Orchestrator (NFVO) is responsible for on-boarding of network services and VNFs, service lifecycle management, and other global resource management tasks.

Virtualization is the current best bet for revolutionizing the agility, performance, and cost of practically every function in the end-to-end network. Starting at the edge, whether it is a residential subscriber, an enterprise customer, or a mobile base station, to the evolved packet core where it can enforce classification and prioritization policies, forward traffic, and perform recovery, to the data center where it can be deployed in a wide range of roles, including security, monitoring, and load balancing.

Software-defined networking (SDN) shares many of the same objectives as NFV, primarily to liberate the telco and cloud operators from the bonds of expensive and inflexible proprietary hardware. SDN focuses on separating data and control plane functions and abstracting the underlying infrastructure for applications and network services.

According to ONF's "SDN Architecture Overview," enterprises and carriers gain unprecedented programmability, automation, and network control, enabling them to build highly scalable and flexible networks that adapt to changing business needs and network conditions. As such, the two technologies are complimentary. SDN can take advantage of NFV to virtualize functions such as monitoring, management, traffic analysis, and load balancing.

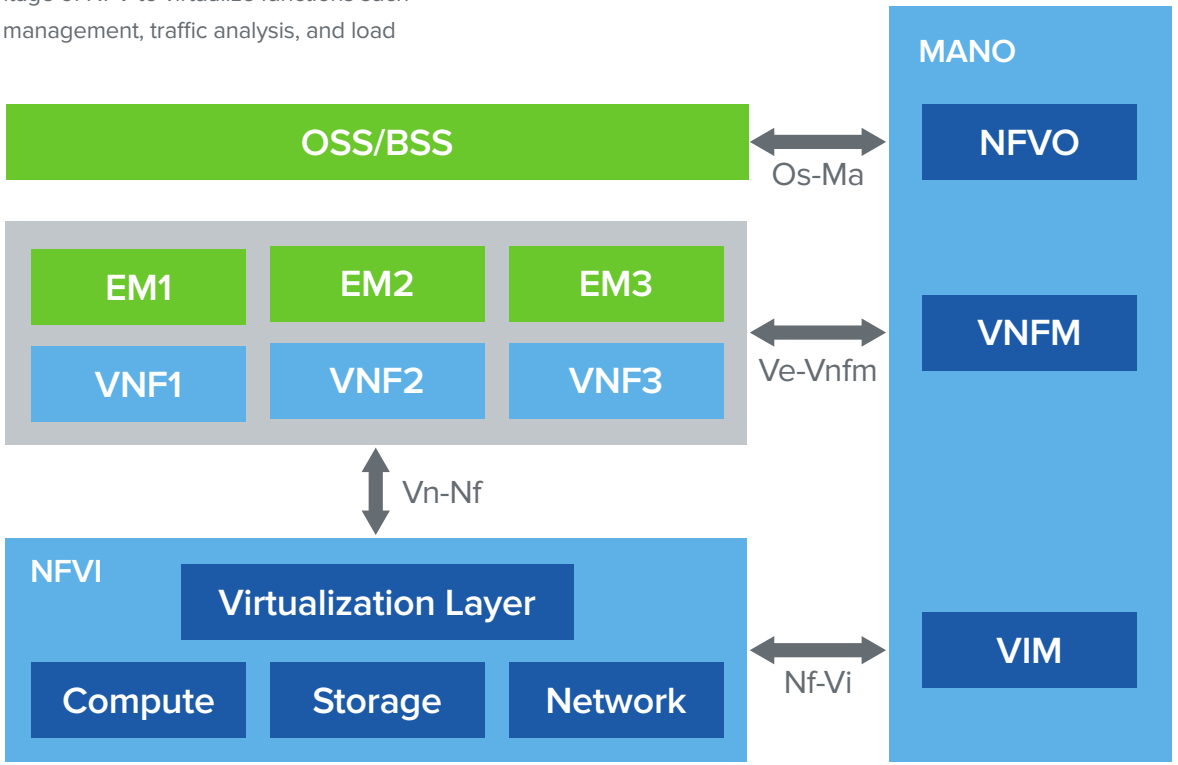

Figure 1: Network Functions Virtualization architecture

Not surprisingly, in a recent survey of service providers, 97% of the respondents planned to deploy SDN, and 93% planned to deploy NFV. The top drivers for adopting SDN were to support cloud services and business access. The top drivers for adopting NFV were service scalability and the profitability of a software-based solution running on commercial off-the-shelf (COTS) servers[2].

