

New CPU Generation Improves Speed and Cost-Effectiveness for vBNG

Tests of virtual broadband network gateways (vBNG) using 3rd generation Intel® Xeon® Scalable processors show up to 20% better performance and up to 56% improvement in performance per dollar processor spending compared to the previous CPU generation.



Increasing broadband internet demand is driving communications service providers (CoSPs) to expand and virtualize their installed base of broadband network gateways (BNGs). BNGs provide user line aggregation, policy enforcement, and packet forwarding functionality. Throughput and cost-effectiveness are virtual BNG metrics that are important to CoSPs because these systems need to handle significant data volumes. netElastic is a pioneering vBNG leader that has tested its software on servers using 3rd generation Intel® Xeon® Scalable processors and recorded a dramatic increase in performance and performance per dollar.



Meeting Growing Broadband Demand

Demand for broadband networking is growing, both in terms of increases in data flows and numbers of new users, as trends toward work from home and remote learning accelerate. CoSPs originally built their access networks around city hubs or business areas. This network design then expanded to residential areas as demand grew. The rapid shift of the past year has changed that dynamic. Now CoSPs are working to deliver more fiber-to-the-home (FTTH) services at high data rates.

CoSPs are working to meet the needs of current customers, understanding bandwidth needs will only continue to increase, while also working to scale the network to meet the widescale, rapid growth in getting more people connected. 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.

Intel and netElastic have joined together to demonstrate vBNG cost effectiveness and performance increases when moving to a general purpose server that uses the newest 3rd generation Intel Xeon Scalable processor. This testing is designed to provide some of the proof points as CoSPs consider adopting the new CPUs as they grow their networks to stay ahead of consumer demand and look for ways to reduce operating expenses.

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 CoSP network, including establishing and managing subscriber sessions and providing services such as:

- Authentication, authorization, and accounting (AAA) for each session
- IPv4 and IPv6 address assignment
- Carrier-grade network address translation
- Access control lists
- Policy management
- Quality of service
- BGP, OSPF, and IS-IS routing

Table of Contents

Meeting Growing Broadband Demand	1
netElastic Virtual BNG	1
Test Topology	3
Test Results	5
Packet Throughput Performance	5
Future Ready CPU Features	6
Conclusion	6

