



# INTEL SELECT SOLUTION NFVI FORWARDING PLATFORM

Eoin Walsh

NFV Solution Architect

# Streamlining Network transformation

Optimize and scale infrastructure planning, procurement, and deployment to meet the Packet Processing challenges of the Cloud and 5G

## CONSISTENT CHALLENGES ACROSS ALL NETWORK TRANSFORMATION PATHS:

- Can you get enough data plane performance out of a COTS platform?
- Can you reduce the complexity of NFVI solutions to ease deployment?

**From the foundation of the proven Intel® Select Solutions for NFVI, Intel® Select Solutions for NFVI Forwarding Platform delivers optimized, tested configurations for tough and widely used packet processing workloads**



4G LTE & 5G UPF



BNG & CMTS



IPsecGW



## INTEL® SELECT SOLUTIONS FOR NFVI FORWARDING PLATFORM:

- Optimized IO performance for data plane packet processing
- Verified implementations from trusted vendors

## BENEFITS:

- Simplifies NFVI procurement, validation, and deployment
- Accelerates time-to-market

# What's Different: Industry challenges



MEET AND BEAT  
5G DATA GROWTH



PROVIDE COMMON  
PLATFORM FOR SERVICES



SUPPORT DIVERSE  
USE CASES

EXISTING INFRASTRUCTURE

DEVELOPMENT COSTS

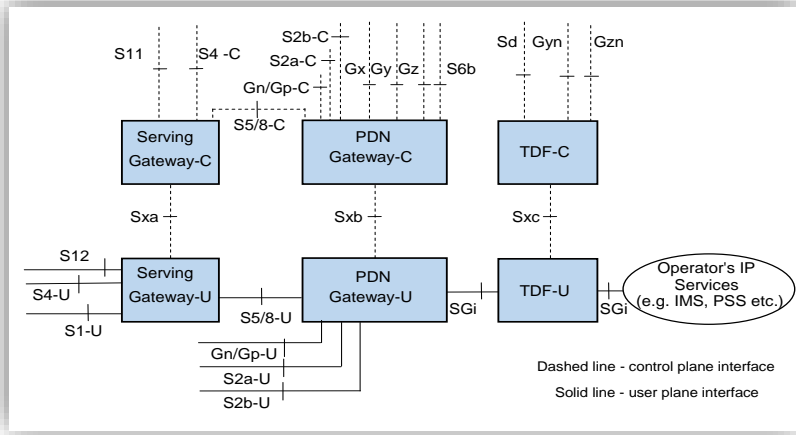
INVESTMENT SCALABILITY

DELIVER TCO EFFICIENCY

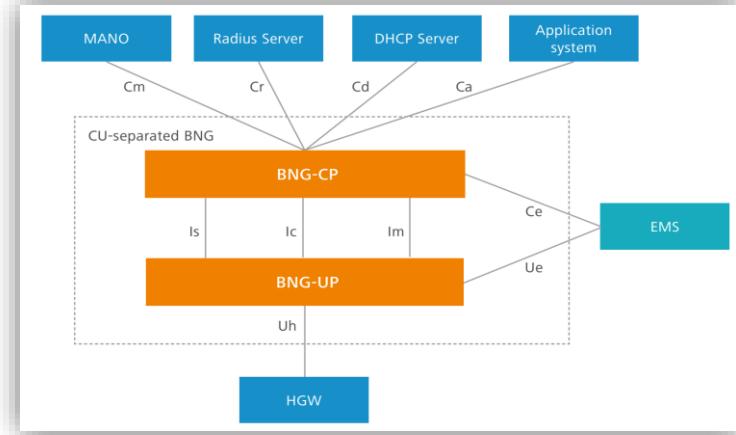
Source: Cisco CNI report

THE GOAL IS TO MAKE IT EASIER FOR NETWORK TRANSFORMATION  
TO SUPPORT BUSINESS NEEDS AND OPPORTUNITIES

# CUPS changing how we look at NFVI



3GPP: EPC/5GCN




BBF: BNG/PE

- 3GPP Rel.14 CUPS defined a disaggregated EPC with separated user plane. 5GCN carries this on
- Fixed line functions like PE & BNG are also adopting a CUPS paradigm (BBF)
- This allows a disaggregation of the node into different software and silicon for control and user plane