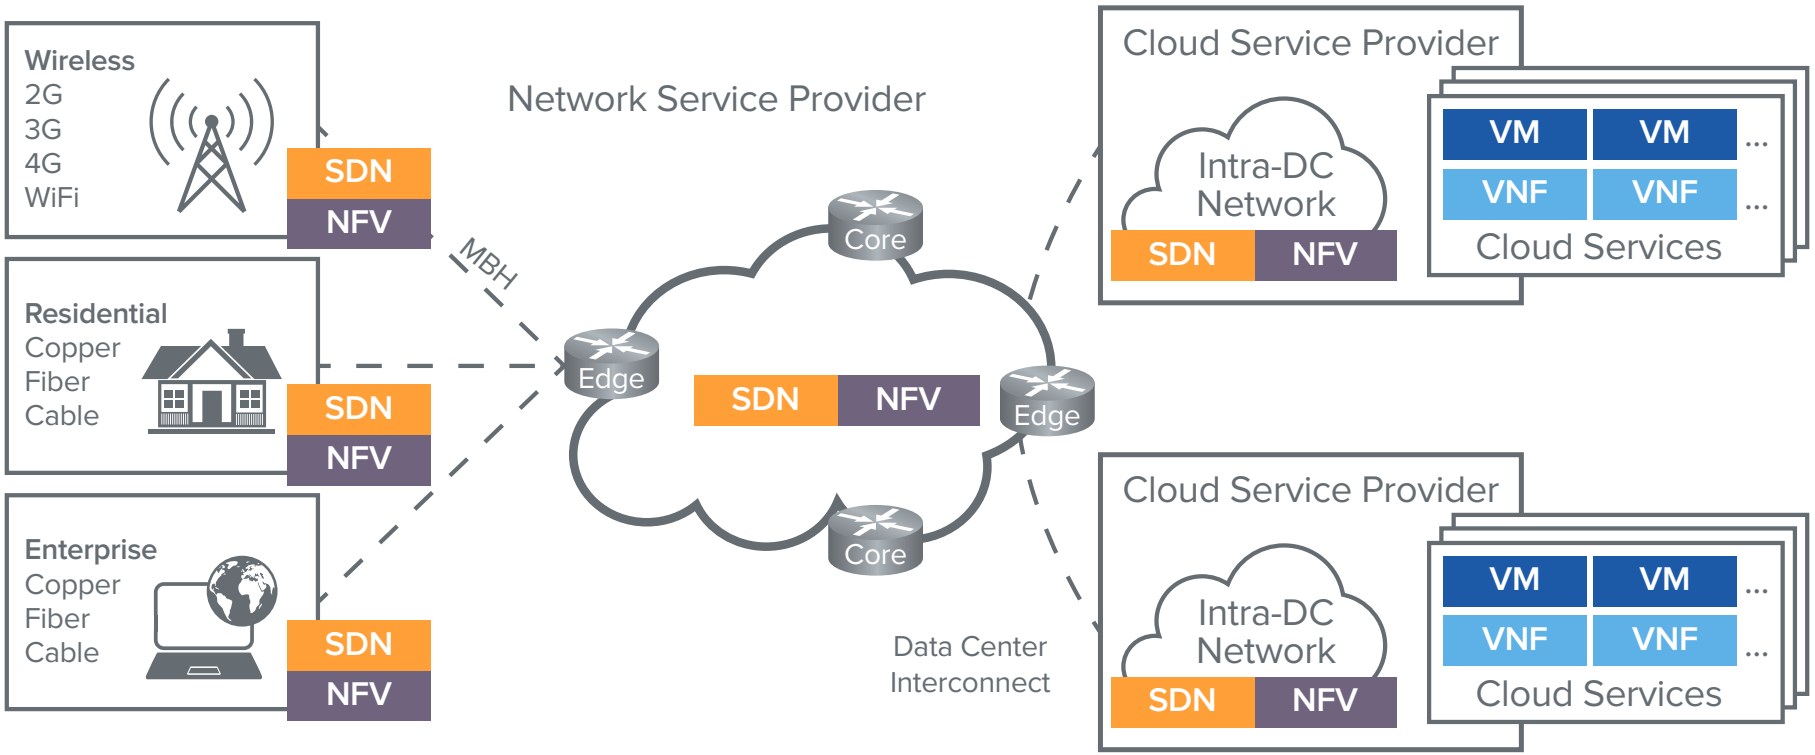

Figure 2: Network and Cloud Service Providers are deeply engaged in NFV and SDN evaluations

2 "SDN and NFV Strategies: Global Service Provider Survey," Infonetics Research, 2014

The Challenge of Moving from PoC to Deployment

Given the potential gains in agility, affordability, and operational simplicity, it should come as no surprise that service providers are actively investigating NFV and SDN. Several have completed proof-of-concept (PoC) trials with a view to deploying NFV in their production network within two to five years.

Some recent trials.

ETSI NFV PoCs:

12 operator led PoCs have been completed successfully and 21 PoCs are in progress. Spirent participated in PoC #9, "VNF Router Performance with DDoS functionality" with Brocade, Intel, AT&T and Telefonica. The details of the PoCs , the use cases and the participants are described in the following link.

http://nfvwiki.etsi.org/index.php?title=On-going_PoCs

November 2014:

Telefónica completes successful multi-vendor SDN proof-of-concept trial (ADVA Optical Networking, Ciena, Huawei, and Infinera)

http://www.ciena.com/connect/blog/Ciena-works-with-Telefonica-on-Successful-SDN-Proof-of-Concept-Trial.html

October 2014:

DoCoMo completes multi-vendor proof-of-concept NFV trials (Alcatel-Lucent, Cisco, Ericsson, Huawei, NEC, and Nokia Networks)

http://www.fiercewireless.com/tech/story/docomo-completes-proof-concept-nfv-trials-6-vendors/2014-10-15

June 2014: TM Forum Live

- Data-driven network performance optimization for NFV and SON (Mycom, TEOCO, and Wipro)
- Dynamic, data-driven management and operations (EnterpriseWEb, Huawi, and Qosmos)
- SDN and NFV while enforcing SLAs over WANs (AT&T, Telecom Italia, Netronome,
- Intel, ServiceMesh, PLUMgrid, Cisco Systems)
- Service bundling in a B2B2X marketplace (Cisco Systems, DGIT, and Liberated Cloud)

May 2014: Network Virtualization and SDN World

- End-to-End vEPC Orchestration in a Multi-vendor Open NFVI Environment (Intel, Cyan, Red Hat, Dell, and Connectem)
- Multi-vendor Distributed NFV (Cyan, RAD, Fortinet, and Certes)
- Unified SDN and Cloud Services (Cyan, Accedian, Arista, Boundary, Canonical, and RYU)

http://sdnworldevent.com/proof-of-concept-demos/

While operators have been engaged in PoC trials since early 2014, for several reasons there has been little success in moving past trials to real-world deployments.

Why DPDK and SR-IOV Matter

Packet processing performance has significantly improved in Intel processor based platforms due to software advances such as DPDK. DPDK is a set of optimized software libraries and drivers that enable high performance data plane performance by eliminating kernel and hypervisor bottlenecks.

SR-IOV enables network traffic to bypass the vSwitch, thereby eliminating the performance bottlenecks introduced by the hypervisor and the vSwitch.

Complexity. As illustrated in Figure 1: Network Functions Virtualization, the NFV architecture defined in GS NFV 001[3] by ETSI involves a shared NFV infrastructure comprising hypervisors, vSwitches, and COTS hardware, orchestrated by the NFV management and orchestration (MANO) through complex interactions between VNFs, the NFVI, the VIM, and the Orchestrator. The resulting new points of failure can affect the quality of experience (QoE), reliability, and availability of network services.

Hardware vs software. For all the advantages of affordability and market agility, NFV and SDN operate in a shared environment running on a COTS platform that will require acceleration techniques such as DPDK and SR-IOV to approach the level of performance, predictability, and scalability of the FPGA- and ASIC-based proprietary hardware appliances of the legacy network.

Multiple vendors. Service providers look to NFV and SDN to escape the dreaded single-vendor lock-in of the legacy network. But the openness of a multi-vendor environment increases the complexity and cost of interoperability testing and vendor integration.

Multi-tenancy. One way NFV helps service providers increase efficiency and lower costs is to allow disparate VNFs or service chains for multiple tenants in much the same way as virtualization is used in the IT world. However, in a shared environment, the data and traffic for each tenant must be protected from interference from other tenants, whether intentional (malicious) or otherwise. Achieving this goal increases complexity and can affect performance.

Dynamic, policy-driven provisioning. The flexibility that NFV and SDN offer for service chaining, demand-based auto-scaling, and dynamic switch/router programming comes at the price of complexity and can pose troubleshooting challenges. For example, if a function in a service chain is auto-scaled, it could be instantiated on a different server, causing user traffic entering Server A to be redirected to Server B and then returned to Server A to traverse the rest of the service chain.

3 ETSI GS NFV 002: “Network Functions Virtualisation (NFV); Architectural Framework”.

As a result of such considerations, despite public and private trials, service providers have reservations regarding the gaps between expectations and implementations. In a recent survey, the top three concerns were end-to-end provisioning across domains, the business case for deployment, and the immaturity of current solutions[4]. Other issues included security and strategies for moving from PoC to full deployment.

NFV has to deliver simultaneously on multiple levels. The trials have demonstrated that it achieves the most foundational requirement, which is functionality. But before service providers will move from evaluation to deployment, PoC trials will have to demonstrate that NFV can support carrier-grade performance in a production network.

Concerns about Deploying NFV

Maturity	• As trials demonstrate incremental progress, how long will it take before NFV implementations are ready for deployment? • Will the maturity of NFV solutions track with the requirements of my services roadmap?
Operations	• What is the management burden of an NFV deployment? • Will NFV actually reduce OpEx, and if so, by how much?
Performance	• How does the network performance of NFV compare with legacy hardware-based solutions? • Will the NFV implementation deliver the real-time performance required for time-sensitive traffic such as voice and video?
Scalability	• Will NFVI performance scale to handle the scale of the internet and beyond? • Can NFV meet our current traffic load and anticipated growth margins?
Reliability	• Does NFV offer carrier-grade reliability, including high availability and robust disaster recovery • If I adopt NFV, will my network meet my service level agreement (SLA) commitments even under peak load?
Security	• Does the implementation have adequate security built in?
Profitability	• To take advantage of the economies of scale promised by NFV, I need a large deployment, but while determining CapEx savings is fairly straightforward, how do I estimate the scalability of OpEx? Is it linear? • Will NFV deliver the anticipated CapEx and OpEx savings to maintain profitability in the face of the expectations and pace of the twenty-first century?