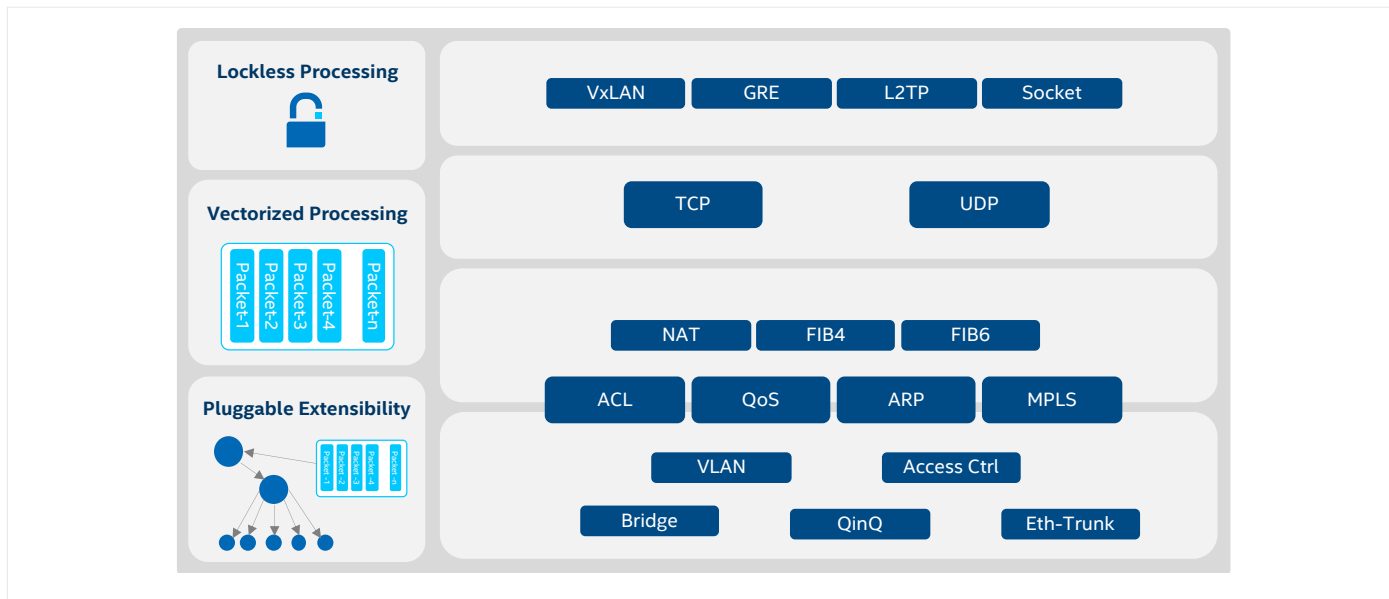


Figure 1. netElastic vBNG architecture.

To support real-time packet processing and high data throughput, netElastic designed the vBNG with the following features (see also Figure 1) 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 its own separate hosts or virtual machines. This architecture allows the data plane to be scaled independently, providing optimal deployment flexibility. The vBNG’s SDN architecture also improves network flexibility and simplifies infrastructure management with software-based control.
- Data Plane Development Kit (DPDK): netElastic vBNG uses DPDK combined with netElastic’s optimized code in the data plane to route packets around the operating system and improve packet 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.

netElastic’s approach enables the vBNG to provide CoSPs with optimal 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. Service providers can also avoid large upfront expenses with netElastic’s “pay-as-you-grow” licensing options.

3rd Generation Intel Xeon Scalable Processors

- **Flexibility from the edge to the cloud**, bringing AI everywhere with a balanced architecture, built-in acceleration, and hardware-based security.
- **Part of a complete set of network technology from Intel**, including accelerators, Ethernet adapters, Intel® Optane™ persistent memory, FlexRAN, OpenNESS, Open Visual Cloud, and Intel® Smart Edge.
- **Engineered for modern network workloads**, targeting low latency, high throughput, deterministic performance, and high performance per watt.
- **Enhanced built-in crypto-acceleration** to reduce the performance impact of full data encryption and increase the performance of encryption-intensive workloads.
- **Hardware-based security** using Intel® Software Guard Extensions (Intel® SGX),¹ enhanced crypto processing acceleration,¹ and Intel® Total Memory Encryption.¹



	DUT 1	DUT 2 ²
PROCESSOR	2X Intel® Xeon® Gold 6238R processors (28 cores, 2.2 GHz, 165 W TDP)	2X Intel Xeon Gold 6330N processors (28 cores, 2.2 GHz, 165 W TDP)
MEMORY	384 GB DDR4 @2933MTs	256 GB DDR4 @2666MTs
NETWORK ADAPTERS	4x Quad-port 25 Gbps Intel® Ethernet Network Adapter XXV710	4x Quad-port 25 Gbps Intel® Ethernet Network Adapter XXV710
MICROCODE	0x5003003	0xd000270
INTEL® HYPER-THREADING TECHNOLOGY	On	On
INTEL® TURBO BOOST TECHNOLOGY	Off	Off
BIOS	Dell Inc. Version: 2.2.10	SE5C6200.86B.0020.P22.2103231313
SYSTEM DDR MEM CONFIG: SLOTS / CAP / SPEED	12x 32 GB 2933 MTs DDR4 DIMM	16 X 32GB 2667MHz RAM
# NODES	1	1
# SOCKETS	2	2
CORES/SOCKET, THREADS/SOCKET	28/56	28/56

Table 1. DUT server configuration.

Test Topology

netElastic worked with Intel to benchmark the performance improvement of running its vBNG on 3rd generation Intel Xeon Scalable processors by testing that performance alongside that of a server based on 2nd generation Intel Xeon Scalable processors. The devices under test (DUT) included:

- DUT 1: Server utilizing dual Intel Xeon Gold 6238R processors that are part of the 2nd generation Intel Xeon Scalable processor product family.

