

Using vBNG for 5G and FTTH Network Demands

Testing from Intel and netElastic demonstrates high throughput with low packet loss

Authors

Ai Bee Lim

Intel Corporation

Shivapriya Hiremath

Intel Corporation

David Williams

netElastic

Kerry Zhang

netElastic

Paul Lancaster

Red Hat

Telecom network data volumes are expected to continue to grow rapidly, requiring communications service providers (CommSPs) to continue to scale their networks to accommodate this growth with a minimal impact on operating and capital expenses. Virtualization is a critical part of this strategy as it dramatically expands network flexibility and scalability with a minimal impact on expenses when compared with using fixed-function network appliances.

To demonstrate the performance of a virtualized solution, Intel and netElastic jointly tested the netElastic virtual broadband network gateway (vBNG) on an Intel® Select Solution for NFVI Forwarding Platform. The testing showed that it is possible to achieve performance ranging between 270 Gbps and 380 Gbps, depending on packet size, while maintaining 0.001% packet loss.¹ It also demonstrates how the use of virtualized dense server configurations can enable CommSPs to scale their networks to accommodate the rise in internet traffic from fiber-to-the-home (FTTH), 5G, and internet of things (IoT) and achieve new levels of throughput.

netElastic Virtual BNG

The VNF used in the tests was the netElastic Virtual BNG, which is designed to provide the same services as an appliance-based BNG in a CommSP network, including establishing and managing subscriber sessions and providing services such as:

- Authentication, authorization, and accounting (AAA) for each session
- IP address assignment
- Carrier-grade network address translation
- Multi-protocol label switching/label distribution protocol L2/L3 virtual private network
- Access control lists
- Policy management
- Quality of service

To support real-time packet processing and high data throughput, netElastic designed the vBNG with the following features (see also Figure 1) designed to boost throughput:

- Software defined networking (SDN) architecture: The netElastic vBNG separates the control plane and data plane functions. This allows each to be run on their own separate hosts or virtual machines. This architecture allows the data plane to be scaled independently, providing maximum deployment flexibility. The vBNG's SDN architecture also improves network flexibility and simplifies infrastructure management with software-based control.

Table of Contents

netElastic Virtual BNG 1

Intel® Select Solutions for NFVI
Forwarding Platform Hardware... 2

Verification Is a Key Part of
Intel Select Solutions' Value
Proposition 3

Test Topology 3

Test Results 5

Conclusion 6

- Data Plane Development Kit (DPDK): netElastic vBNG uses DPDK combined with netElastic's optimized code in the data plane to help maximize throughput.
- Dynamic device personalization (DDP): vBNG uses the DDP technology built into the Intel® Ethernet 700 Series Network Adapters used in the tests. The Intel Ethernet

700 Series Network Adapters offer a programmable pipeline that allows for customization to meet network needs. The vBNG also leverages DDP to dynamically allocate CPU cores based on actual network demands, delivering high performance when needed and freeing up resources when demand is lower.