4 Dr. Jim Metzler, The 2015 Guide to SDN and NFV, 2015

The issue at hand is the lack of predictability of how the NFV-enabled network will respond under real-time, real-world conditions. The virtualized network might allow the service provider to respond to the market with agility, but that is of little use if there is no way to ascertain that the corresponding service will deliver the performance, availability, scalability, and security required to assure profitability.

Agility without predictability is just chaos.

Testing is the key to providing predictability. Virtualization has the potential to transform the legacy network to accommodate the demands of the twenty-first century, but testing methodologies based on physical test endpoints are not adequate to provide actionable information. It comes down to the location of test system endpoints. To understand the depth of the problem, consider the topology of test methodologies using physical test endpoints.

A legacy network consists of functional silos implemented on purpose-built proprietary hardware appliances such as a router, switch, firewall, intrusion prevention system, or other device targeted to a specific network function. As such, the device under test (DUT) or system under test (SUT) functions as a black box. What happens inside the box is opaque, but that is not the concern of the test or the service provider. All that matters is that the SUT properly and efficiently responds to the stimuli presented to it in the form of data and control plane traffic, commands, and network conditions.

In this environment, the physical test platform mirrors the legacy network elements—a dedicated physical device. It brackets the SUT by replicating the functions of all the other components of the end-to-end network, presenting user and control plane traffic at a scale appropriate to the test case to assess the performance, availability, scalability, and security of the SUT.

"Agility without predictability is just chaos."
— Neil Holmquist, Sr. Director, Product Marketing & Mmgt Cloud & IP, Spirent Communications

Test Methodologies Using Physical Test Endpoints

To illustrate the conventional test topology and methodology, consider these typical test cases used to validate physical devices—data plane validation, control plane compliance validation, and management plane validation.

Data Plane Validation

A few decades ago network designers were in a similar predicament to those considering NFV now. How do I objectively evaluate the performance of a device? The IETF developed a set of benchmarking methods to validate data plane performance, including RFC 2544, RFC 2889, RFC 3918, and RFC 5180.

In this test, one test port sends traffic of varying frame sizes and frame rates to the SUT, which processes it and forwards it to another test port, which collects key performance indicators (KPIs), including throughput (bits per second), latency (milliseconds), and frame loss (frames per second).

Control Plane Compliance and Scalability Testing

The control plane is more complex than the data plane. In this case, a physical test system emulates network nodes running control plane protocols, establishes sessions, exchanges routes, and generates traffic flows that simulate real subscriber behavior. The SUT processes control messages and forwards traffic to the terminating test port. The test ports validate the ability of the SUT to support control plane sessions at high scale and verify that the control plane traffic received from the SUT is compliant with protocol standards.

Management Plane Validation

Near-instantaneous fault detection, recovery, and convergence are essential for a carrier-grade network. In this case, the test system emulates two routes to the same destination and generates traffic. The test system then causes a failure on the primary route or node and measures the time it takes the network to recover and direct traffic to the back route or node.

Figure 3: Data plane test topology

NFV/SDN Test Methodologies

As we consider the examples of conventional test methodologies, a few things become apparent. First, to generate the control and data plane traffic and capture the results required to validate a solution, the test system brackets the DUT or SUT, serving as the endpoints for the solution or function being evaluated. Second, in the case of conventional devices and networks, the endpoints of a test topology are the ingress and egress ports of the SUT. The connective tissue between the test system and the SUT is always a cable (or over-the-air RF signal in the case UE/enodeB testing).

But when we look at Figure 1: Network Functions Virtualization architecture, we see that the physical DUTs shown in the conventional testing examples are instantiated and executed as VNFs in the virtualized network. In addition, in the virtualized environment we have new components, such as the NFVI and the NFV MANO, and new interfaces, such as between the VNFs, NFVI and the NFV MANO components.

When do you Need a Virtualized Test Solution?

In the virtualized network, standards development organizations such as the ETSI NFV ISG are leading the way, not only in defining the NFV architecture and requirements, but also methodologies on what to test and how to test.

The new components of the NFV architecture introduce points of failure that don't exist in the legacy network and therefore must be tested for functionality, performance, availability, scalability, and security. But we can't run a cable from a test point to a specific VNF to measure its performance. Instead, one or more test endpoints in a test will be a VNF.

Virtualized test solutions (test VNFs) are software-only offerings that run on commercial off-the-shelf x86-based servers. Test VNFs execute on a hypervisor or container-based NFVI and are used to validate other VNFs, the NFVI components, the NFV MANO, and E2E network services. Test VNFs, like their physical counterparts, bracket the VNF or NFVI under test, originate user and control plane traffic, and verify whether the received traffic is compliant with protocol standards and expected service level agreements (SLAs). A combination of virtual and physical test solutions is required to validate NFV and SDN environments as in the DCI WAN example in Figure 4.

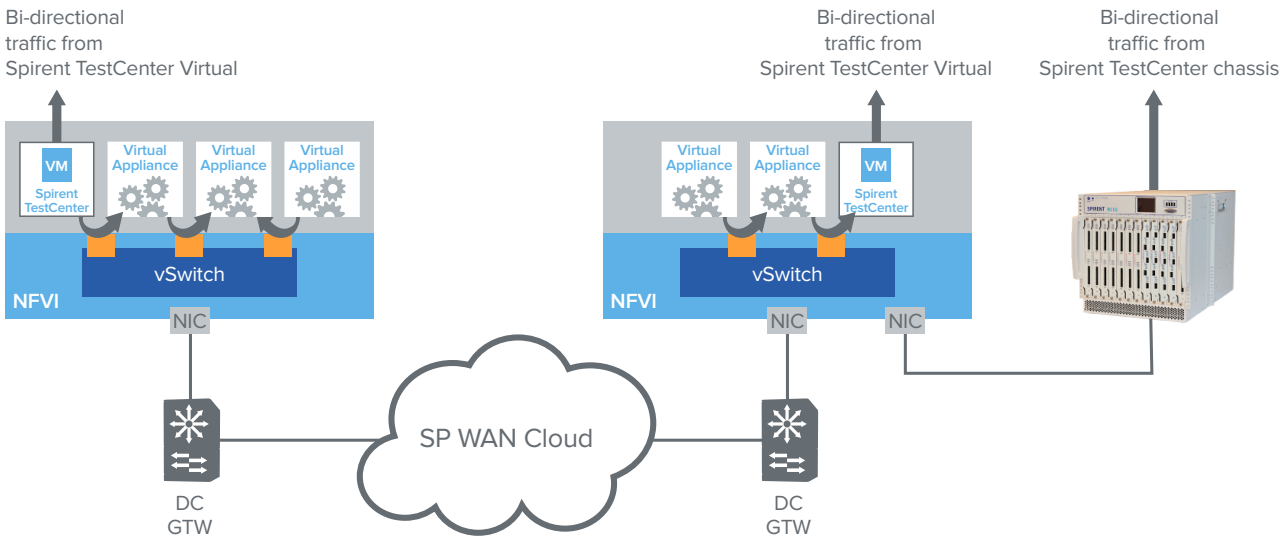

Figure 4: Physical and virtual test solutions used for validating DCI across SP WAN

If we re-frame Figure 1: Network Functions Virtualization to accommodate virtualized test endpoints, we get the arrangement shown in Figure 5: The virtualized test bed.

The NFV environment needs to be tested on several levels:

- Validate the NFV infrastructure
- Validate VNFs for functionality and scale
- Validate the SDN controller for functionality and southbound protocols
- Validate service chaining, auto-scaling, and policy-driven use cases

Some test cases will involve virtual test functions that all reside within the NFVI. Others, as shown in Figure 4, will involve both virtual and physical test functions. The exact choice of physical or virtual test points depends on the nature of the service being virtualized, as described in the table.