- DUT 2: Server based on dual Intel Xeon Gold 6330N processors that are part of the 3rd generation Intel Xeon Scalable processor product family.

The testing was done by Intel in February and March 2021. Table 1 shows the server configuration and Table 2 shows the software stack used.

The object of the test was to highlight the generational improvements of the Intel Xeon Gold 6330N processor. Both servers used Intel® Ethernet 700 Series Network Adapters, for DPDK-accelerated networking, and Intel® Solid State Drives (Intel® SSDs).

OS	RED HAT 8.2
KERNEL	4.18.0-240.10.1.el8_3.x86_64
WORKLOAD	netElastic vBNG-xeon 1.3.20.B18
COMPILER	gcc (GCC) 8.3.1 20191121 (Red Hat 8.3.1-5)
INTEL® ETHERNET NETWORK ADAPTERS XXV710 DRIVER	2.8.10-k
INTEL® ETHERNET NETWORK ADAPTERS XXV710 FIRMWARE	Firmware-version 6.02
QEMU/LIBVIRT	libvirt 6.0.0/QEMU 4.2.0
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 PROFILE³	iommu=pt intel_iommu=on,eth_no_rmrr vfio_iommu_type1.allow_unsafe_interrupts=1 pcie_acs_override=downstream selinux=0 enforcing=0 usbcore.autosuspend=-1 nmi_watchdog=0 default_hugepagesz=1G hugepagesz=1G hugepages=160

Table 2. DUT software configuration.




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

The memory in DUT 2 was 256 GB compared to 384 GB in DUT 1. This difference provided a small disadvantage to DUT 2, but the exact amount of this disadvantage was not calculated.

The overall test topology is shown in Figure 2 and Figure 3. Both DUT 1 and DUT 2 featured two vBNG virtual machines each loaded into their own non-uniform memory access (NUMA) zone. NUMA is a technology that configures the CPU and memory in a way that allows VMs or containers to share local memory for improved performance. Each NUMA zone utilized 24 cores (48 logical cores per socket were used in the tests) of the processors in DUT 1 and DUT 2.

In addition, 80 GB memory was used on both DUTs 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;

however in this testing, only 64,000 sessions per vBNG were utilized 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 each device under test. This means that each server had compute headroom to add future services and currently allocated for OS operations for this testing.

