WHITE PAPER

Communications
Data Center Solutions

Network Services Benchmarking:
Accelerating the Virtualization of the Network

Network Services Benchmarking (NSB), part of the OPNFV* Yardstick project, offers service providers, OEMs, software vendors, and systems integrators a framework for characterization and benchmarking of virtual network functions (VNF), network functions virtualization infrastructure (NFVI), and network services.

Table of Contents

The Need for Community-Driven Benchmarking

Across the industry, the deployment of virtualized networks and functions has been challenged by a lack of common standards and industry-accepted benchmarks for conformance to carrier-grade requirements. This has made it difficult for communications service providers (CoSPs) and other stakeholders to effectively evaluate the network solutions on the market, compare how different vendor solutions may impact their existing network requirements, and estimate total cost of ownership (TCO) when it comes to network transformation.

At the same time, more effective tools to characterize network performance—along with the ability to dimension network workloads and model the impact of stress vectors such as traffic throughput—would give independent software vendors (ISVs), original and telecommunications equipment manufacturers (OEMs/TEMs), and systems integrators (SIs) the ability to make (or implement) best-in-class products that meet the needs of network operators, and to bring those solutions to market more quickly.

To meet these challenges, Intel contributed the Network Services Benchmarking (NSB) project along with industry partners and the Open Platform for NFV (OPNFV) community. The NSB framework features were added to the Yardstick tool to support both network functions virtualization infrastructure (NFVI) and virtualized network functions (VNFs) characterization with the goal of facilitating deterministic and repeatable benchmarking. NSB has the flexibility to help CoSPs, ISVs, VNF vendors, OEMs/TEMs, and SIs find performance bottlenecks and evaluate VNF/NFVI solutions using standard benchmarks, a key factor in enabling network functions virtualization (NFV) deployments to help networks become more cost-efficient and agile. NSB also added representative workloads under SampleVNF project to demonstrate the characterization and benchmarking using Yardstick tool.

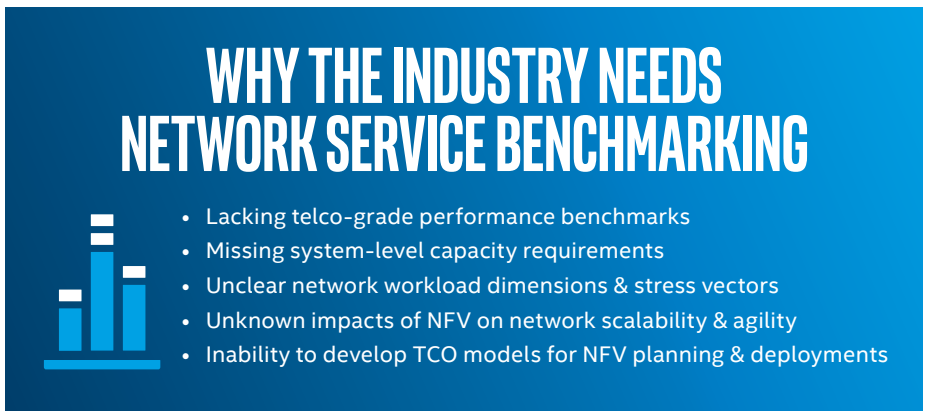

Figure 1. NSB solves the industry problems: five steps

NSB solves the industry problems listed above in the section titled "Why the Industry Needs Network Service Benchmarking" in five steps. Ultimately, the goal of NSB is not simply to measure performance at the packet-processing level. It's about giving CoSPs, VNF vendors, and others a test architecture to identify service-level benchmarks they can use to create a better product and improve quality of service for their customers. This document introduces and describes the NSB project, outlines the benefits to various stakeholders, then discusses the process for using NSB within your own product planning and NFV adoption. It ends with a call to become more involved in the Yardstick project.

Continue reading to learn how NSB can positively influence your network's resources, efficiency, and capabilities.

What Is Network Services Benchmarking?

The goal of Network Services Benchmarking (NSB) is to extend OPNFV Yardstick framework to perform real-world VNF and NFVI characterization and benchmarking with repeatable and deterministic methods. The NSB extensions to Yardstick were introduced to identify performance bottlenecks telecommunications services providers may encounter in VNF development and/or deployment to achieve overall system performance.

By using NSB and an external traffic generator, a user can benchmark and test many different parameters in an automated fashion, while keeping a record of all the tests done. This becomes extremely useful when trying to understand the behavior of a VNF within a system or to evaluate the VNF against other vendors or versions. It can also give service providers additional data with which they can decide whether or not to deploy a new VNF in their networks.

