

# IMPLEMENTING A HIGH-PERFORMANCE BNG WITH INTEL® UNIVERSAL NFVI PACKET FORWARDING PLATFORM TECHNOLOGY

**netElastic**  
systems

## Authors

**Kerry Zhang**  
netElastic

**David Williams**  
netElastic

**Rich Sabin**  
netElastic

**Eoin Walsh**  
Intel

**Paul Mannion**  
Intel

**Andrew Duignan**  
Intel

**Ma Jianwei**  
Intel

**Yan Feng**  
Intel

The broadband network gateway (BNG) is located at the edge of a metropolitan area network (MAN), giving users access to various broadband and data services through a fixed network. It is also the point where carriers and communications service providers (CoSPs) can control network access behaviors. With the growth of software-defined networking (SDN) and network functions virtualization (NFV) in communication networks, the virtual BNG (vBNG) will be an important network element (node) in the next generation central office (NGCO), which hosts various edge access gateway nodes and provides access to local services.

The vBNG helps reduce network operator deployment costs by facilitating a transition from expensive, fixed-function ASIC-based hardware devices to low-cost, commercial off-the-shelf (COTS) servers. This white paper describes the outstanding performance of vBNGs implemented by netElastic\* vBNG on an Intel® Xeon® Platinum 8276L processor with the Intel® Ethernet Network Adapter XXV710-DA2 standard network interface card (NIC), featuring dynamic device personalization (DDP). The test results show that in full-service mode, a single 25 gigabit Ethernet (GbE) port on the server forwards 256 byte packets at near wire speed, while the entire server achieves packet forwarding throughput rates up to 300 Gbps for a typical traffic profile.

To verify the performance of the vBNGs, the test environment employs an OpenStack\* platform and an Ixia\* simulation device to construct a network functions virtualization infrastructure (NFVI) test framework for broadband access services. The DDP feature increases system performance by offloading data traffic processing rules through profiles, isolating and hashing BNG protocols and data packets, and alleviating typical packet processing bottlenecks by efficiently allocating data plane packets to processor cores.

The Intel® select edge universal data plane platform can be used to enable this vBNG performance as it is based on the Intel® Xeon® processor family and is optimized for the edge of the operator network.

## Table of Contents

1.0 Solution Brief.....	3
1.1 OpenStack* Platform .....	3
1.2 netElastic* vBNG Overview .....	3
1.3 Tester (Ixia*) .....	3
2.0 OpenStack* Platform.....	4
2.1 Hardware Configuration .....	4
3.0 Network Topology .....	5
4.0 Intel® DDP Feature .....	5
5.0 Benchmark Performance Testing .....	6
5.1 Service Traffic Model .....	6
5.2 Performance Data .....	7
5.2.1 Performance Comparison of Upstream Traffic over PPPoE via a Single 25GbE Port 8.....	7
5.2.2 Forwarding Performance of 16 x 25GbE Ports for the Entire Server .....	7
6.0 Intel® Select Edge Data Plane Platform .....	8
7.0 Conclusion .....	9
Appendix A: BIOS Configurations .....	9

# 1 Solution Brief

## 1.1. OpenStack\* Platform

OpenStack is an open source platform that uses pooled resources to build and manage private and public clouds. OpenStack's architecture consists of many open source projects, and there are stable core services that handle compute, networking, storage, identity, and images. These six core services form the infrastructure that allows the rest of the projects to handle dashboarding, orchestration, bare-metal provisioning, messaging, containers, and governance.

The core services include:

- Nova\*: On-demand access to compute resources
- Neutron\*: Network connectivity as-a-service
- Swift\*: Object storage service
- Cinder\*: Block storage service
- Keystone\*: Identity service for authentication and authorization
- Glance\*: Image service to discover, register, and retrieve virtual machine

## 1.2. netElastic\* vBNG Overview

The Broadband Network Gateway (BNG) is the access point for subscribers to connect to the broadband network. BNG aggregates traffic from subscriber sessions from an access network, and routes it to the service provider's network. In addition to providing routing functions, the BNG also communicates with the authentication, authorization, and accounting server to perform session management and billing functions.