Choosing Between Physical and Virtual Test Solutions

Physical	Virtual
• Validate extreme data plane scalability of high-speed Ethernet interfaces (10 GB and above) • Reliably measure timing (latency, delay variation, synchronization) at microsecond accuracy at all data rates • Validate VNFs and physical DUTs for high performance and scalability	• Perform functional testing of VNFs • Validate NFV infrastructures at high scale • Perform developer testing early in the software testing cycle • Perform testing at short notice by quickly downloading and executing a test VNF versus the procurement cycle required for a physical test solution • Support multi-user, multi-site testing without shipping physical test devices to multiple locations • Test in an automated, orchestrated, multi-tenant environment • Perform complex service chaining and auto-scaling validation in an NFV environment where a physical test solution can't isolate and debug individual segments of a service function

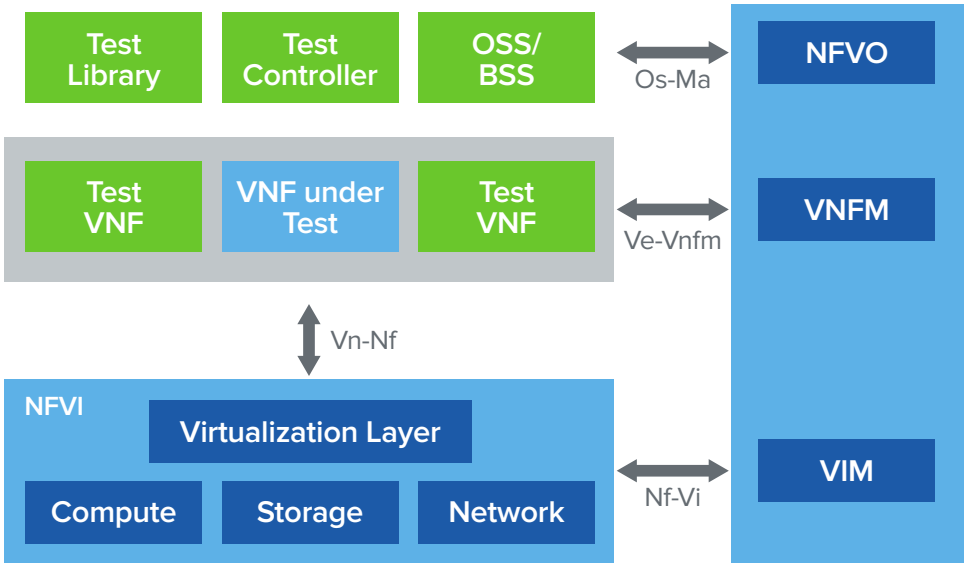

Figure 5: The virtualized test bed

Each domain of the end-to-end network, such as data centers, mobility gateways, or the access/edge networks, faces its unique set of challenges, but in the virtualized network there are use cases that are common across all segments of the industry with common testing challenges. Mobile, cloud, and access/edge telco operators alike are concerned about NFV data and control plane performance, NFVI validation, and network services testing involving service chaining, multi-tenancy, and auto-scaling.

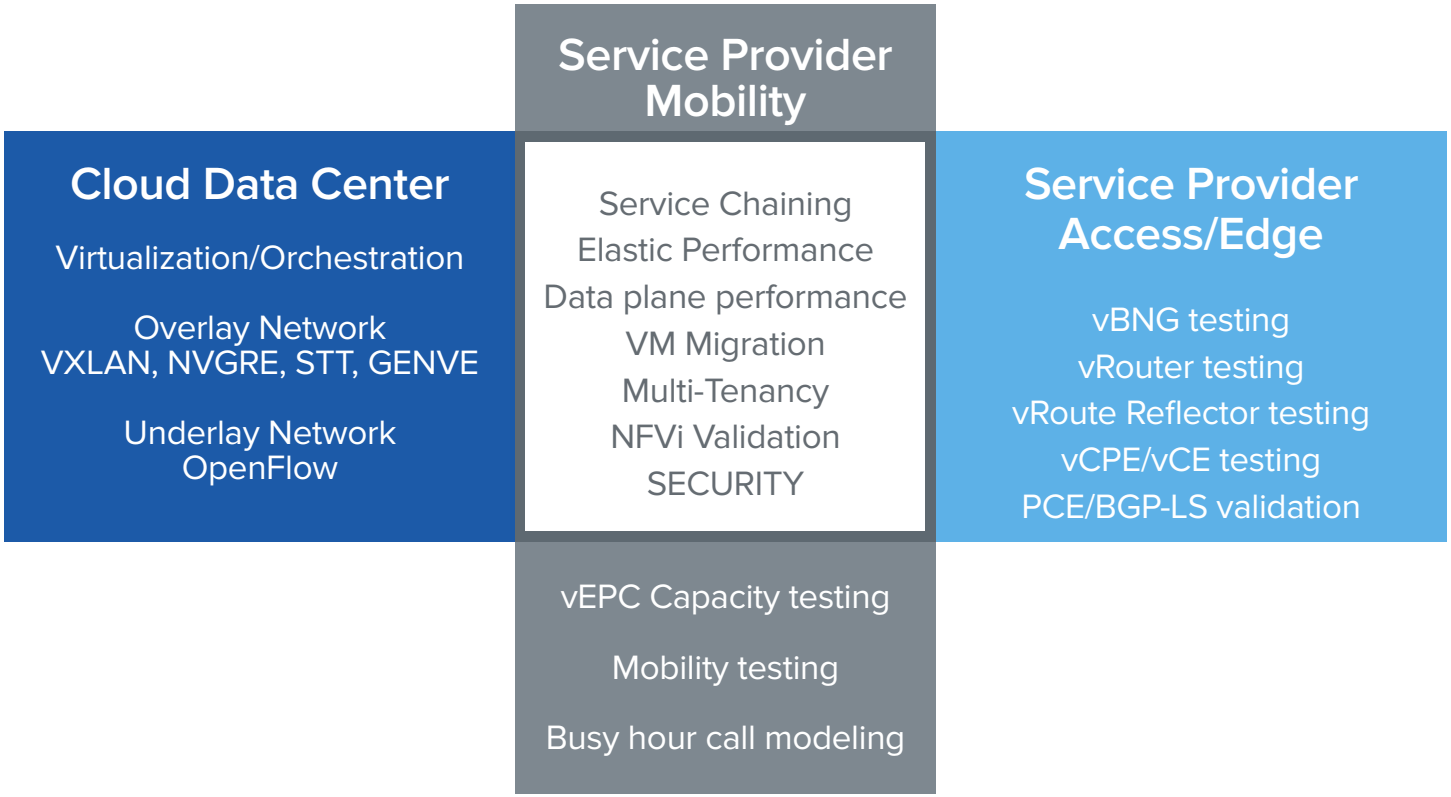

Figure 6: SDN/NFV testing and deployment challenges

What We're Testing

Before we discuss the test methodologies, we should clarify what we are, and are not, testing. NFV allows a service provider to virtualize a variety of network functions, such as routing, load balancing, or WAN acceleration. For decades, vendors and operators have used physical test appliances to validate the protocol state machines and messaging and that is not the focus of this document.

Our focus is to validate the performance of VNFs, the NFVI, and services, and to alleviate reliability concerns resulting from uncertainties introduced by virtualization. Spirent is leading the efforts within ETSI NFV ISG to define pre-deployment methodologies for validating NFV environments. The use cases discussed in the table below are described in greater detail in GS NFV TST001[5]. This document will focus on a subset of the use cases, relating to VNF benchmarking and service chain/auto-scaling validation.