With NSB, providers can characterize the performance of a VNF solution working in different virtual environments; compare existing VNF solutions from different VNF vendors, and plan for capacity and network changes against network requirements. The starting point of NSB is to obtain characterization data—running tests to capture and analyze performance data in the lab. The results of these characterization tests can be used in multiple ways:

- **Benchmarking:** Formalizing characterization methodology allows for more exact comparison of VNFs and NFVIs when selecting a new solution. NSB can also allow companies to adopt external benchmarks to compare the performance of their current solution.
- **Modeling sizing and performance:** Characterization data can be used to predict approximate system performance under increased traffic workloads or feature changes to evaluate new business use cases or prepare for their requirements.
- **Modeling revenue and TCO:** This data can also be used to estimate the bill of materials (BOM), cost, and revenue as part of the planning process.
- **Modeling MANO/OSS compliance:** Characterization can help companies understand how MANO components impact performance with policy-based networking.
- **Optimizing performance:** Lastly, characterization data can highlight areas to focus efforts for performance optimization.

NSB can be used to characterize both VNF network performance metrics (e.g., throughput, latency, jitter, session scale) and NFVI resource utilization statistics under various configurations and user workloads. It enhances the Yardstick test framework with the following capabilities:

- NFV infrastructure benchmarking includes bare-metal, standalone virtualized and managed virtualized (SR-IOV and virtual switching based) infrastructure
- Definition of the Network Service topology, including control plane and data plane interfaces
- Standardized VNF/Network Services (NS) models compliant with the Internet Engineering Task Force* (IETF) definition
- Generalized VNF models written in Python* supporting any type of NFVI defining instantiation, termination, and configuration
- Generalized traffic profiles
- Generic models for traffic generators supporting different vendors

The Yardstick code base extensions includes generic data models of network services based on ETSI specs. It allows operators to gather data such as system agent support for NFVI (e.g., CPU statistics, memory BW, and OVS-DPDK stats); network KPIs (e.g., in packets, out packets, throughput, latency, etc.); and VNF KPIs (e.g., packet in, packet drop, packet fwd, etc.).

VNF Test Cases

NSB test cases have multiple configuration parameters in the OPNFV Yardstick project. These are designed to test virtual network functions such as:

- CG-NAPT (Carrier Grade Network Address and Port Translation)
- vACL (Virtual Access Control List)
- vFw (Virtual Stateful Firewall)
- vPE (Virtual Provider Edge Router)

A typical block diagram outlining the NSB components for deploying these sample VNFs is outlined below in Figure 2.

For further information on the sample VNFs please refer to the SampleVNF OPNFV page.

Supported Traffic Generators

Because NSB is an open solution, every user can add their own traffic generator to achieve the desired traffic profiles. Supported profiles are intended to mimic real world traffic profiles implementing full L2-L3 and L4-L7 traffic generation. However, one goal of NSB is to make testing possible without dedicated testing hardware, which is why NSB supports a number of open-source traffic generators in addition to commercial traffic generators. They include:

- Ixia IxNetwork/IxLoad*
- Spirent Landslide*
- DPDK Based:
  - TRex*
  - Packet pROcessing eXecution (PROX)
  - Pkt-gen