The netElastic vBNG is designed to deliver outstanding performance and scalability on commercially-available hardware. netElastic's optimized data plane and patented packet processing techniques maximize hardware performance, and its decoupling of the control plane and data plane allows for flexible capacity expansion without restrictions from either plane.

Whether network operators are deploying a new rural network or upgrading a large-scale metro POP, the vBNG provides the deployment flexibility to deliver new services faster for a wide range of scenarios, ranging from a very small number of subscribers up to millions of subscribers.

The netElastic vBNG is software-based and runs on white-box servers. This approach and industry-leading scalability can help service providers lower solution cost by 70 percent compared to traditional BNGs.<sup>1</sup>

netElastic also offers a comprehensive, web-based BNG Manager that provides operational visibility, management, and control from a single pane of glass. BNG Manager monitors all interfaces, routes, peak and active subscribers, and overall system health.

For more information, please visit <https://www.netelastic.com/>.

## 1.3. Tester (Ixia\*)

The Ixia IxNetworks-XGS2 testing solution shown in Figure 1 was used to simulate the control plane and data plane services when broadband users access the carrier network.

The test server is a dual socket system, in which 8 x 25GbE ports are connected over PCIe\* links to non-uniform memory access (NUMA) node 0 and 8 x 25GbE ports connected to NUMA node 1. There are 16 x 25GbE ports on the server.

With 8 x 25GbE ports connected to each CPU socket, the tester simulates the point-to-point protocol over Ethernet (PPPoE) user logons and logoffs. It uses 8,000 PPPoE clients per port and generates user upstream traffic. Meanwhile, with 8 x 25GbE ports, the tester also simulates the network side. It distributes open shortest path first (OSPF) routes to the vBNG device, and each port establishes a peer relationship with the vBNG device to distribute 25,000 routes over the OSPF network segment and simulate user downstream traffic.

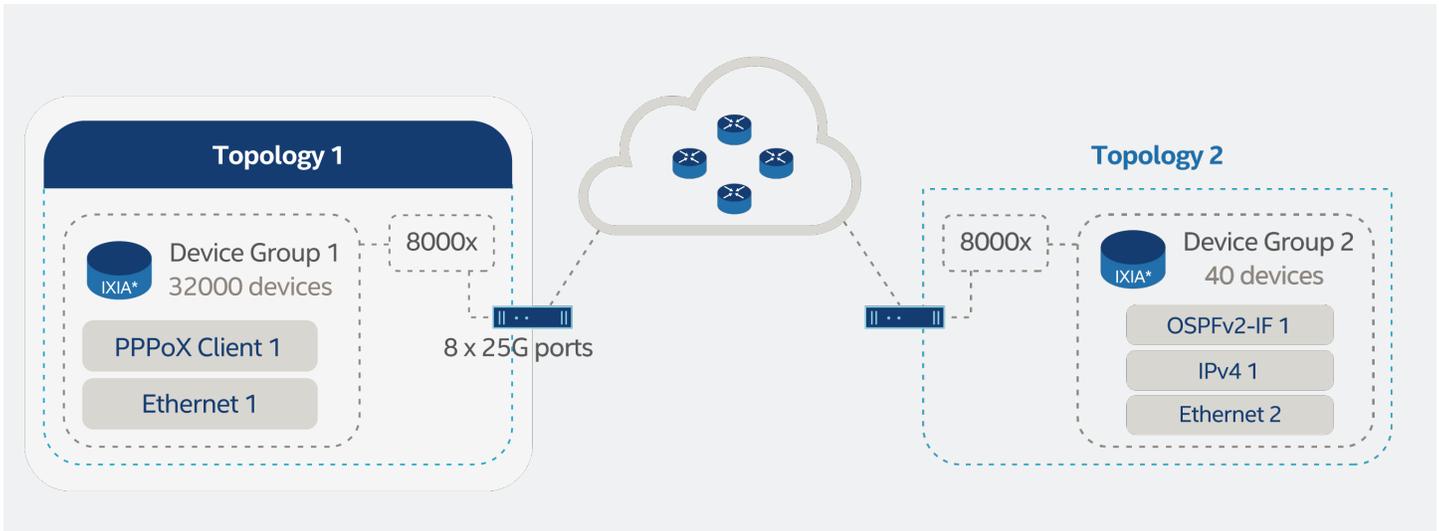


Figure 1: Configuration of the Ixia\* Tester

## 2 System Components - Software and Hardware

This section provides a detailed overview of the hardware and software components used for performance testing of the Intel Xeon processor-based server.