Use Case	Functionality	Speed, turn up rate	Reliability	Scalability
VNF testing– Lifecycle mgmt	VNF Instantiation VNF termination		VNF long duration testing	VNF scale
VNF testing– Data & control plane benchmarking	vBNG, vRR, vRouter, vCPE protocols	Protocol session bring up rate	RFC 2544 style sweep test, iMix, Integrity test, long term control plane stability, error recovery, convergence	Line rate data testing, Protocol scale testing
NFV Infrastructure– Benchmarking	Hypervisor, vSwitch functional testing		VNF cycle testing VNF long duration testing	Hypervisor, vSwitch scale testing, VNF scale testing
Network services– Lifecycle mgmt	NS instantiation & termination	NS instantiation time	NS cycle testing, fault detection and recovery	
Network services– Service chains & autoscale	Service chain & autoscale policy validation	Autoscale time	Scale up/down, in/out cycle test, NS validation post autoscale	NS chain scale testing

Table 1: Test methodologies for validating the most significant NFV use cases

5 ETSI GS NFV TST 001: "Network Functions Virtualisation (NFV): Pre-deployment validation of NFV Environments"

There are two goals when benchmarking VNFs or network services. One finds maximum performance given specific resources. The other finds the resources required to achieve a target performance level.

Goal 1: Maximum performance benchmark. Find the maximum performance of the VNF given a specific NFVI configuration.

In this test, the VNF and the vSwitch, if applicable, are assigned a fixed set of resources. As data plane traffic increases the stress on the VNF, the test captures performance metrics to find the maximum level of performance the VNF can deliver without errors.

This test can used to measure the maximum performance of the VNF under the fixed conditions or to validate that the VNF can achieve published performance metrics.

Goal 2: Resource allocation benchmark. Find the amount of resources required to achieve a specific performance level.

This test specifies a given set of performance values, similar to a service level agreement (SLA). As data plane traffic increases stress on the VNF, the VNFO allocates additional resources, such as processor cores or memory, to find the mix of resources required to maintain the performance requirements.

In addition to specifying the SLA, a desired level of resource allocation can be defined before testing. For example, if the goal of the test is to find the number of CPU cores and memory that must be allocated to a VNF to achieve the SLA of 10 gbps forwarding performance, the target maximum CPU core utilization might be 80 percent.

Best practices. Modify only one variable between iterations. For example, a test run could vary the number of processor cores per iteration while keeping all other resources constant. A second test run could vary memory allocation while keeping the number of processor cores fixed.

Maximum Performance Test Variable	Fixed/Variable per test Iteration/Measured
Virtualization layer, vSwitch, physical resources (including NICs)	Fixed
CPU Cores allocated to VNF	Fixed
Memory allocation to VNF	Fixed
Data plane acceleration	Fixed
vSwitch resources	Fixed
Multi-tenancy	Fixed (single tenant)
Core utilization	Measured
Performance metrics	Measured

Target Performance Benchmark Test Variable	Fixed/Variable per test Iteration/Measured
Virtualization layer, vSwitch, physical resources (including NICs)	Fixed
CPU Cores allocated to VNF	Measured
Memory allocation to VNF	Measured
Data plane acceleration	Measured
vSwitch resources	Measured
Core utilization	Measured
Number of VNFs required to achieve performance	Measured
Performance metrics	Fixed

VNF Benchmarking

The traffic that traverses a VNF is subject to reliability, QoE, and predictability requirements. These values are defined in the various information elements of the VNF Descriptor (VNFD) and stipulated to NFV consumers in the SLA. Data plane benchmarking evaluates these qualities of a VNF.

The test VNFs originate full mesh traffic toward the VNF under test (VNFUT) and evaluate the ability of the VNFUT to correctly forward the traffic by analyzing the frames received from the VNFUT. Basic analysis metrics include short-term, long-term, and average packet delay and packet delay variation, number of sequencing errors, and a comparison of offered versus measured bandwidth. Advanced analysis uses layer 7 workflows that are representative of real services. Analysis metrics include service reliability, service render latency, service errors, and service availability.

Basic Traffic Sweep Test Methodology

Goal: Benchmark the forwarding plane performance of the VNFUT.

Test iterations. This test runs multiple iterations while varying frame size or frame rate per iteration. The values for size and rate may vary according to the network function under test. A typical L2/L3 test for a vRouter would use the following values.

Frame rate in frames per sec (fps)	10, 100, 1000, 10000, 100000 . . . up to the forwarding performance target of the VNF
Frame size in bytes	64, 65, 128, 256, 578, 1024, 1280, 1518, 9022

The test runs fully meshed traffic between all ports connected to the VNFs for 120 seconds. The frame rate starts at 10 fps, ramping up the frame rate for each iteration, and then the test is repeated using the next frame size until the frame size set is exhausted. The result polling rate or test duration can be adjusted to match thwe performance capabilities of the VNFUT.

Results. Metrics are recorded at a polling rate of once per second and include:

- Received bandwidth on the Test VNF ports
- Total sequencing errors (frame loss, duplicate frames, out of order frames, reordered frames, late frames)
- Maximum and average frame delay and frame delay variation
- Utilization of resources allocated to the VNF (processor cores and memory blocks)

To achieve actionable results that are comparable across tests, the results take into consideration the underlying hypervisor and the efficiency of resources such as processor cores and memory blocks used by the VNFUT.

The median bandwidth received by the destination test port is reported per iteration, and also as a percentage of the offered bandwidth and per processor core.

The maximum packet delay and packet delay variation are also reported in the same fashion.

Median received bandwidth	Maximum packet delay	Maximum packet delay variation
Median received bandwidth / offered bandwidth	Maximum packet delay / offered bandwidth	Maximum packet delay variation / offered bandwidth
Median received bandwidth / number of processor cores	maximum packet delay / number of processor cores	Maximum packet delay variation / number of processor cores

Absolute metrics such as sequence errors or dropped packets are reported as a total count.

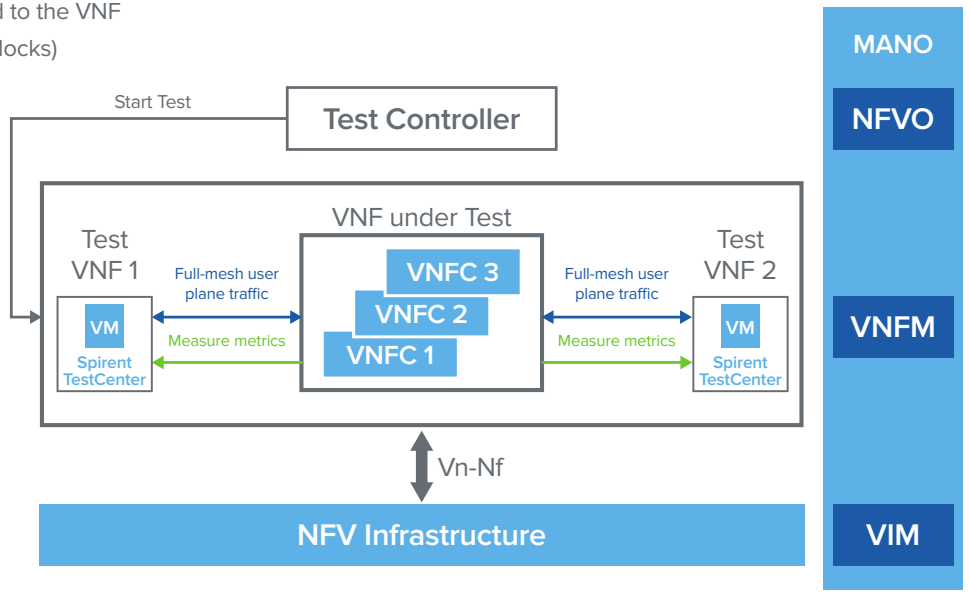

Figure 7: VNF traffic sweep test

Long Duration Traffic Test Methodology

Goal. Determine the stability and reliability of the VNF over time.