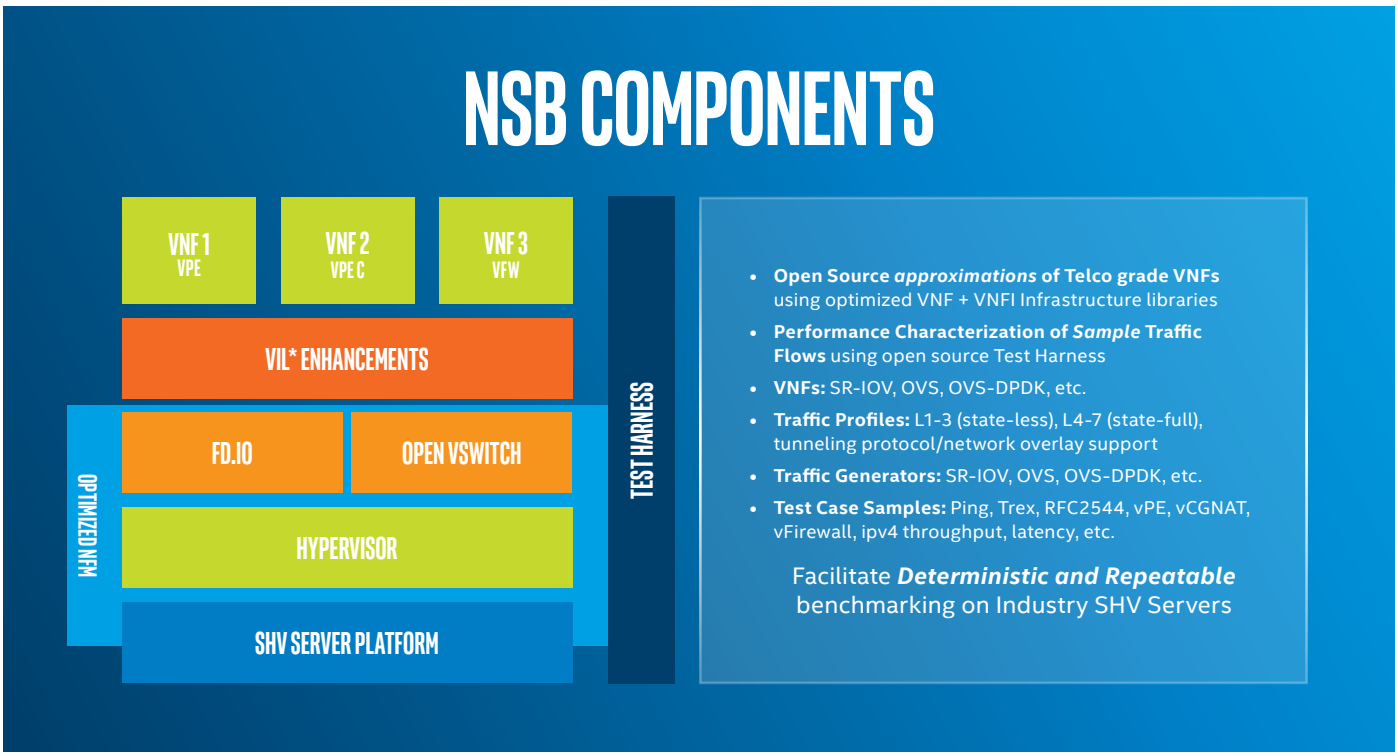

Figure 2. NSB block diagram for sample VNF deployment

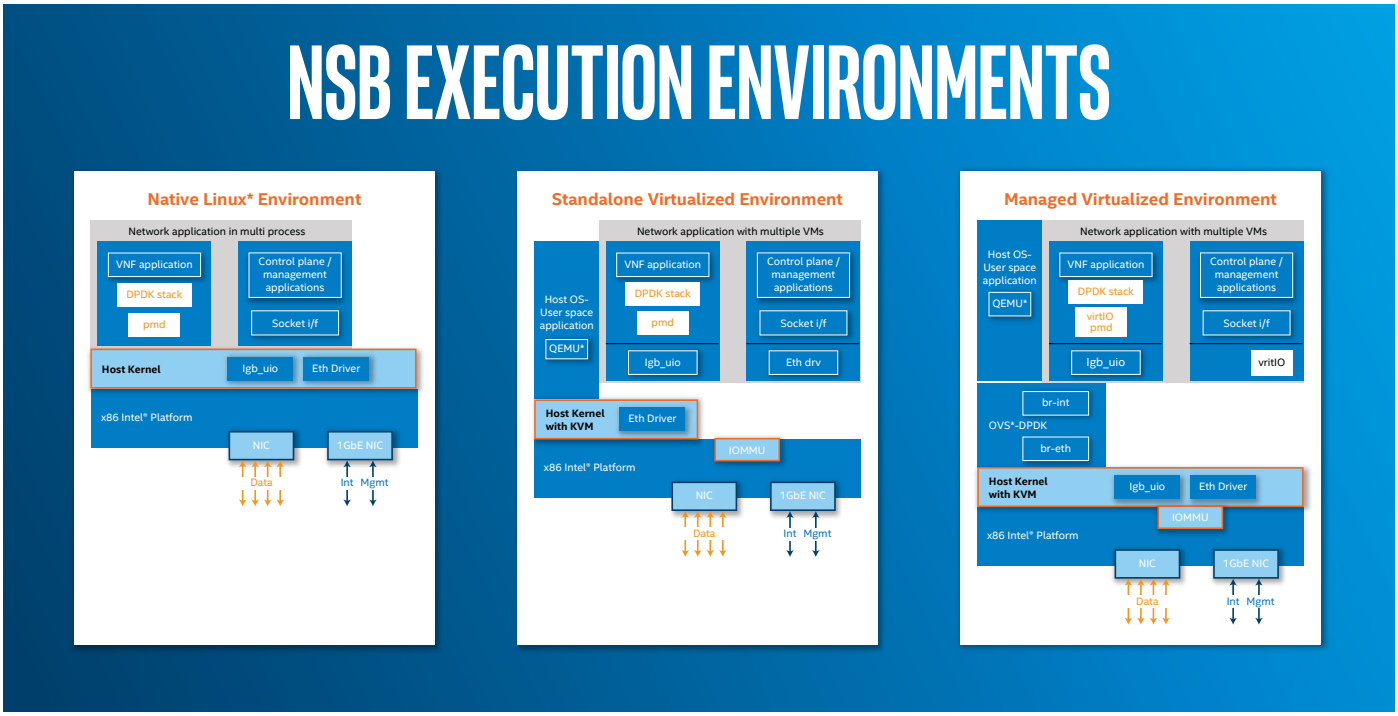

Figure 3. NSB execution environments: Bare metal, standalone virtualized, and managed virtualized (with OpenStack as an example)

Supported Execution Environments

For both VNF and NFVI performance benchmarking, NSB can be applied in bare metal environments (Native Linux*), standalone virtual environments (PCI pass-through, SR-IOV, OVS-DPDK), or managed virtualized environments (e.g., OpenStack*). Figure 3 above gives an overview of the three execution environments that NSB provides to users for performance benchmarking.

NSB offers the capability to interact with both hardware and software-based external traffic generators, for triggering and validating traffic patterns according to user-defined traffic profiles.

VNF Performance Benchmarks

The following test methodology, also shown in Figure 4, gets adapted in NSB. The VNFs will be made to run on all the above three different environments (Native Linux*, Standalone Virtualized and Managed Virtualized), and for each characterization test case, simultaneously three different KPIs—Network KPI, NFVI KPI, and VNF KPI—will be collected. This allows operators to understand baseline VNF performance and investigate performance "hot spots" to understand the infrastructure overheads and optimize the VNF. The framework tool enables evaluation of both scale-up and scale-out benchmarking scores of a given VNF/services.