### 2.1 Hardware Configuration

Packet throughput performance was tested on 19th June 2019 and measured on an Inspur\* server platform NF5280M5 (<https://www.inspursystems.com/product/nf5280m5>) with two CPU sockets. Each CPU socket ran a vBNG instance and was connected to four Intel® Ethernet Network Adapters seated in PCIe 3.0 x 8 slots. Each NIC has two 25GbE ports. The detailed hardware and software components are shown in Tables 1 and 2.

Server	Inspur* Server Platform NF5280M5
CPU	Intel® Xeon® Platinum 8276L processor
Number of CPUs	2
Number of cores per CPU	28 (24 cores for data plane forwarding)
Memory	DDR4 @ 2400 MHz, 112 GB (64 GB used)
Network interface card (NIC)	Intel® Ethernet Network Adapter XXV710-DA2 with DDP enabled
Microcode	0x4000017

Table 1. List of Hardware Components

vBNG	netElastic* vBNG
Guest OS	CentOS* 7.5.1804
OpenStack*	OpenStack* Pike

Table 2.: List of Software Components

### 3 Network Topology

The test environment, shown in Figure 2, used OpenStack as the cloud platform test environment, consisting of three servers as the basic cloud environment. The vBNG ran on an OpenStack Nova (Compute) service and was connected with the Ixia IxNetworks-XGS2 through the Intel Ethernet Network Adapter.

The deployment and network settings of OpenStack are shown in Figure 3.

### 4 Intel® DDP Feature

Dynamic device personalization (DDP) is an advanced feature of the Intel Ethernet Network Adapter XXV710-DA2. As a general-purpose, hyper-converged standard NIC, it provides partial programmability, and with the downloadable firmware profile, helps parse various communication network protocols and steers the packets directly into the correct vBNG core. Combined with flow director (FDir) and hash (RSS) functions, the NIC enables hardware offloading of the packet parsing and distribution functions, thus dramatically improving network performance.