Duration. Achieving predictable performance is critically important in shared NFV environments. Select a specific combination of frame size and frame rate from the basic traffic sweep test that yielded zero frame loss and run a full mesh traffic test for a duration that matches the deployment needs and VNFUT capabilities. Typically the duration is six hours or longer.

Results. Review the reported metrics to determine whether the performance of the VNFUT is consistent throughout the test run. A one to two percent variation in performance over time is acceptable.

iMIX Sweep Test Methodology

Goal. Validate maximum performance against expected traffic conditions.

Traffic mix. This test uses the same methodology as the basic traffic sweep test, but with a mix of frame size and sequence that reflect aggregate traffic found on most public networks.

iMIX a		iMIX b		iMIX c	
Size (bytes)	Percent of bandwidth	Size (bytes)	Percent of bandwidth	Size (bytes)	Percent of bandwidth
64	58.33%	90	58.67%	90	50.67%
594	33.33%	92	2.00%	594	23.66%
1518	8.33%	594	23.66%	1518	15.67%
		1518	15.67%	9022	10.00%

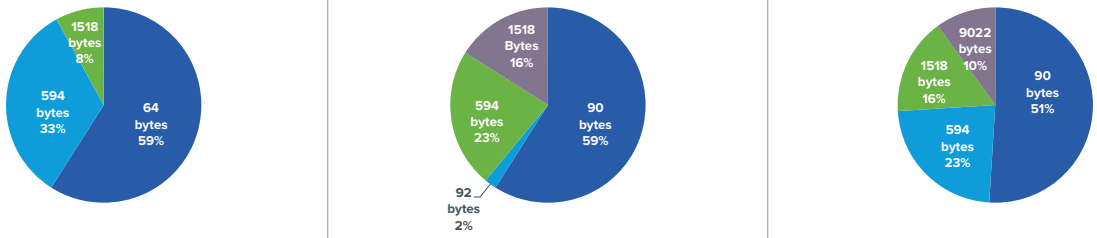

Figure 8: iMIX packet distributions

The IMIX Genome (RFC 6985) specifies the repeating sequence of frame sizes for each mix to achieve maximal repeatability. However, you can create a custom mix of frame sizes matched to a specific deployment using the methods described in RFC 6985 Section 1.

Size (bytes)	Custom Code Letter
64	A
90	B
92	C
594	D
1518	E
9022	F

IMIX	# Frames	Repeating Sequence
a	60	AAADD AAADDE AAADD AAADDE AAADD AAADDE AAADD AAADDE AAADD AAADDE AAADD AADD
b	50	BBBBBBDDEE BBBBBBDDE BBBBBBDDEE BBBBBBDDE CBBBBBDDE DDE
c	50	BBBDE BBDEF BBBDE BBDEF BBBDE BBDEF BBBDE BBDEF BBBDD BBDDF

Control Plane Benchmarking Methodology

Goal. Benchmark the control plane scale and performance of VNFUT.

The implementation of standards based protocols (such as BGP, OSPF, ISIS, LDP, RSVP) is identical in VNFs and physical network functions. As a result, the control plane benchmarking methods of VNFs are similar to those of physical network functions. The traditional benchmarking methods determine the maximum supported scale (number of simultaneous control plane sessions) and performance (rate of bring up of sessions) per port.

Unlike physical devices, where a bulk of the data plane processing is offloaded to FPGAs, ASICs and off-board processors, the compute cores in NFV environments are responsible for both fast path packet processing and the processing of control sessions and messages. Therefore, at any given time the scalability of the control plane can be influenced by the data plane load (fast path packets). VNFs that split control and data plane processing between different cores are likely to perform better than VNFs that share cores for those functions.

Performance and scalability targets. VNF vendors benchmark their VNFs for the maximum control plane scale and performance and the needed NFVI resources to achieve the performance levels. Network operators who deploy the VNFs have specific performance goals they want to achieve. For example an operator may have specific goals for a vPE such as supporting x BGP sessions, y PPPoE sessions, and z BGP routes per session while supporting data forwarding at n gbps. Goal seeking mechanisms help the operator determine the number of the vendor VNFs and NFVI resources required to meet his objectives.

Test setup. Using the published control plane benchmarks of the VNF vendor, the operator instantiates the appropriate number of VNF components (VNFCs) required to meet performance objectives. In this example, Spirent TestCenter VMs emulate CE and core routers running BGP and PPPoE, and establish the desired number of BGP and PPPoE sessions with the vPE under test. They also originate bi-directional user plane traffic at the desired forwarding rate.

Pass/fail. If the control plane sessions are formed successfully and user plane forwards traffic at the desired rate without errors or drops, the VNF provided by the vendor meets the operator needs. If not, the test moves to goal-seeking.

Goal seeking.
For each iteration of the test, modify the NFVI resources of the VNFUT in a stepwise manner. Change one variable at a time while keeping other NFVI parameters constant, within permitted constraints, and repeat the test until the desired performance is obtained. Employing a stepwise increase in VNF or NFVI resources maximizes performance levels.

Incremental configuration changes:

- Modify the number of VNFCs (VMs) of the VNF
- Modify the cores allocated per VNFC
- Modify the memory allocated per VNFC
- If possible and needed, change the allotment of cores between control and data plane processing for the VNFC
- Enable or disable acceleration techniques such as DPDK or SR-IOV if applicable

Results. When the VNFUT achieves the desired performance levels, record the VNF flavor used and number of VNFCs, number of cores per VNFC, number of cores allotted to control and data plane processing, amount of memory allocated per VNFC, and core utilization.

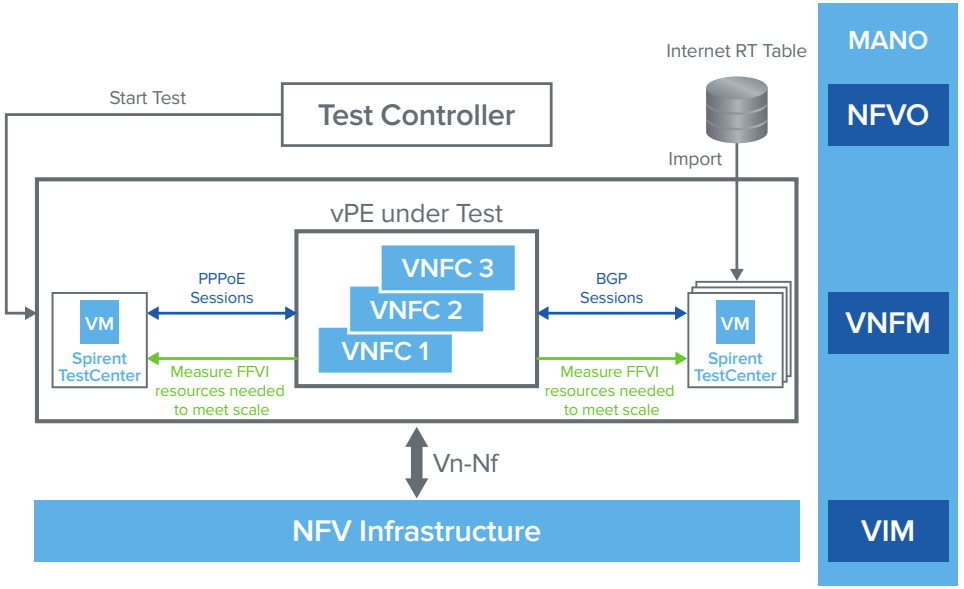

Figure 9: Control plane benchmarking test

Network Services Test Methodologies

A network service (NS) comprises a chain of service functions (forwarding graph of virtual or physical NFs).