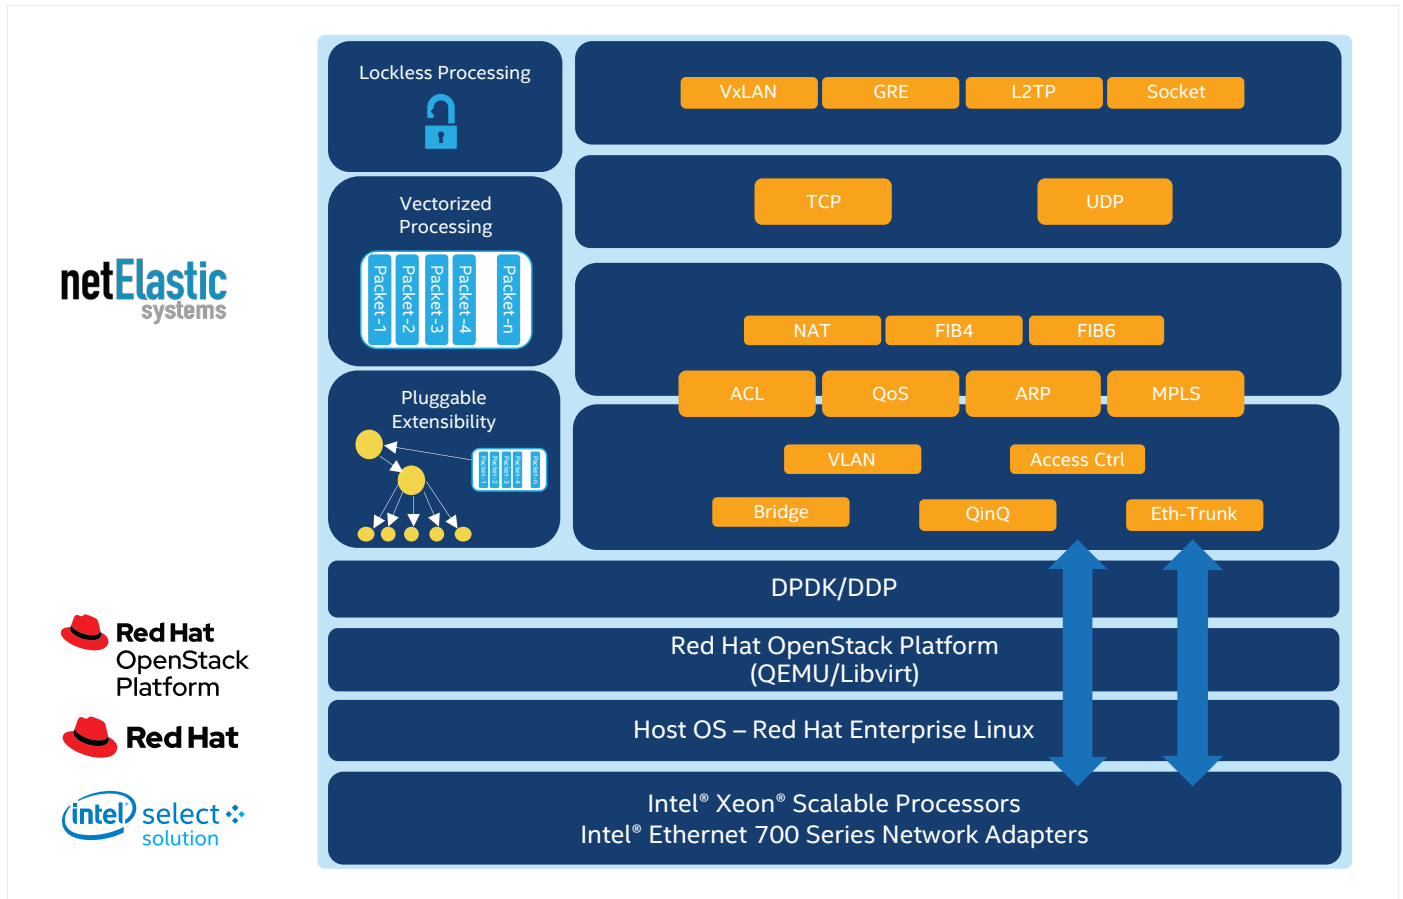


Figure 1. netElastic vBNG Architecture

netElastic's approach enables the vBNG to provide CommSPs with maximum deployment flexibility to deliver new services quickly and efficiently, whether it's deploying a new rural network or upgrading a large-scale metro point of presence (POP). The vBNG can be deployed for very small subscriber bases (e.g., a few thousand people) all the way up to millions of subscribers. Service providers can avoid large upfront expenses with netElastic's "pay-as-you-grow" licensing options.