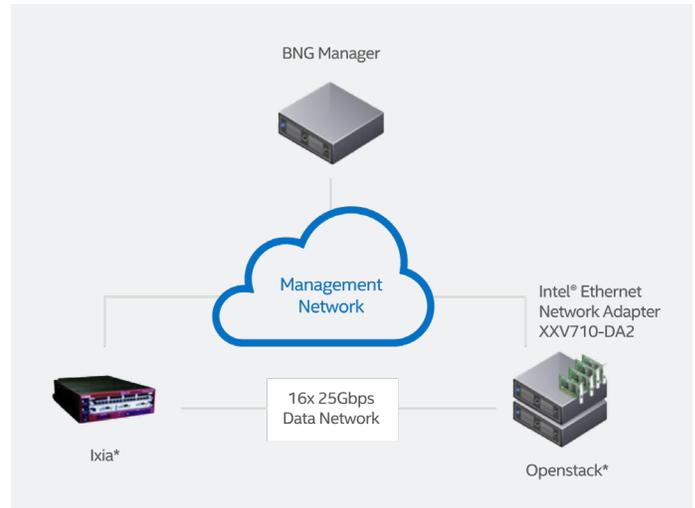


Figure 2: Topology of the Test Environment

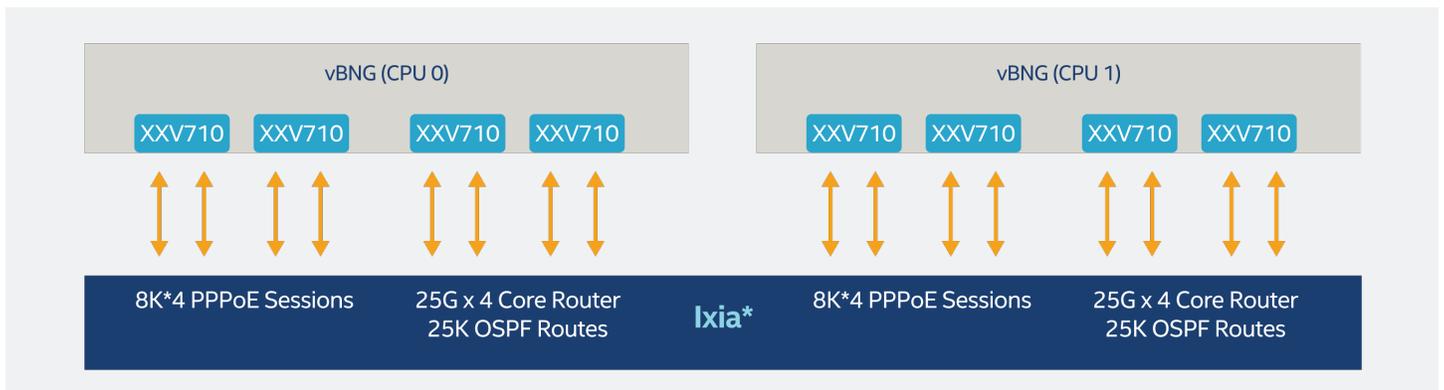


Figure 3: Schematic Diagram of Test Platform Configuration

- Improved Programmability

Packet processing can be customized to meet specific service needs. For example, by downloading different profiles (e.g., QinQ), the NIC can support both broadband network gateways and wireless core gateways.

- Optimized Performance

The identification and diversion of network packets is made possible by the use of hardware. For instance, the hardware allows for the separation of control packets from protocol packets and hashes data packets for traffic classification using multiple queues of the NIC. In addition, PPPoE control plane packets can be easily filtered at the NIC and sent on a different path to the data plane packets, thus greatly enhancing the possibilities for Control and User Plane Separation (CUPS)-enabled systems. DDP effectively reduces packet processing delays, increases throughput, and improves the utilization of multi-core processors.

- Rapid Deployment

Intel provides industry-tested DDP profiles, which support various protocols such as PPPoE, GTP-C, GTP-U, and L2TP. These profiles can be easily downloaded through generic Ethtool or Data Plane Development Kit (DPDK) drivers.

• Examples

In general, generic NICs can only identify a certain number of protocols and cannot identify packets using some communication protocols, such as PPPoE. If there is no DDP feature, the inner layers of PPPoE packets will not be parsed. The NIC can only determine PPPoE packets are Layer 2 packets by default. With DDP enabled, the NIC can classify deeper into the packet to the transport layer (Figure 4), and the aforementioned FDir and RSS functions can be used to enable hardware offloading to improve performance.