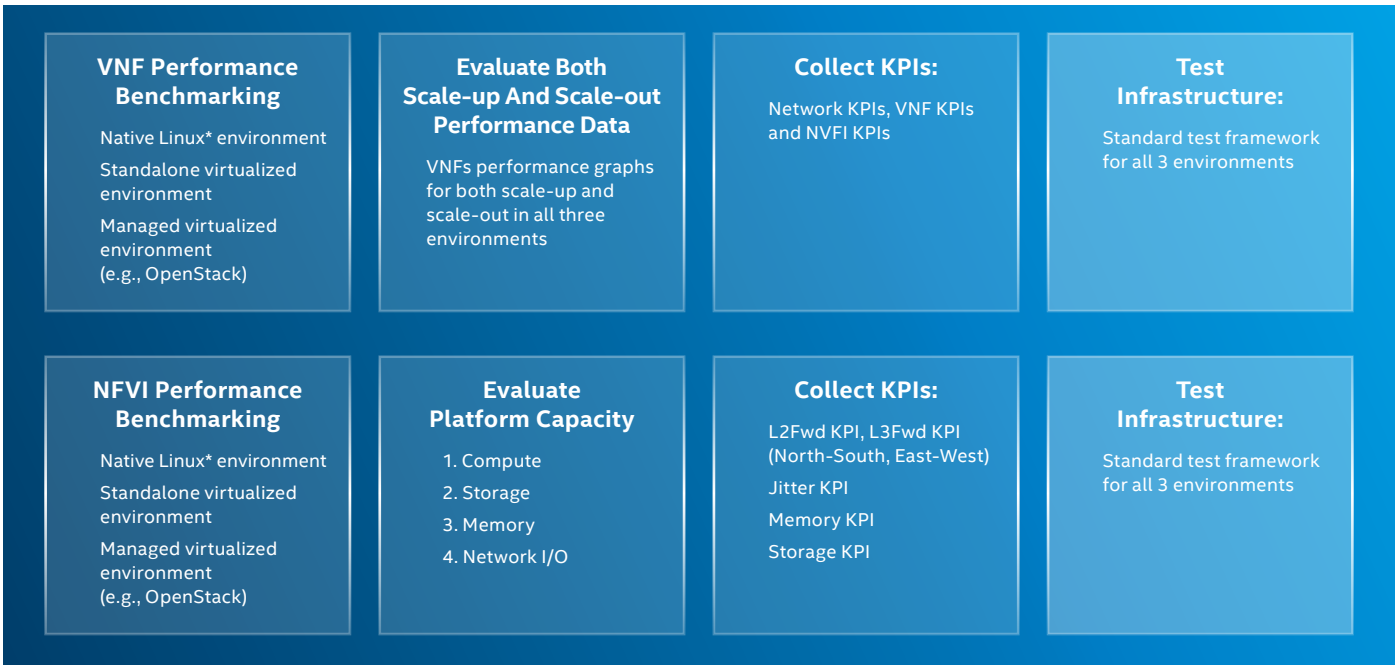

Figure 4. NSB testing methodologies

Network Metrics

- Zero-loss packet throughput as percentage of line rate (test defined by IETF RFC 2544)
- Packet throughput with a threshold packet loss (by default = 0.02%)
- Network KPIs as function of system load (packets/sec), measuring latency, delay, jitter, and loss statistics
- Traffic verification, validation of "goodput" (correct packet throughput) in test cases

NFVI Metrics

- CPU, memory, network, and storage utilization as stats collected and archived as time series
- System-level performance stats, such as total simultaneous sessions as function of traffic intensity

VNF Metrics

- **Compute:** Sustained and burst bandwidth, error/resend rate, dropped packet rate, CPU utilization, L2Fwd/L3Fwd (North-South and East-West)
- **Memory:** Memory utilization, kernel time, user space time, Cache utilization
- **Network I/O:** Throughput, latency, packet loss

NFVI Performance Benchmarks

NSB helps operators understand the NFVI "dimensioning" problem (i.e., size of NVFI required to support network service). It does this by measuring NFVI under test against a fixed suite of sample VNFs to understand the sensitivity of NFVI configuration to hardware, operating systems, and other variables.

Sample NFVI Test Cases

Network NFVI test cases include:

- Packet forwarding (no touch): "no-op" VNF, baseline
- L3 Forwarding: packet forwarding modifying MACs
- Multi-flow L2 Forwarding: 200K flows, SoftSwitch stress-test
- MPLS tagging: protocol conversion, adding/removing MPLS tag, packet length variation
- ACL: flow matching Access Control List, complex packet filtering
- LB/5-tuple lookup: 5-tuple based flow matching table lookups for load balancing
- Buffering: packet flow buffering for at least 125ms, stresses cache and memory
- BNG: ARP, QinQ, LB, Routing, GRE, MPLS
- BNG + Qos
- vPE: ACL filtering, flow classification, routing (LPM lookup), metering, policing, and marking
- Iw-AFTR: lightweight Address Family Translation Router: IPv4 <—> IPv6

Who Will Benefit from NSB?

VNF vendors, software vendors, equipment manufacturers, and service providers all need better VNF characterization tools to help them fine-tune and optimize the performance of their NFV infrastructure and VNFs, as well as to determine efficient utilization of their infrastructure resources. NSB solves this problem, helping a variety of users perform characterization and benchmarking of VNFs. With any of these use cases, NSB can be seamlessly integrated into the service lifecycle of an organization's validation activities.

NFV/VNF Vendors

Understanding system-level capacity requirements is critical for deploying agile and scalable NFV environments that are competitive. Using NSB can help network solutions vendors develop products that perform to spec. At the same time, it can assist them in meeting the quality and conditions their customers expect—using characterization data to model the impact of traffic workloads and performance on specific infrastructure. This not only helps them optimize network performance to the needs of their customers, but it also gives them a better understanding of reliability, performance, and function.

Operating System/Independent Software Vendors (OSVs/ISVs)

Having confidence in software performance and speed on a wide range of network conditions is critical for OSVs/ISVs. Using NSB to model use cases can help them understand the impact of dynamic traffic workloads and perform capacity planning against given network requirements.

Telecommunications/Original Equipment Manufacturers (TEMs/OEMs)

NSB can help manufacturers understand how their hardware performs under a variety of traffic loads and network functions. With this knowledge, manufacturers can more precisely target performance optimization efforts and deliver high-quality products that have been tested in a range of highly-regulated environments. NSB can also help manufacturers model operating costs and project revenue.