Intel® Select Solutions for NFVI Forwarding Platform Hardware

Intel® architecture-based servers deliver requisite performance for virtualization across a family of CPUs that can be deployed cost-effectively in different locations throughout the network. As higher performance services have become virtualized, the capability of the hardware platform to support increased data throughput has also expanded.

Intel has developed the Intel® Select Solutions for NFVI Forwarding Platform reference design targeted at applications that are reliant on high capacity packet forwarding to meet service level expectations. The hardware for this test was compliant with a version of the reference design that has been verified on certified Red Hat Enterprise Linux 7 hardware and Red Hat OpenStack Platform 13. These reference designs provide a recipe for optimized solutions that offer the verified software and hardware foundation for high-throughput packet processing virtualized solutions.

The Intel Select Solutions for NFVI Forwarding Platform are designed for virtualized applications. In addition to a vBNG, some example network service workloads include 5G user plane function (UPF), cable modem termination system (CMTS), evolved packet core (EPC), and secure gateways. The Intel Select Solutions for NFVI Forwarding Platform are suitable for deployment in access central office (CO) and remote CO, and all the way into centralized core site network locations.

The Intel Select Solutions for NFVI Forwarding Platform are designed to help maximize the input/output throughput density per system. There are two configurations of the Intel Select Solutions for NFVI Forwarding Platform:

- The plus configuration is tailored for high performance and high density for maximum I/O packet processing.
- The base configuration is a value/performance-optimized offering.

Verification Is a Key Part of Intel Select Solutions' Value Proposition

Intel Select Solutions are carefully and thoroughly verified by Intel to help ensure the platform is proven, stable, and reliable. Intel is committed to continuous improvement of its verification capabilities and introduces new NFV use cases to produce the hardened hardware, firmware, and software foundation that is verified for CommSPs to use to build their production workload for large scale deployment. As a result, this reduces early integration issues, simplifies evaluations, and accelerates time to market.

Test Topology

The netElastic tests utilized a server based on the Intel Select Solutions for NFVI Forwarding Platform plus configuration.

These servers feature a hyperconverged infrastructure in a 1U or 2U rack-mounted configuration, with solution components and configurations selected to ensure maximum network throughput.

The system under test (SUT) was based on a 2nd generation Intel® Xeon® Gold 6258R processor. These CPUs operate at 2.7 GHz, which can be increased to 4.0 GHz with Intel® Turbo Boost Technology mode turned on. The devices also have more processor cache for added performance in frequency-fueled workloads. The increased frequencies and processor cache deliver built-in value.


The reference design also features Intel® Ethernet 700 Series Network Adapters for DPDK-accelerated networking, and Intel® Solid State Drives (Intel® SSDs). Intel® Optane™ DC persistent memory can be added to the configuration to provide large memory resources that increase the amount of data that can be held in close proximity to the processor.

This NFVI platform has the technology and features to meet the specific needs of NFVI applications by delivering an I/O optimized data path design into each non-uniform memory access (NUMA) node.

The testing was conducted by Intel in March 2020. Complete hardware and software details are shown in Tables 1 and 2.

PROCESSOR	2X INTEL® XEON® GOLD 6258R PROCESSORS (28 CORES, 2.7 GHZ)
Memory	12x 32 GB 2933 MTs DDR4 DIMMs (Total of 384 GB)
Network Adapters	4x Quad-port 25 Gbps Intel® Ethernet Network Adapter XXV710
Microcode	0x5000002c
Intel® Hyper-Threading Technology	On
Intel® Turbo Boost Technology	Off
BIOS	Release Date: 05/15/2019
System DDR Mem Config: slots / cap / speed	12 slots / 32 GB / 2933 MTs
Platform	Intel® Select Solution for NFVI Forwarding Platform-aligned platform
# Nodes	1
# Sockets	2
Cores/socket, Threads/socket	28/56