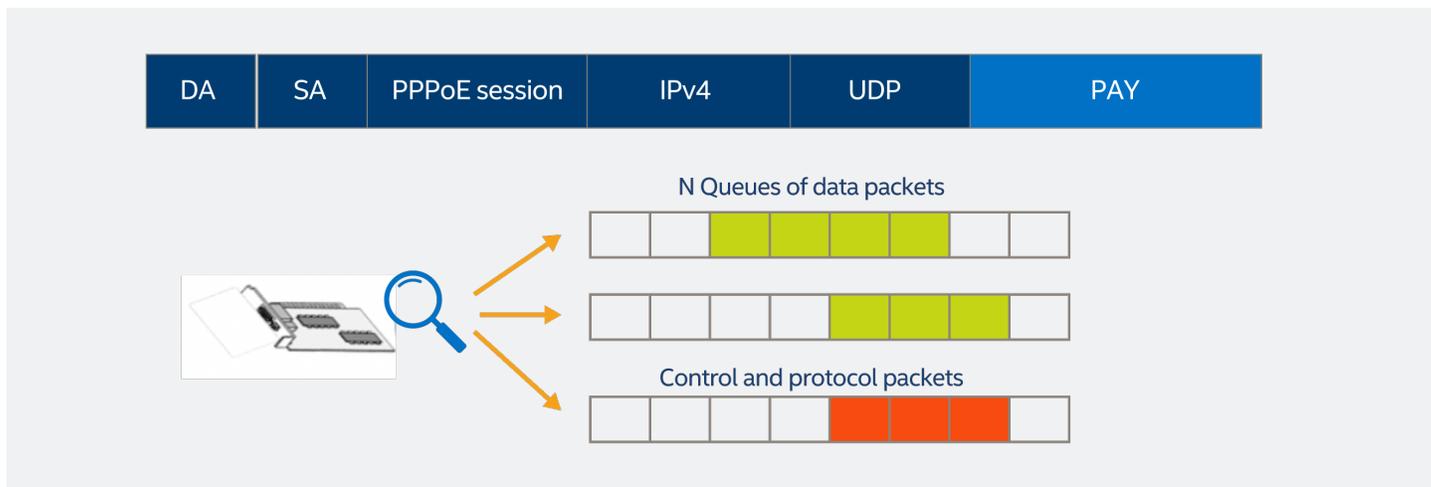


Figure 4: Packet Processing Mode with the DDP Feature

## 5 Benchmark Performance Testing

The vBNG is deployed at the near edge to enable the broadband network access and ensure quality of service (QoS) for user access on top of the virtualized architecture. The netElastic vBNG was tested using a typical traffic profile model to measure its performance in real-world scenarios.

### 5.1 Service Traffic Model

Data plane packet performance was tested using a service traffic model similar to a real network configuration, with a specific number of user sessions, access control lists (ACLs), QoS, and routing rules to simulate an actual networking environment.

The server had 8 x 25GbE ports used for user access (UNI) and 8 x 25GbE ports used for the network (NNI). There were 8,000 users per UNI port, and each port was full-duplex.

For QoS testing, there was one committed access rate (CAR) defined per user, which was applied to both upload and download speeds, as shown in Table 3.

Item	Entries/Number	Comments
User Sessions	8,000 users/Access Port	Up to 32,000 users per port are supported.
ACLs	4,000/template	Four templates used
Routes	50,000 routes	Global for all users
Uplink QoS	8,000/port	One CAR entry per user
Downlink QoS	8,000/port	One CAR entry per user

Table 3: Configuration List of Service Model

## 5.2 Performance Data

The benchmark performance data is characterized in two parts:

1. The upstream traffic performance via a single 25GbE port
2. The overall server performance for a full 8 port configuration

The test used a 256 byte packet size.

### 5.2.1 Performance Comparison of Upstream Traffic Over PPPoE via a Single 25GbE Port

In total, there were four uplink scenarios in this benchmark test:

- Scenario 1: 32,000 PPPoE users and 50,000 routes were configured on a single 25GbE uplink port.
- Scenario 2: 4,000 ACL rule were enabled.
- Scenario 3: User QoS policy was added to the configuration.
- Scenario 4: The 4,000 ACL rule and the user QoS policy were both used in addition to the standard pipeline.

The performance results for these scenarios with DPP enabled (w/DPP) vs. disabled (w/o DPP) are shown in Figure 5.