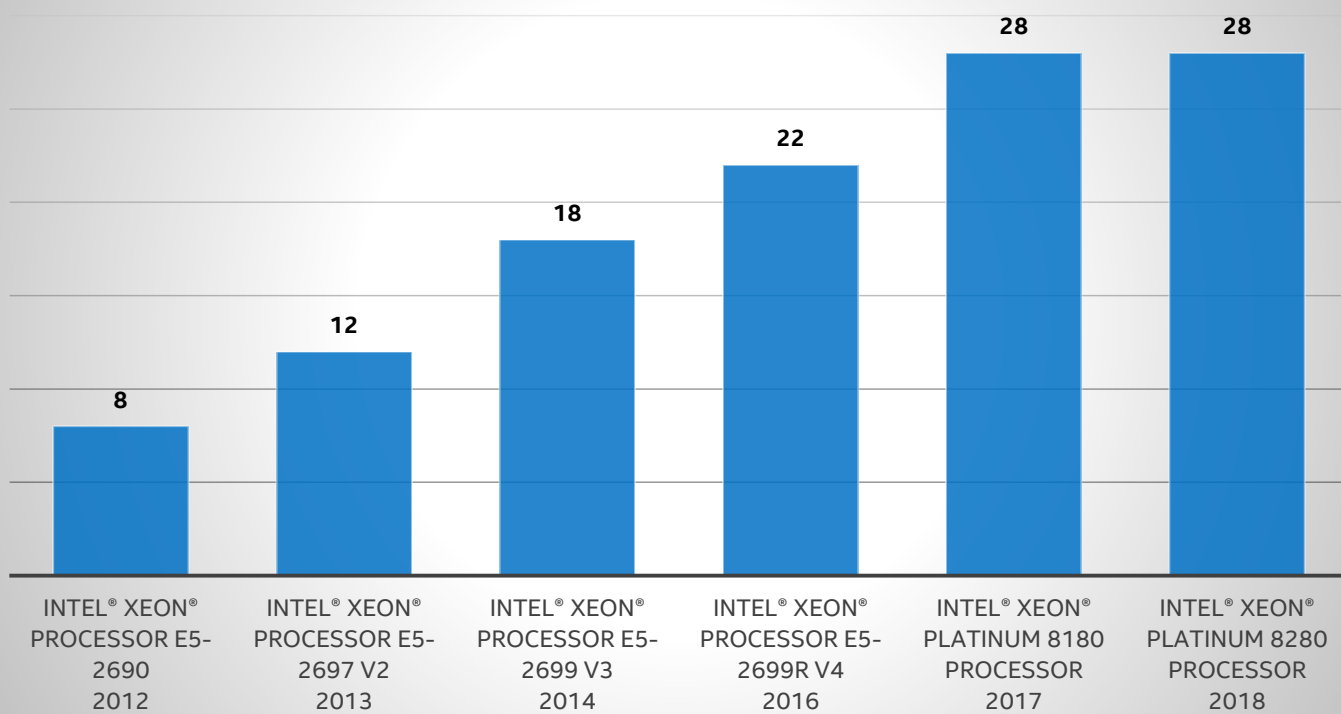
# Comms Use Case Example

## Tradition



Traditionally, Comms Control Plane co-exists with Data Plane Processing. This is non-optimal configuration, as Control Plane functionality requires less throughput processing, and the server is architected asymmetrically, which leaves unconsumed CPU headroom.