Table 1. System under test hardware



OS	RED HAT 7.7
Kernel	3.10.0-1062.12.1.el7.x86_64
Workload	netElastic vBNG-xeon-release-3.2.11
Compiler	4.8.5 20150623 (Red Hat 4.8.5-39) (GCC)
Intel® Ethernet Network Adapters XXV710 driver	2.8.10-k
Intel® Ethernet Network Adapters XXV710 firmware	Firmware-version 6.02
QEMU/Libvirt	libvirt 4.5.0/QEMU 2.12.0
Huge pages	160x 1 GB huge pages configured (80 GB per Virtual Machine)
Boot configuration	Max performance with virtualization (Intel® Hyper Threading Technology Enabled; Intel® Turbo Boost Technology Disabled)
Boot Profile	<pre> BOOT_IMAGE=/vmlinuz-3.10.0-1062.12.1.el7.x86_64 root=/dev/mapper/ rhel-root ro crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet iommu=pt intel_iommu=on,eth_no_rmr vfio_iommu_type1.allow_ unsafe_interrupts=1 pcie_acs_override=downstream selinux=0 enforcing=0 usbcore.autosuspend=-1 isolcpus=2-95 rcu_nocbs=2-95 nohz_full=2-95 kthread_cpus=0,1 irqaffinity=0,1 nmi_watchdog=0 default_hugepagesz=1G hugepagesz=1G hugepages=160 LANG=en_US.UTF-8 </pre>

Table 2. System under test software

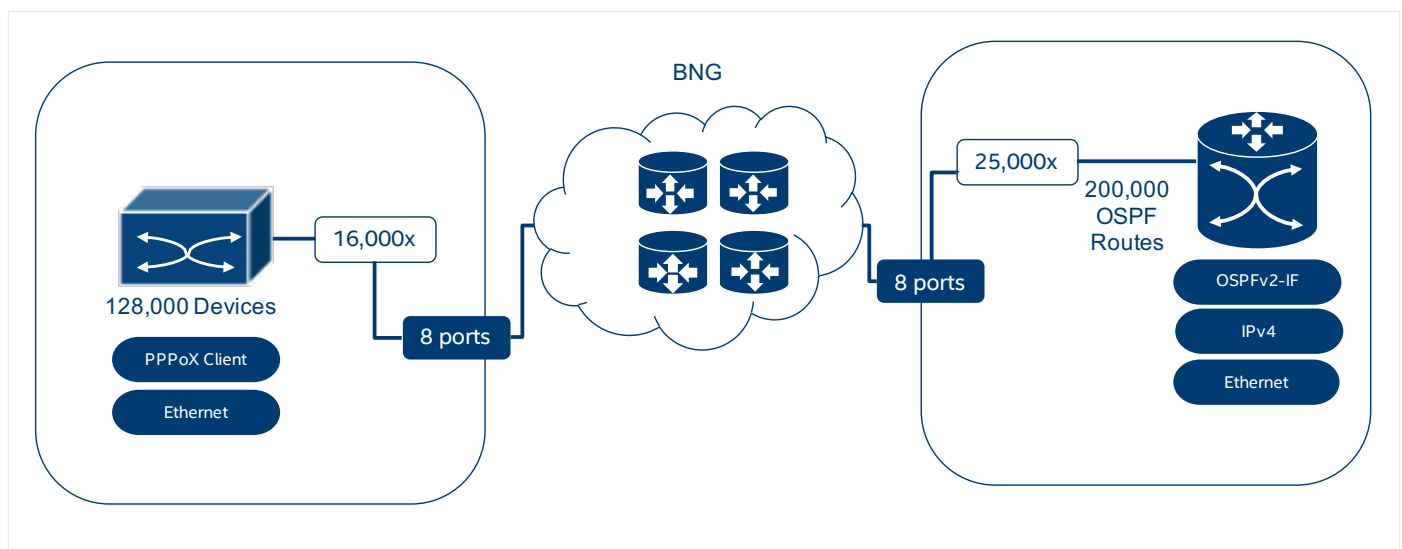
The overall test topology is shown in Figure 2. The vBNG system under test consisted of two vBNG virtual machines loaded on a single test server with two NUMA zones. Using NUMA configures the CPU and memory in a way that allows multiple virtual machines or containers to share memory locally for improved performance. Each NUMA zone consists of a 28-core Intel Xeon Gold 6258R CPU (48 logical cores per socket were used in the tests) and 80 GB memory with huge pages memory backing. Two quad-port 25 Gbps Intel Ethernet Network Adapters XXV710 were aligned to each CPU socket, thus providing eight 25 Gbps ports into the vBNG.