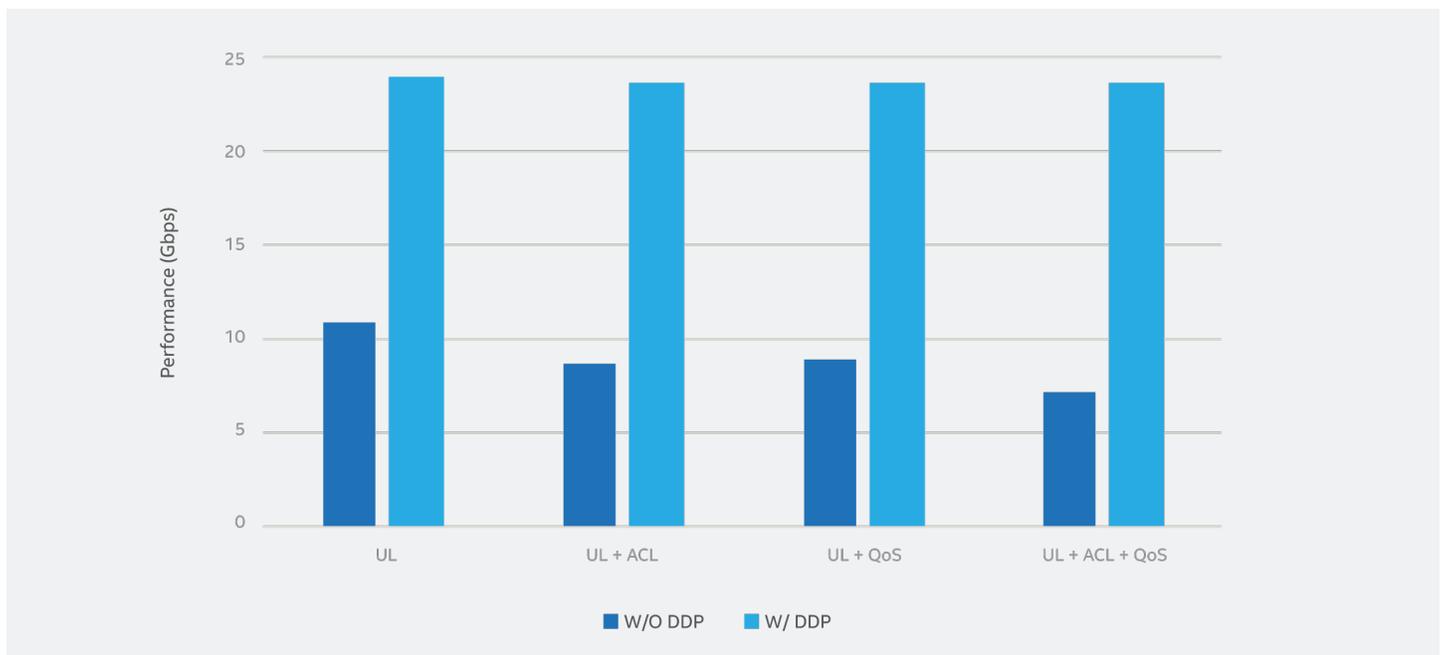


Figure 5: Comparison of Upstream Traffic Throughput via a Single 25GbE Port (256 Byte Packets )

The test results show a single core in the test platform processes significantly less traffic than the 25GbE port can handle when the DPP is disabled, resulting in underutilized port capacity. In addition, even when CPU resources are available, additional CPU cores cannot be utilized if DPP is disabled. When the DPP is enabled, the DPP feature can divert data packets (received at a 25GbE port) to multiple CPU cores for parallel processing, thereby maximizing the port bandwidth and CPU resource utilization. With the help of DPP, the forwarding performance for upstream traffic improves significantly to be nearly wire speed, even when ACL rules and QoS policy are enabled.

### 5.2.2 Forwarding Performance of 16 x 25GbE Ports for the Entire Server

The objective of this test was to verify the data plane forwarding capacity of the entire server, using eight Intel Ethernet Network Adapter XXV710 DA-2. The test results, based on the four scenarios defined in the prior section, are shown in Figure 6.