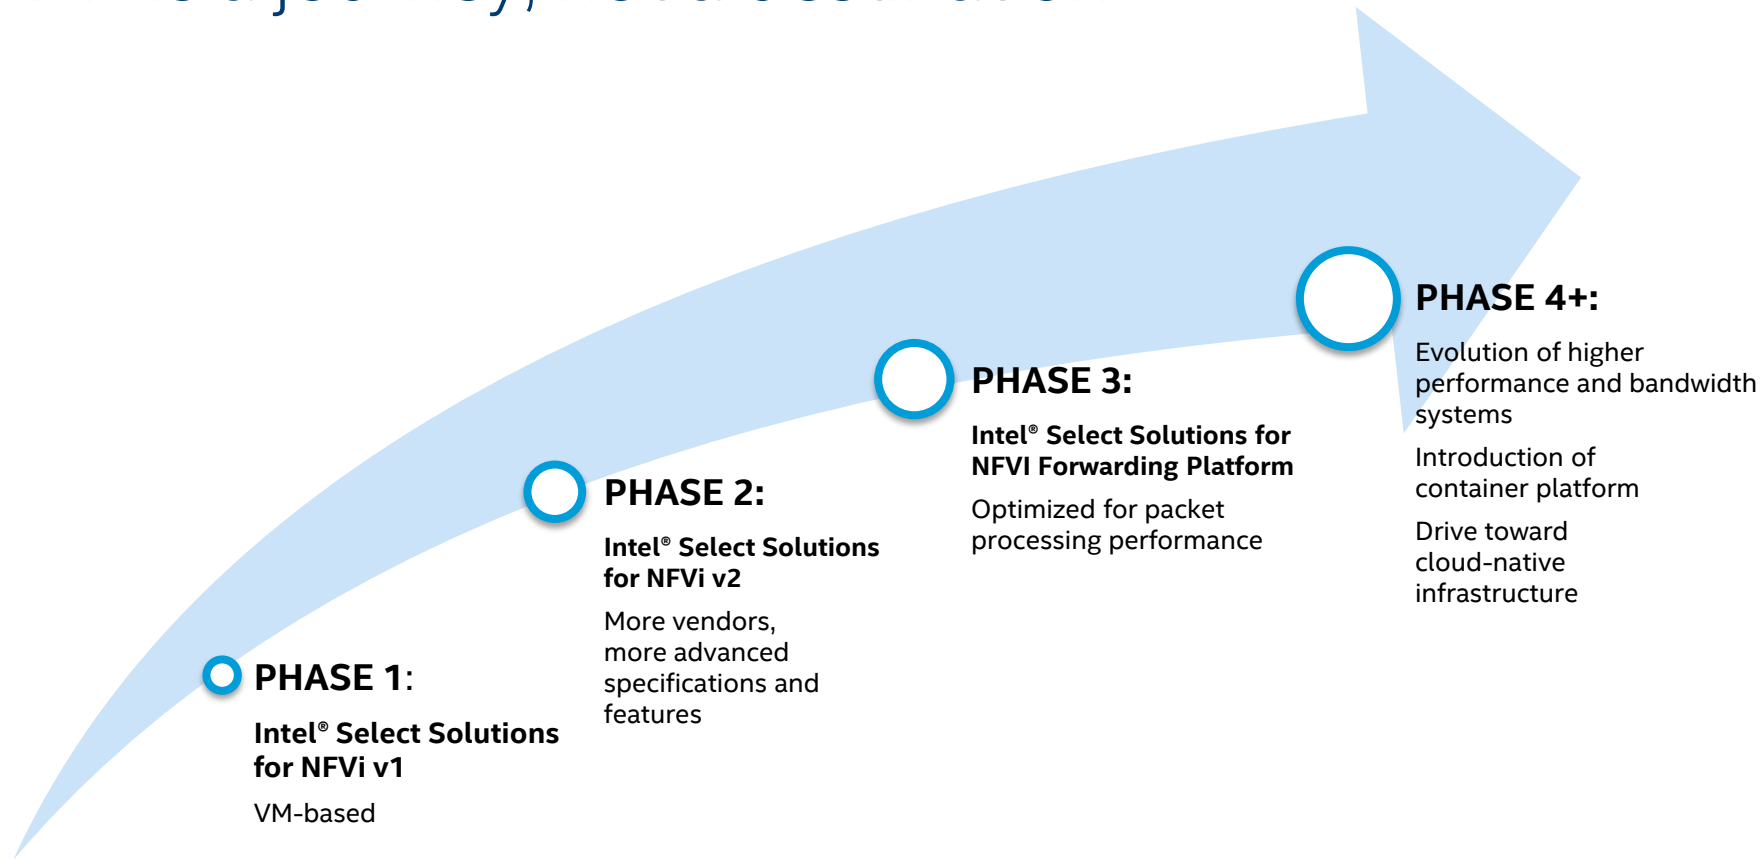
## Top bin Cores per Generation



Balanced IO across sockets enabling dual DP server to maximize Data Plane processing.