Communications Services Providers (CoSPs)

A key challenge for CoSPs is the need to compare the performance of commercial VNFs, both during the initial purchase consideration and as software applications within the network and traffic patterns on the network mature. NSB helps CoSPs gather the concrete performance characterization data they require to effectively compare solutions from VNF vendors, including modeling total cost of ownership (TCO) and bill of materials (BOM) requirements.

Additionally, NSB can assist CoSPs in understanding existing system-level capacity requirements to model mobile application loads, as well as to know how the infrastructure will react as traffic workloads scale. This helps them provide a top subscriber quality of service, speed new services to market, and lower operating costs.

Example Use Cases for NSB

This section covers an example characterization of a virtual firewall VNF application service.

Concurrent TCP Connection Capacity

NSB can be used to test what limitations are imposed on a virtual firewall application through measuring certain KPIs, including modeling concurrent connections. This test is done as per the rfc3511 section 5.2.3 procedure. The goal is to characterize session management (memory usage), as well as the impact of the scanning policies on the CPU usage. Variables used can include number of CPUs, number of ports, and type of interface (SR-IOV) and OVS-DPDK.

Maximum HTTP Throughput

NSB can be used to test how allocated memory and number of CPUs affect how a virtual application processes incoming traffic via defined policies. The maximum HTTP transfer rate test is carried out as specified in rfc3511 section 5.7.3 procedure. The goal of the test is to get information on virtual applications capability to achieve the throughput rate at a given CPU utilization based on policy scanning parameters and various HTTP transaction sizes.

Maximum Transactions Per Second

NSB can be used to understand how a virtual application behaves when trying to allocate memory as fast as possible. This test is done as per rfc3511 section 5.3.3 procedure.

What Is the NSB Process?

Step 1: Select the Stack	Choose hardware & software for VNF characterization based on general industry need and/or specific business use case.
Step 2: Integrate VNFs	Integrate VNFs into all three environments (Native Linux, Standalone Virtualized, Managed Virtualized) and the test harness.
Step 3: Validate & Debug	Run basic tests across the three environments to ensure the VNF integration is bug-free, and characterization results are meaningful.
Step 4: Optimize Infrastructure	Implement best-known configuration of BIOS, OS, Infrastructure & Intel® architecture EPA features to optimize VNF performance for the NFVI environment.
Step 5: Gather KPIs	Gather network, VNF, and NFVI KPIs to understand scale-up and scale-out, or platform capacity (compute, storage, memory, network I/O).

NSB can be utilized for characterization within different layers of the stack, dependent on goals. NFVI can be evaluated against reference VNFs to insights into sensitivity to hardware, operating system, configurations, and more. VNFs can be characterized to determine performance within selected infrastructure environments and scale-up or scale-out metrics. VNFs operating as a single application can be measured against a reference NFVI configuration to understand dimensioning issues.

Before beginning the process, identify the problem to be solved. For example, a CoSP may be considering VNF vendors and want to know if the VNF will work with their existing infrastructure.

Next, define the hardware and software stack to be used in the VNF characterization tests. This can be driven by industry need or the specific VNF use case the user wishes to learn more about.

Sample Software Component

To characterize the virtual application with NSB framework, multiple system under test (SUT) setups may be used depending on the outcome desired. For example, testing the virtual application in a standalone virtualized environment (SR-IOV, or OVS-DPDK configurations) as opposed to testing OpenStack (Pike) deployed on the application. Because the test harness server resides outside the SUT server, it may be used across multiple SUT configurations.

How Can I Get Involved with the NSB Project?

Since NSB was launched, the goal has been to develop a common testing framework to accelerate the growth of the NFV ecosystem as we prepare for the arrival of 5G. As Intel continues working with the OPNFV community in the Yardstick project, we are actively encouraging participation and collaboration from all stakeholders in the community.

CoSPs can participate by contributing KPIs, traffic profiles and other metrics to ensure the viability of test scenarios as virtualized networks evolve, mature and scale. Developers can improve the OPNFV platform by contributing to the Yardstick project and continuously testing compatibility across hardware and software configurations. Vendors can expand NSB by contributing new scenarios, use cases, and combinations of components.

To learn more, connect with the Intel® Network Builders, an ecosystem of industry stakeholders focused around the same goal: to help the industry accelerate the adoption of NFVI, VNFs, SDN, and other transformational network technologies by using objective characterization to pave a path to predictable NFV deployments.

Get More Information

For more information about Intel and Network Services Benchmarking, visit **https://networkbuilders.intel.com/blog/network-services-benchmarking-from-vision-to-reality**

1018/DO/PMG/PDF

Please Recycle

338277-001US