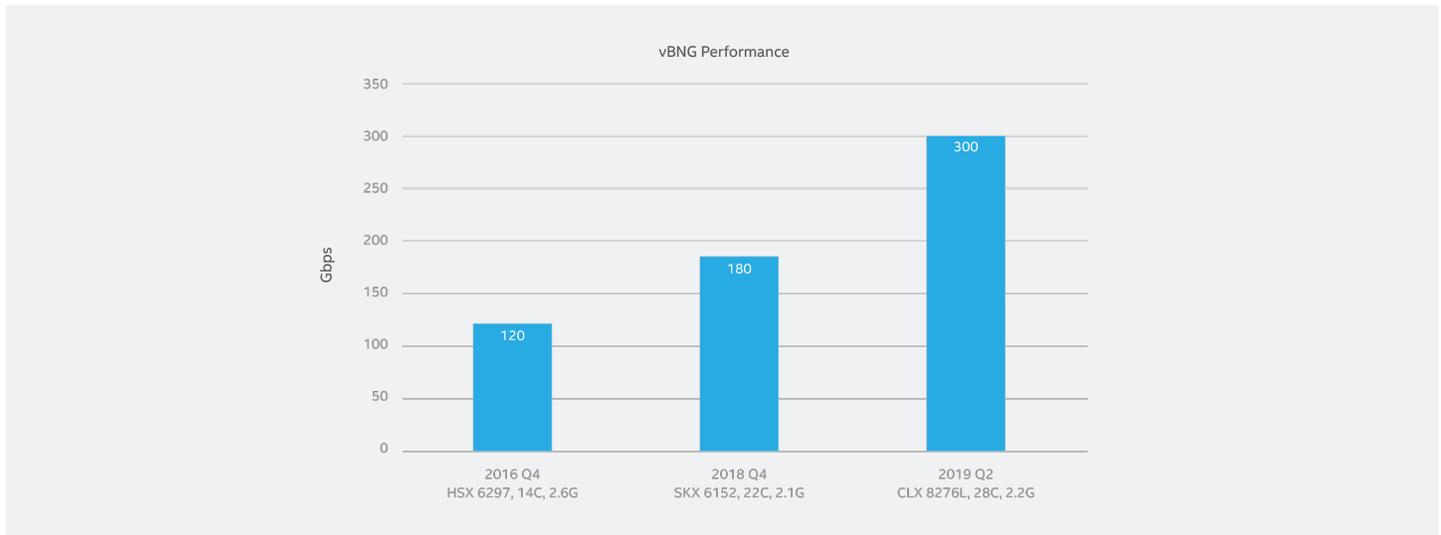


Figure 6: Server Performance of Three Generations of vBNGs

With the advanced Intel Ethernet Network Adapter XXV710 DA-2 DDP capability, the server can forward packets up to a rate of 300 Gbps with QoS policy and ACL rules enabled.

The use of DDP removes the need to use one CPU core as a load balancer, which could prove to be a performance bottleneck, and therefore allows system performance to scale linearly with the number of cores (as seen in Figure 6). The CPU cores previously used for load balancing are freed to be used for compute workloads, rather than infrastructure work.

## 6 Intel® Select Edge Data Plane Platform

Intel select edge universal data plane platform was defined in collaboration with CoSPs and ecosystem partners to unlock the value of an I/O-balanced platform architecture which maximizes network I/O throughput per NUMA node on a standard server. This platform is a customized NFVI solution for 4G/5G core user plane functions that require high packet forwarding throughput, such as vBNGs and virtual cable modem termination systems (vCMTs). This multi-service platform is well-suited for edge deployment scenarios.