The NFVO, in collaboration with the VNF manager, the VIM, and the OSS/BSS, manage the lifecycle of one or more NS[6]. The NFVO has the end-to-end view of resource allocation and serves as the single point of access for all requests from the OSS. The NFVO handles the lifecycle of NS and VNF forwarding graph. The VNF manager handles the VNF lifecycle from an application viewpoint.

In a shared NFV environment, multiple network services are executing on the same server, each at its own stage of the lifecycle. Some will be instantiating, scaling, or terminating while others are executing in a steady state. Lifecycle testing is essential to determine whether lifecycle changes of one NS is affecting other NS.

Network Services test methodologies validate the successful instantiation and termination of network services, measure the time needed to instantiate network services and ensure the successful completion of autoscaling. The NS test methodologies assume that the constituent VNFs of the network service under test (NSUT) have already been validated prior to the execution of the NS test.

Network Service Lifecycle Phases

NS on-boarding	Submit network service descriptor (NSD) to the NFVO to be included in the catalog. Validate the integrity and authenticity of the NSD and the presence of mandatory elements and required external interfaces.
NS instantiation	Perform a pre-instantiation validation and feasibility check, identify and reserve resources, instantiate VNFs, if necessary, and the connectivity network required. Connect the required VDUs to the connectivity network.
NS scaling/updating	Validate the request, check feasibility, determine scaling action (increase/decrease resources, instantiate/terminate VNF), allocate/de-allocate resources and connectivity.
NS termination	Validate request, terminate VNFs if necessary, delete resources, delete internal connectivity network, remove from catalog.

"Lifecycle testing is essential to determine whether lifecycle changes of one NS is affecting other NS."

6 The phases of the NS lifecycle are described in ETSI GS-MAN-00[i.8].

Network Services: Functional Test and Instantiation Time Methodology

Goal. Perform a functional validation of network services and measure the time required to activate services. This metric is an important measure of QoE experienced by customers. The NFVO manages dynamic instantiation and activation of network services.

In this test, an originating Test VNF sends traffic to the newly instantiated NSUT and ensures that the NSUT forwards the traffic correctly to a terminating Test VNF.

Test setup. Link three VNFs on an NFV server as a service chain, for example, a VNF forwarding graph consisting of a vCE, vFW, and vWAN accelerator. The test methodology assumes the constituent VNFs have already been instantiated before test execution.

Bracket the NSUT with Spirent TestCenter physical test ports. Because of the synchronization and microsecond accuracy required, physical test devices should be used.

Test procedure. The test controller instructs the NFVO to complete the instantiation of the network service at time T=t1. The NFVO notifies the test controller after it completes the NS instantiation. (See Annex C.3 of GS NFV-MAN 001 for detailed NS instantiation flows.)

The Spirent TestCenter ports generate appropriate bi-directional L2-L7 traffic toward the NSUT at a frame rate that matches the performance target of the NSUT. For example, a service function chain consisting of a vFirewall, vADC and vWOC receives application traffic; a service function chain consisting of a vCPE and a vBNG receives L2-L3 traffic traffic . It is recommended that many test passes are run, at different frame sizes. The exact set of frame sizes are dependent on the NSUT.

The test devices continue each pass of the test until time T=t2, when the service frames are detected at the terminating test device, after successful processing by the NSUT.

Pass/fail. If the results of the test report no errors such as reordered frames, data integrity errors, or CRC errors, the functional instantiation test passes. Otherwise, it fails.

Test results. The test devices record the QoE metrics defined in the monitoring_parameter field of the Network Services Descriptor and plot the time needed to complete network service activation for each combination of frame rate and frame size as indicated by the value [t2 – t1].

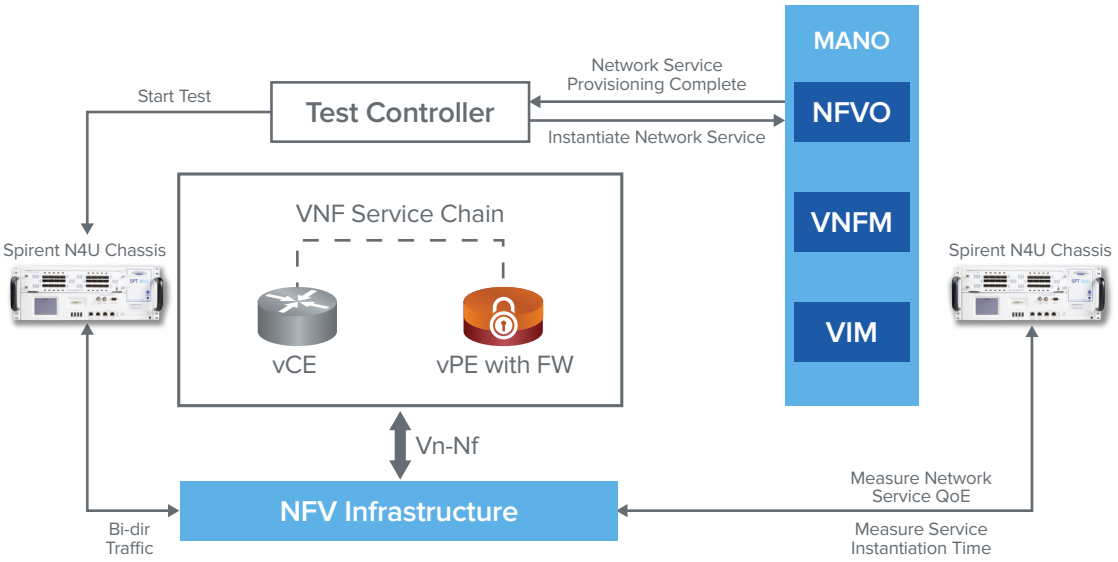

Figure 10: NS instantiation functional test

Network Services Instantiation: Auto-Scale Validation Test Methodology

One of the most significant drivers for NFV is the ability it provides to network operators to allocate resources when needed and contract resources when not needed. A network service can dynamically react to a sustained spike in customer traffic by scaling up or scaling out. Similarly, during periods of reduced customer traffic, it can scale down or scale in. Such elastic scaling capabilities prevent over-provisioning of network resources.

Scale up	One or more VNFs constituting the NS are allocated additional NFVI resources such as compute, memory, or storage
Scale out	Additional VNFs are instantiated in either the same server or another server, to handle the increased load
Scale down	One or more VNFs constituting the NS are allocated a lesser amount of NFVI resources such as compute, memory or storage
Scale in	Some of the VNFs constituting the NS are terminated in response to reduced load

Table 2: Network services auto-scaling

Goal. Validate the successful completion of auto-scaling and the maintenance of customer SLAs both during and after the completion of auto-scaling.

Test setup. For this test case example, set up a NS with a VNF forwarding graph that includes a vCE and vRouter on the same server. The NS provides end-to-end circuits with guaranteed SLAs. The test methodology assumes the NS has been successfully instantiated and is complaint with its performance target before test execution.

Bracket the NSUT with Spirent TestCenter physical test ports. Because of the synchronization and microsecond accuracy required to measure compliance to SLAs, physical test devices should be used.

Test procedure. At time T=t1, the test devices initiate an increase of traffic load (sustained increase or a traffic burst) sufficient to trigger auto-scale mechanisms at a time t1 + Δt. (The exact means by which the test devices obtain this knowledge is outside the scope of this document.) In response, the VNFs, VNF Manager and/or the NFVO initiate and execute auto-scaling.

Starting at time T=t1, the test devices also monitor the NS performance for adherence to SLAs. The sampling period is NS-dependent. In this example, once every 100 ms is recommended.

When the NS is able to support the higher scale without SLA degradation, the test devices record time T=t2.