Use Cases  
C  
IG  
MTS  
NAT  
ryption

# NFVI is a journey, not a destination



# Intel® Select Solution for NFVI v2 & NFVI Forwarding Platform are Variations on a Theme

## Intel® Select Solutions for NFVI v2 Configuration



General NFVI compute platform

## Network Traffic




Intel® Processor




Intel® Network Adapters

## Data Storage



Intel® Optane™ DC Persistent Memory



Intel® SSD with NVMe\*



Intel® SSD

## Crypto/Compression Acceleration



Intel® QAT

## I/O OPTIMIZED FOR MULTI-SERVICES PLATFORM

## Intel® Select Solutions for NFVI Forwarding Platform Configuration



Optimized for Packet Processing

## Max IO – High Speed Packet Processing



Intel® Processor



Intel® Network Adapters



Intel® SSD

### OPTIONAL

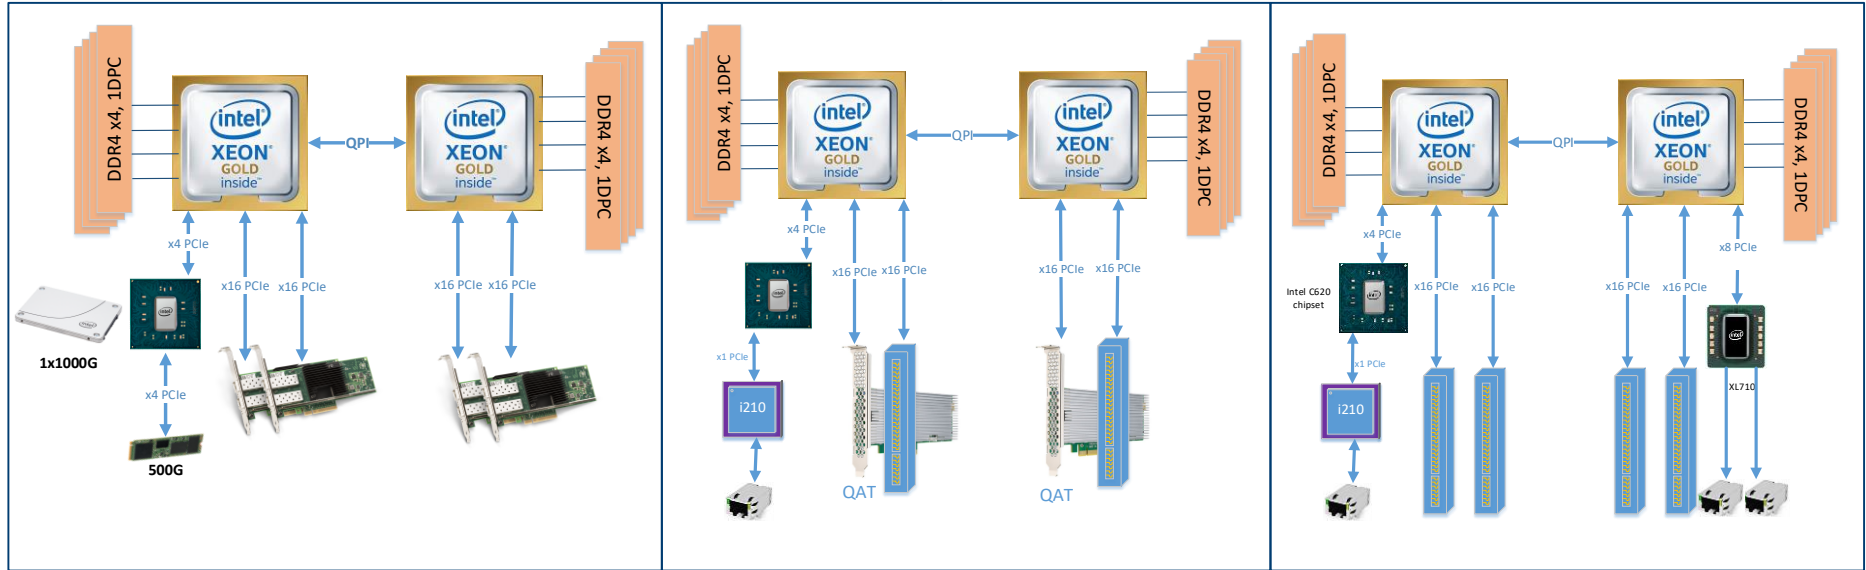
- Intel® Optane™ DC Persistent Memory
- Intel® SSD with NVMe
- Intel® QAT