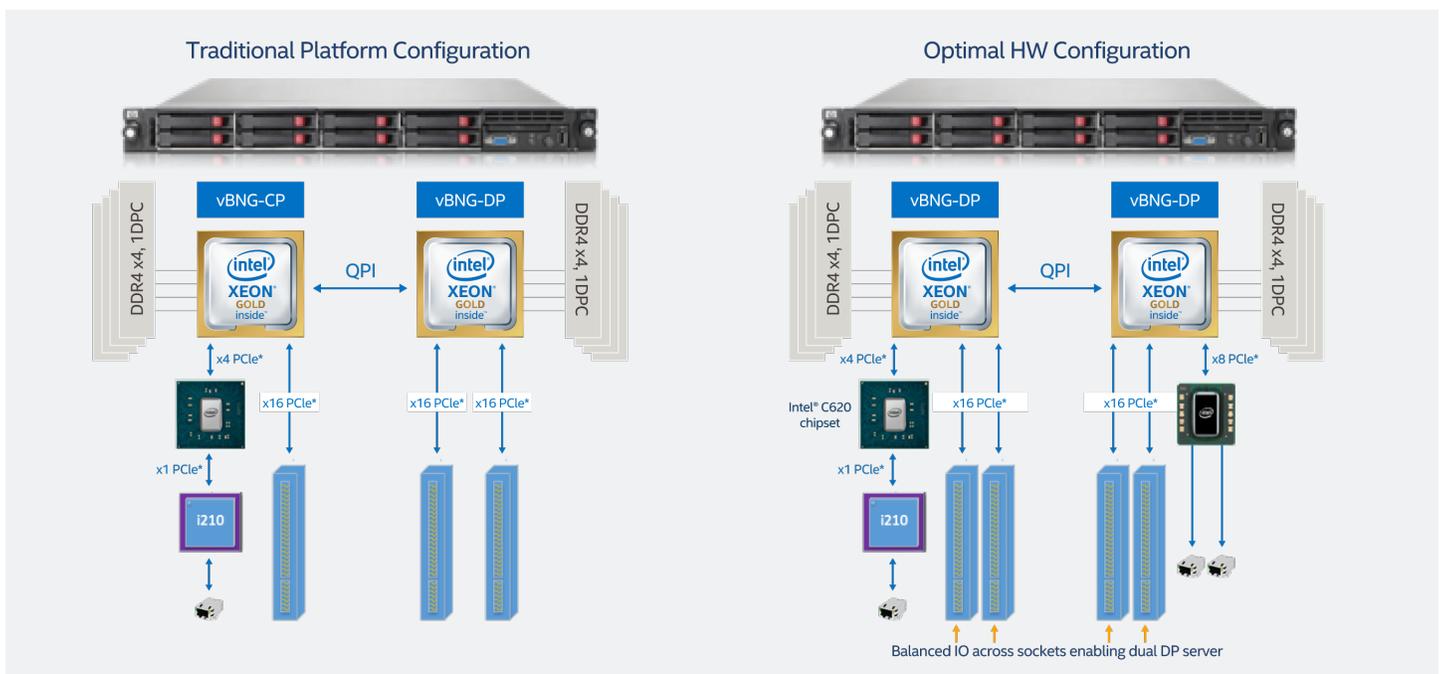


Figure 7. Traditional and Optimal Intel® Hardware Configuration

The optimized Intel select edge universal data plane platform shown in Figure 7 overcomes the deficiencies of an asymmetric design of the traditional platform configuration, which does not maximize I/O capacity. Balanced I/O across sockets enables the implementation of CUPS, which can boost data plane performance for many high I/O workloads, such as vBNG, vCMTS, etc.

## 7 Conclusion

netElastic and Intel have worked to maximize vBNG performance and offer seamless scalability on Intel® processors used in standard OEM and telecom equipment manufacturer (TEM) server designs. Running on servers with optimized NIC and CPU configurations, vBNGs can reduce the cost-per-bit and improve return on investment (ROI), helping CoSPs cost effectively transition to NFVI at the edge.

## Appendix A: BIOS Configurations

<b>Processor Configuration</b>	
Intel® Hyper-Threading Technology	Disabled
Intel® Virtualization Technology	Enabled
MLC Streamer	Enabled
MLC Spatial Prefetcher	Enabled
DCU Data Prefetcher	Enabled
DCU Instruction Prefetcher	Enabled
Direct Cache Access (DCA)	Enabled
<b>Power and Performance</b>	
CPU Power and Performance Policy	Performance
<b>Uncore Power Management</b>	
Uncore Frequency Scaling	Enabled
Performance P-limit	Enabled
<b>CPU P State Control</b>	
Enhanced Intel SpeedStep® Technology	Enabled
<b>CPU C State Control</b>	
CPU C-State	C0/C1
<b>QPI Configuration</b>	
Intel® QPI Frequency Select	Auto Max
<b>Memory Configuration</b>	
Enforce Recommended Memory Speed	Enabled
Memory Operating Speed Selection	2133

<b>Memory RAS and Performance Configuration</b>	
Memory RAS Configuration	Maximum Performance
NUMA Optimized	Enabled
Cluster-on-Die	Disabled
Early Snoop	Enabled
Patrol Scrub	Enabled
Correctable Error Threshold	10
Memory Correctable Error Enabling	Enabled
<b>Integrated IO Configuration</b>	
Intel® VT for Directed I/O	Enabled

1. <https://www.netelastic.com/index.php/products/vbng>.



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](http://www.intel.com/benchmarks).

Performance results are based on testing as of 19th June 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at [\[intel.com\]](http://intel.com).

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

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

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. © Intel Corporation