Test results. The test results include several parameters, including the VNFs that constitute the NS, the trigger that caused auto-scaling, the traffic load prior to auto-scaling, and the traffic load after auto-scaling. The test devices also periodically record NS performance metrics and NFVI utilization metrics during the interval between t1 and t2, and report the duration of the auto-scaling process as indicated by the value [t2 – t1].

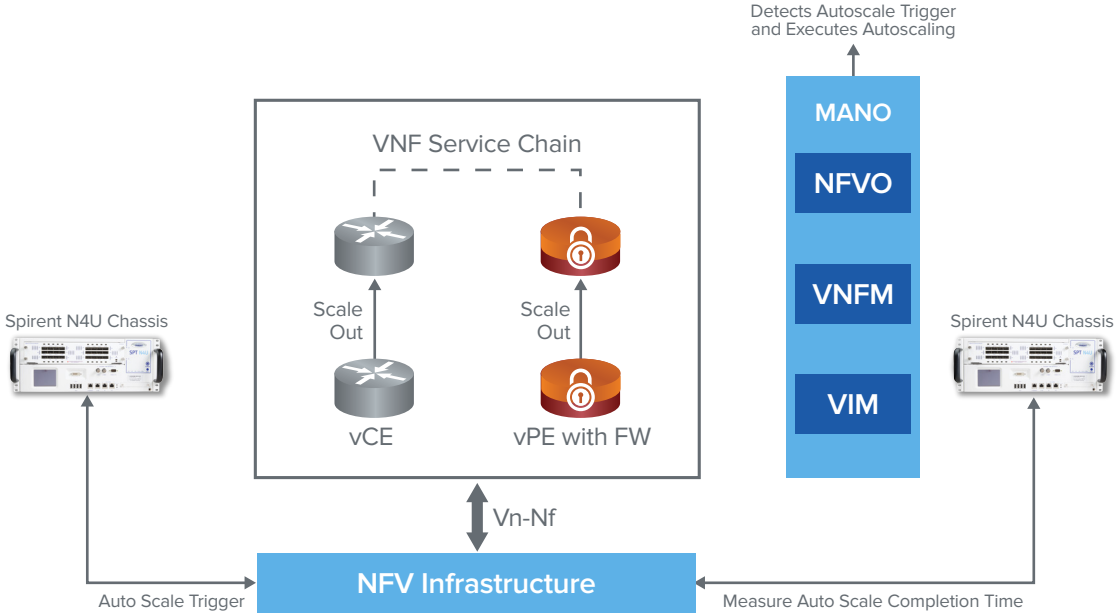

Figure 11: Network Services Scaling – Validating auto-scaling

Spirent: Leader in Virtualized Testing

With test solutions that address all the areas of concern shown in Figure 6: SDN/NFV testing and deployment challenges, including mobility data center, and access/edge, Spirent is the leader in validating network performance, availability, scalability, and security in virtual and physical networks.

Leading Role in Standards Development Organizations

Spirent plays an integral role in the ETSI NFV ISG, leading the development of guidelines for pre-deployment and post-deployment validation of NFV environments. Spirent also plays a leading role in the ONF Testing Council, making significant contributions toward OpenFlow controller and switch validation.

Ease of Use and Portability

Spirent's virtual and physical test platforms have an identical look-and-feel and support seamless interoperability. Test configurations and scripts are portable across the two platforms. Test traffic can originate and terminate on physical test devices only, virtual test devices only, or on any combination of virtual and physical test devices. Spirent TestCenter Virtual executes on a wide range of hypervisors including VMWare ESXi and KVM/QEMU. It is also compatible with open source cloud management systems, such as OpenStack.

Performance and Scalability Validation of VNFs and Network Services

Spirent solutions simplify the often daunting task of benchmarking hundreds of different VNF types and flavors by providing exhaustive support for L2-L7 data plane and control plane testing.

Spirent's virtual solutions are optimized using DPDK—enabling users to achieve significant improvements in data forwarding performance, while utilizing a lower number of compute cores for generating the test traffic.

Spirent's test methodology offerings also enable users to easily benchmark vSwitches, vRouteReflectors, vPE Routers, vBNG, vCPE routers, vFirewalls, vWAN Accelerators, vIDS, vIPS, and every element of the vEPC.

Product/Capability	NFV Application
Spirent TestCenter	Spirent TestCenter provides measurement solutions for next generation networks—from traditional performance testing to the rigorous analysis of virtualization, cloud computing, mobile backhaul, and high-speed Ethernet. Spirent TestCenter combines nanosecond level accuracy with extreme scale and high port density to test data center fabrics, and virtual appliances such as vBRAS, vPE, vCE, and vSTB devices.
Spirent TestCenter Virtual	Spirent TestCenter virtual is a software module that resides on virtual machines and servers. It extends and complements the capabilities of Spirent TestCenter to benchmark VNFs, NFV environments and cloud management platforms. With Spirent TestCenter Virtual, vendors and cloud service providers can validate virtual switches, routers, and firewalls using the same tests that have been used to validate physical devices for years. Spirent TestCenter virtual is capable of high performance using DPDK drivers and is supported on multiple hypervisors including: • VMWare ESXi 4.0/4/1/5.3/5.5 • KVM/QEMU on Fedora, CentOS and Ubuntu
Spirent Avalanche Virtual	Spirent's Avalanche virtual is a software cloud L4-L7 test solution that is based on the industry leading Spirent Avalanche platform. It is designed to test and measure the performance, availability, security, and scale of virtualized cloud environments. Spirent Avalanche virtual is compatible with multiple hypervisors including: • VMWare ESXi 4.0/4/1 • KVM/QEMU • Xen Server • Hyper-V
Spirent Landslide Virtual	Spirent Landslide™ virtual is a complete suite of advanced test elements for the mobile networks and services of tomorrow. Spirent Landslide Virtual emulates the control and data traffic of mobile subscribers moving through the network while using carrier and OTT services. The solution also incorporates a complete suite of mobile core, Diameter and IMS network nodes and interfaces. This enables complete end-to-end network validation or isolation of virtualized EPC, Wi-Fi controller, and Authentication, Authorization, & Accounting, Policy and charging functions. Control and reporting is supported by a web UI or orchestrated via REST API.

Acronyms

ASIC	Application Specific Integrated Circuit
BGP	Border Gateway Protocol
COTS	Commercial Off The Shelf
DCI	Data Center Interconnect
DPDK	Data Plane Development Kit
FPGA	Field Programmable Gate Array
ISIS	Intermediate System to Intermediate System
LDP	Label Distribution Protocol
MANO	Management and Orchestration
NF	Network Function
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
NFVO	NFV Orchestrator
NSD	Network Services Descriptor
NSUT	Network Services Under Test
ONF	Open Networking Foundation
OSPF	Open Shortest Path First
PoC	Proof of Concept
PPPoE	PPP over Ethernet
QoE	Quality of Experience
RSVP	Resource Reservation Protocol
SLA	Service Level Agreement
SR-IOV	Single Root I/O Virtualization
VIM	Virtualized Infrastructure Manager
VNF	Virtualized Network Function
VNFC	VNF Component
VNFD	VNF Descriptor
VNFUT	VNF Under Test
WAN	Wide Area Network

About Spirent

About Spirent Communications

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics and security, serving developers, service providers, and enterprise networks.

We help bring clarity to increasingly complex technological and business challenges.

Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information, visit:
www.spirent.com

AMERICAS 1-800-SPIRENT
+1-800-774-7368
sales@spirent.com

US Government & Defense
info@spirentfederal.com
spirentfederal.com

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

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

For more information on SDN and NFV Testing, please visit: https://www.spirent.com/Solutions/SDN-NFV-Testing.

spirent™
Promise. Assured.

Rev B | 09/18