# Example GW Configurations

5G UPF

CMTS/IPsec GW

BNG/PE

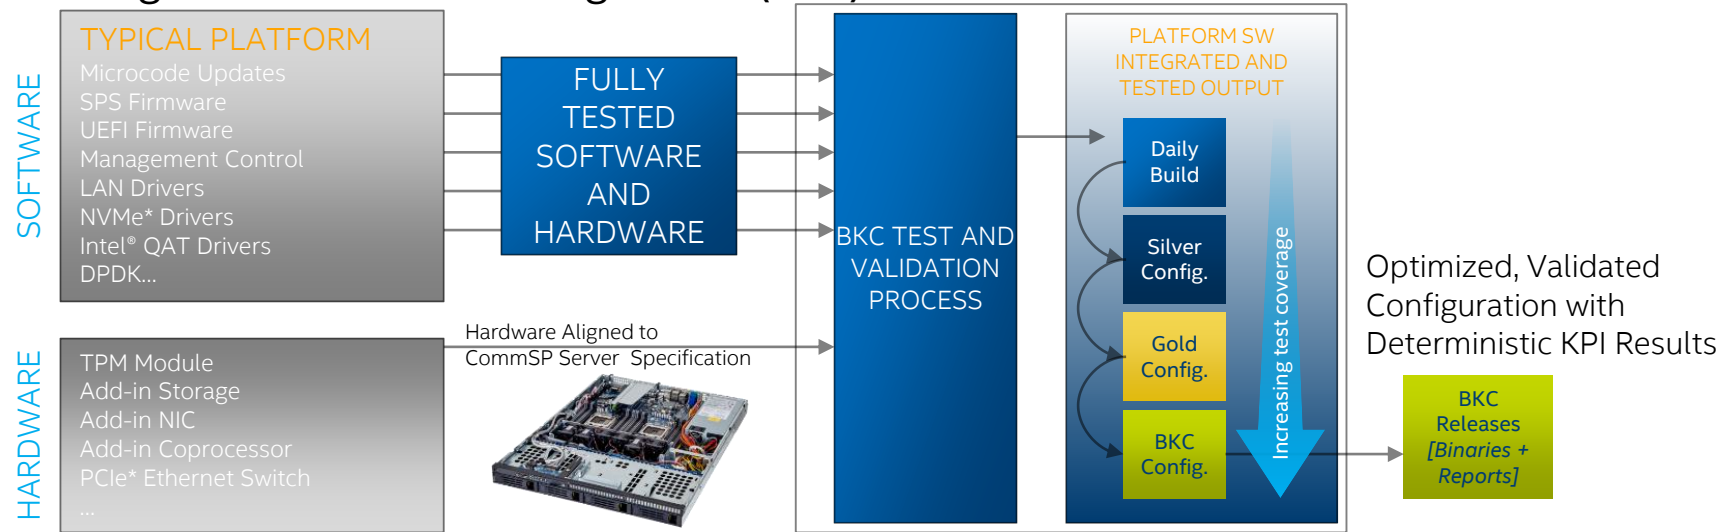


Telco Gateway UP are all variations on a common platform Architecture



# The Process Behind an Optimized and Well-Engineered Foundation

## Defining the Best Known Configuration (BKC)



PROCESS RESULTS ENSURE CONSISTENCY AND STABILITY TO SIMPLIFY EVALUATION, ORDERING, AND DEPLOYMENT

# Intel® Select Solutions for NFVI

## Forwarding Platform with Red Hat Enterprise Linux 7.6 and Red Hat OpenStack Platform 13 on 2<sup>nd</sup> Gen Intel® Xeon® Scalable Processors

## HARDWARE BILL OF MATERIALS (BOM)

Ingredient	Controller Node Requirement	Base Requirement	Plus Requirement
CPU	Intel® Xeon® Gold 5218 CPU @ 2.3 GHz, 125 W or 5218N CPU @ 2.3 GHz, 16C/32T, 105 W (or higher SKU)	Intel® Xeon® Gold 6230 CPU @ 2.1 GHz or 6230N CPU @ 2.3 GHz, 20C/40T, 125 W (or higher SKU)	Intel® Xeon® Gold 6252 CPU @ 2.1GHz, 24C/48T, 150W (or higher SKU)
Memory	192 GB – Required 2667 MHz 384 GB – Recommended	384 GB – Required	384 GB – Required
Intel® Optane™ DC Persistent Memory	Not Required	2-1-1 Topology 128 GB DIMM 2LM mode (512 GB Total) - Recommended	2-2-1- 128 GB 2LM – Recommended 2-2-2- 128 GB 2LM – Recommended
NIC	2x 25GbE Intel® XXV710 or Intel® E810-CQDA2	2x Quad port 25GbE Intel® XXV710 or 4x Dual port 25GbE Intel® XXV710 or 2x Intel® E810-CQDA2	4x Quad-port 25GbE Intel® XXV710 or 4x Intel® E810-CQDA2
LOM	Required	Required	Required
Storage	Required	Required	Required