Each vBNG supports a maximum of 128,000 sessions; in this testing, only 64,000 sessions per vBNG were set up due to a limitation of the packet generator. Additionally, on the network side, 25,000 open shortest path first (OSPF)

routes per port were set up, providing a total of 200,000 OSPF routes in this testing. The user network interface (UNI) indicates a user request into the service provider's network. The network-network interface (NNI) represents the network routing within the service provider's network. The traffic generator emulates subscriber traffic and the network traffic for testing.

Downstream traffic is defined as the network traffic flowing between the network to the users, while upstream traffic is defined as data flowing from the users to the network.

In this testing there were a total of 16 logical cores unused across the two CPUs that made up the system under test. This means that the server had 16% compute headroom to add future services or use for other applications.

**Figure 2.** Test topology for performance tests conducted by Intel and netElastic

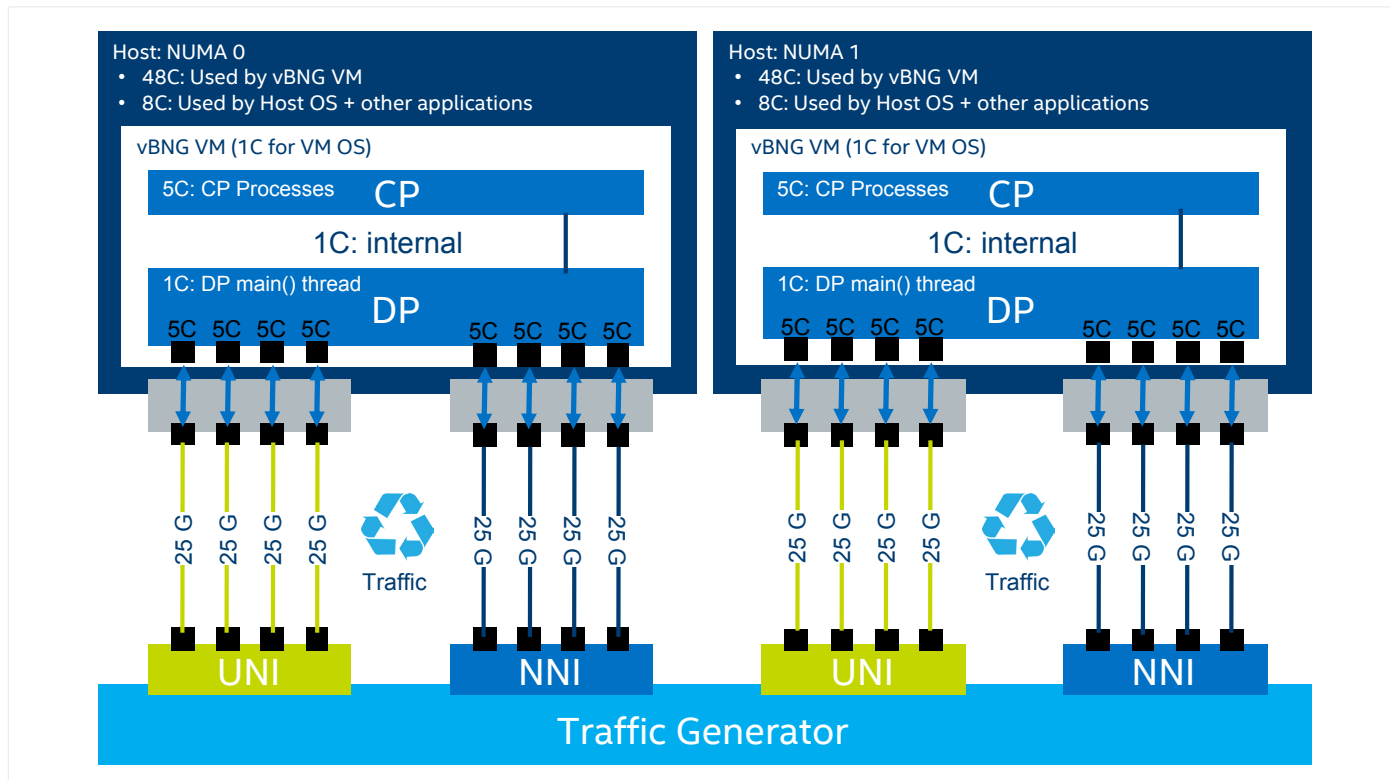


Figure 3. Test node block diagram

Test Results

Figure 4 shows total performance for four packet sizes (both in Gbps and millions of packets per second (Mpps) compared to the theoretical maximum throughput available from the

sixteen 25 GbE ports per server. The test results show that the netElastic vBNG can achieve between 270 Gbps and 380 Gbps of throughput depending upon packet size with 128,000 attached devices. All tests were conducted to meet the IEEE RFC2544 standard for 0.001% packet loss.




Figure 4. netElastic vBNG throughput in Gbps and Mpps (higher is better).¹

The test set up was designed to closely recreate a real-life deployment and to provide an assurance to CommSPs that they can use the system where they need at least 270 Gbps and 128,000 connected devices. As a significant portion of the subscriber traffic can be expected to carry streaming data, social media, and other large packet-based traffic, the actual throughput in a real-world network should be between 334 Gbps and 367 Gbps as the most common packet sizes will be between 384 bytes and 512 bytes long. The 512 byte packet size performance obtained in the tests is about 96% of the theoretical maximum, a significant throughput performance of over 360 Gbps.¹