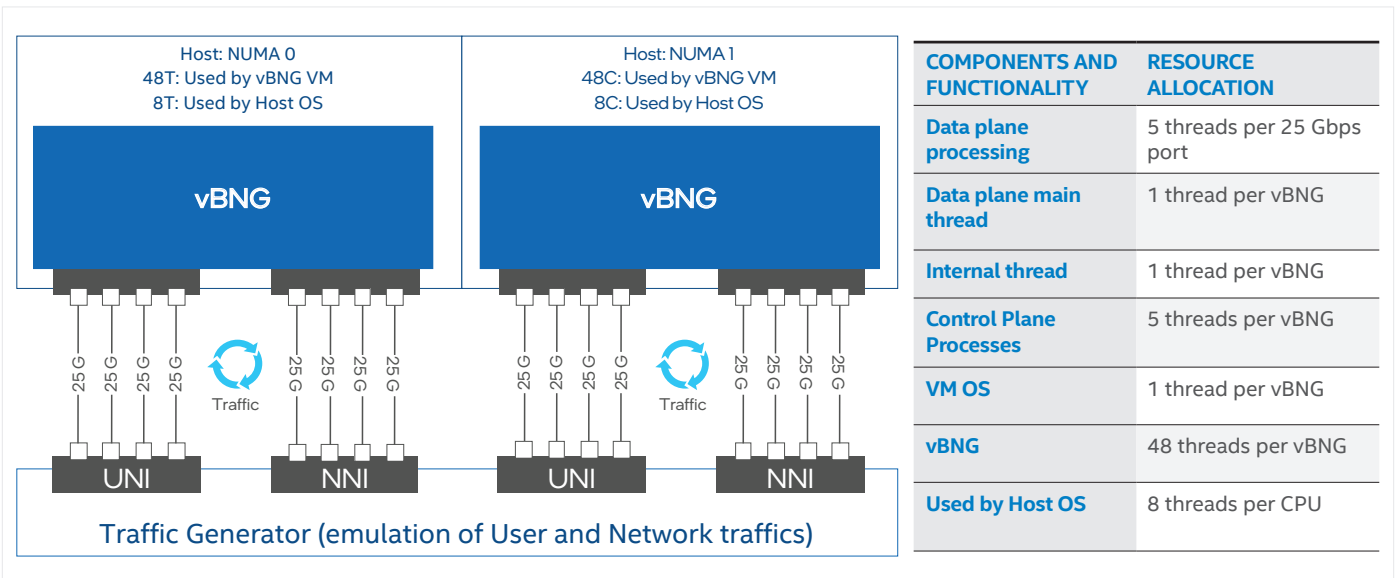


Figure 3. Block diagram of DUT test configuration with data flows.

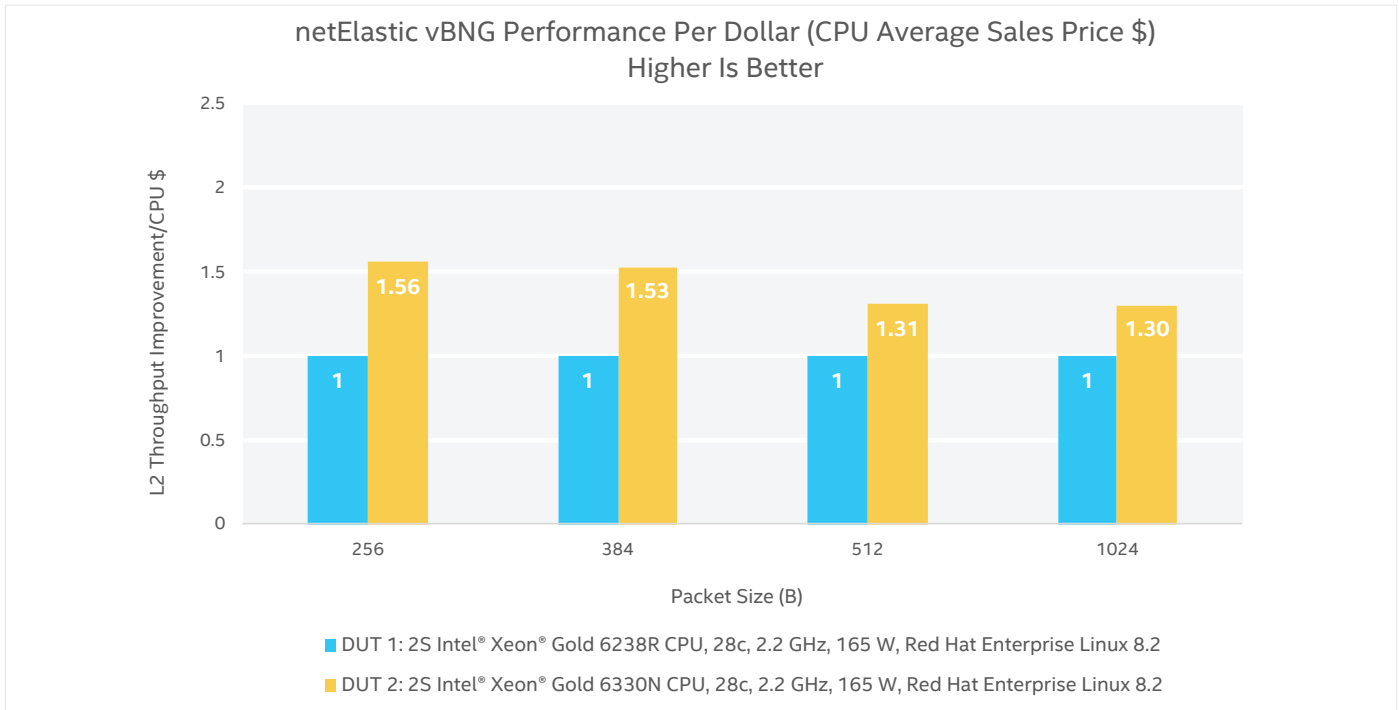


Figure 4. Comparison of 0.001% packet loss performance per dollar of DUT 1 and DUT 2.⁴

Test Results

The test results show an up to 56% increase in compute performance per dollar and up to a 20% increase in packet throughput performance.