## KEY PERFORMANCE METRICS

### Performance Baseline

- MLC
- Jitter
- Power

### Network with DPDK

- L3FWD RFC2544 zero Packet loss
- OVS-DPDK

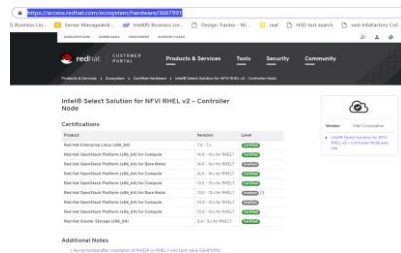
### vBNG

- DL/UL Throughput (MPPs)
- DL/UL cycle per packet cost

### Cyclictest - Latency

## RED HAT CERTIFICATION

1. <https://access.redhat.com/ecosystem/hardware/3667991>
2. <https://access.redhat.com/ecosystem/hardware/3667951>
3. <https://access.redhat.com/ecosystem/hardware/3668051>



## SOFTWARE STACK

Operating System, Software and Drivers	Red Hat Enterprise Linux* v7.6
Host OS	kernel 3.10.0-933.el7.x86_64
Libvirt	4.5.0 <sup>1</sup>
QEMU	2.12.0 <sup>1</sup>
QAT	L4.4
CollectD	5.8 <sup>1</sup>
NVMe	1.0
DPDK <sup>2</sup>	18.11
i40e	2.3.2-k
i40evf	3.2.2-k
ixgbe	5.1.0-k-rh7.6
ixgbev	4.1.0-k-rh7.6
Ice <sup>3</sup>	TBD
avf (i40evf) <sup>3</sup>	TBD

## PLATFORM TECHNOLOGIES

Platform Technologies	Enable/Disable
Intel® VT	Intel® CPU VMX Support Enable
	Intel® IO Virtualization Enable
Intel® Boot Guard	Intel Boot Guard Enable for Base Enable <sup>3</sup> for Plus
Intel® TXT	Intel® Trusted Execution Technology Enable <sup>3</sup>

## BIOS

BIOS	Enable/Disable
BIOS Profile	Max Performance
Turbo	Disable
C-states	
Uncore Frequency Scaling	Disable

<sup>1</sup>Available as part of Red Hat OpenStack Version 13

<sup>2</sup>Available as extra packages

<sup>3</sup> Launch plan target RHEL7.7/RHOSP13 and plan to refresh to RHEL8.1 w RHOSP16

# Key benchmark requirements

	KEY KPI TARGETS	BASE CONFIGURATION	PLUS CONFIGURATION
Red Hat Hardware Certification	Red Hat Enterprise Linux* 7 Hardware Certification and Red Hat OpenStack* 13 Certification	Required	Required
HW/SW BOM	HW/SW BOM requirement	Required	Required
Performance Baseline	MLC (local)	Latency <90ns BW >100000MB/s	Latency <90ns BW >100000MB/s
	MLC (remote)	Latency <150ns BW >34000 (remote)	Latency <150ns BW >34000 (remote)
	Jitter	<10000	<10000
Network Performance Baseline	DPDK L3FWD with RFC2544 0.001% packet loss	90% line rate at 256B packet size	90% line rate at 256 B packet size
Latency	Cylicttest (24hr) without stress	1 CPU : <5 µs (min), <10 µs (avg) 8 CPU : <10 µs (min), <15 µs (avg)	1 CPU : <5 µs (min), <10 µs (avg) 8 CPU : <10 µs (min), <15 µs (avg)
vBNG v19.03 Workload	Overall server power	<430 W	<650 W
	Per-CPU power	<100 W per socket	<160 W per socket
	Overall vBNG throughput RFC2544 0.001% packet loss	150 Gbps / server	310 Gbps / server
	Store Forward Latency (Avg)	50 µs	50 µs

# Legal Disclaimers

- All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.
- 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
- 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].
- 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 disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.
- Intel, the Intel Logo and other Intel Marks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.
- Other names and brands may be claimed as the property of others.
- Copyright © Intel Corporation. All rights reserved.