The number of cores per 25 GbE port was an important configuration question during the tests. After experimenting with a wide range of combinations, the Intel and netElastic team determined that optimal performance for this server came from using five logical cores per 25 GbE port.

Conclusion

With internet traffic growing, the need for more bandwidth has never been greater, which makes the vBNG a critical piece of network infrastructure. With that in mind, the goal of this testing was to document the performance netElastic vBNG software can achieve running on the Intel® Select Solution for NFVI Forwarding Platform. The test results verified that innovations from Intel and netElastic have enabled them to significantly exceed their previous vBNG throughput levels, which will allow CommSPs to scale their networks to accommodate the tremendous growth in internet traffic.² The test results achieved both new levels of throughput and numbers of connected users that can be scaled up by using additional servers. The results also showcase the performance density per server that can be accomplished with the Intel Select Solution for NFVI Forwarding Platform and the netElastic vBNG solution.

Learn More

netElastic: www.netelastic.com

Intel Select Solutions for Network: builders.intel.com/intelselectsolutions/network

Red Hat: www.redhat.com



Notices & Disclaimers

¹ See Tables 1 and 2 for configuration details. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. Refer to <http://software.intel.com/en-us/articles/optimization-notice> for more information regarding performance and optimization choices in Intel software products.

² For more complete information about performance and benchmark results, visit <http://www.netelastic.com/wp-content/uploads/2019/11/White-Paper-High-Performance-BNG-on-Intel-NFVI-Platform.pdf>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See Tables 1 and 2 for configuration details. No product or component can be absolutely secure.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

0720/DRO/H09/PDF

Please Recycle

343027-002US