Figure 4 shows the performance of each CPU relative to its cost based on average sales price (ASP). The tests showed compute increases per dollar across all packet sizes with a more than 56% gain in performance per dollar at a 256 byte packet size, which is the most stressful test case in terms of millions of packets per second.

Packet Throughput Performance

Figure 5 shows total layer 2 forwarding performance for four packet sizes in Gbps from DUT 1 and DUT 2 as well as the theoretical maximum throughput available from the sixteen 25 GbE ports per server. The test results show that using DUT 1 the netElastic vBNG achieved between 219.3 Gbps and 381.4 Gbps of throughput depending upon packet size with 128,000 attached devices. Using DUT 2, the netElastic vBNG achieved between 264.9 Gbps and 383.3 Gbps of throughput depending upon packet size with 128,000 attached devices. All tests were conducted in conformance with the IEEE RFC2544 standard for 0.001% packet loss.




Figure 5. netElastic vBNG 0.001% packet loss throughput in Gbps.

The improvement in throughput was 20% at 256-byte packet size and 18% at 384-byte packet size. At higher packet sizes, the performance increases were much smaller, which was expected since the performance was already very close to theoretical maximum performance.

The test setup was designed to closely re-create a real-life deployment and to provide an assurance to CoSPs that they can use the system where they need at least 268 Gbps and up to 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 1,024-byte packet size performance obtained in the tests is about 98% of the theoretical maximum, a significant throughput performance of over 383 Gbps achieved in one server.

Future Ready CPU Features

In addition to the improvements uncovered in the test results, moving to 3rd generation Intel Xeon Scalable processors provides added features that make these systems future ready. Support for faster I/O is important because it is a significant contributor to packet throughput. The tests in this paper utilized PCI Gen 3.0 interfaces, but newer Intel® Ethernet adapters can be used with 3rd generation Intel Xeon Scalable processors, leveraging PCI Gen 4.0 interfaces with throughput up to 100 Gbps per port. Servers using these processors can utilize up to 52/57 TB of physical/virtual addressability (PA/VA) memory, an improvement over the 46/48 PA/VA from the prior CPU generation.

Conclusion

To meet increasing demand for broadband internet services, CoSPs are utilizing virtual BNGs to deploy lower cost systems that are flexible and scalable. Compute performance and compute performance per dollar are important criteria in the decision to invest in a new system. The tests described in this paper show that by moving to servers powered by 3rd generation Intel Xeon Scalable processors, vBNGs can deliver better cost effectiveness and overall improvement in packet throughput. The 3rd generation Intel Xeon Scalable processor family has the processing power for today's compute-intensive applications and flexible performance to scale as network requirements evolve. By embracing these new CPUs, netElastic is offering its customers a high performance vBNG solution that meets the feature and performance needs of CoSPs.

Learn More

[netElastic](#)

[Intel® Network Builders](#)

[3rd generation Intel Xeon processors](#)



Notices & Disclaimers

¹ This technology is not supported when using Intel Optane persistent memory.

² Preproduction System

³ isolcpus, rcu_nocbs, nohz_full, kthread_cpus, irqaffinity are also part of the boot profile depending on the CPU set in use.

⁴ Based on pricing as of March 30, 2021, from <https://ark.intel.com/content/www/us/en/ark/products/199345/intel-xeon-gold-6238r-processor-38-5m-cache-2-20-ghz.html?wapkw=6238R> and <https://ark.intel.com/content/www/us/en/ark/products/212286/intel-xeon-gold-6330n-processor-42m-cache-2-20-ghz.html>. Prices were \$2,625.00 for Intel Xeon Gold 6238R processor and \$2,029.00 for Intel Xeon Gold 6330N processor.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

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

Your costs and results may vary.

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

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© 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.