

Solution Brief

Communications Service Provider



Intel-HPE Verified Reference Configuration (VRC) for Network Function Virtualization Infrastructure (NFVI) and Secure Access Service Edge (SASE) on Red Hat* OpenShift Container Platform



Authors

Intel®

Jonathan Tsai, Ai Bee Lim, Timothy Miskell, Vaishnavi Saravanan

HPE

Table of Contents

Introduction	1
Intel® VRC for NFVI and SASE.....	1
Partner Solution	2
Solution Software.....	3
Performance.....	3
Summary	6

Introduction

Intel Verified Reference Configuration (VRC) is a family of workload-optimized infrastructure solutions based on the Intel® Xeon® Scalable processor family, which targets today's complex workloads. This document describes a reference implementation for the 5th Gen Intel® Xeon® Scalable Processor, which includes a selection of HPE ProLiant Compute DL380 Gen11 servers and Red Hat* OpenShift Container Platform.

When network operators, service providers, cloud service providers, or enterprise infrastructure companies choose an Intel Verified Reference Configuration (VRC) for Network Function Virtualization Infrastructure (NFVI) and Secure Access Service Edge (SASE) deployment based upon on 5th Generation Intel® Xeon® Scalable Processor, they will be able to deploy various network-function-virtualized applications more securely and easily than ever before.

The reference implementation helps end users simplify design choices by bundling hardware and software pieces together, making high performance more predictable. End users will spend less time, effort, and expense evaluating hardware and software options. The solution is architected to deliver consistent performance for workloads based on in-house benchmarks and best-known configurations and methodologies.

This document features a workload-optimized stack tuned to fully utilize the Intel® architecture foundation.

Intel® VRC for NFVI and SASE

Intel Verified Reference Configurations for NFVI and SASE are based on a multi-node architecture that consists of a controller, a cloud node, and storage at a minimum. The Intel Verified Reference Configuration for NFVI and SASE provides all required resources to implement a software-defined infrastructure that resides within each cloud server instance and is controlled by the hypervisor. The Controller Node is intended to be used for control, signaling, and management by implementing the Virtual/Container Network Function (VNF) Management (VNFM) and Virtualization Infrastructure Management (VIM). Thus, it may not require additional local storage and hardware acceleration.

NOTE: Refer to <https://networkbuilders.intel.com/solutionslibrary/intel-vrc-for-network-function-virtualization-infrastructure-nfvi-v4-and-sase-on-red-hat-openshift-container-platform-rhocp> for solution specification details.

Partner Solution

HPE ProLiant Compute DL380 Gen11

The HPE ProLiant Compute DL380 Gen11 Server with 5th Gen Intel® Xeon® Scalable processors is engineered to unlock more value from your data and drive innovation across the edge to the cloud. Designed for today's hybrid computing environments, this future-ready server delivers:

- Optimized performance for your workloads
- Intuitive cloud operating experience
- Trusted security by design
- Extended protection across the HPE partner ecosystem. Added protection leverages Security Protocols and Data Models (SPDM) for component authentication, with trust established through HPE iLO 6 verification.



Figure 1: HPE ProLiant Compute Compute DL380 Gen11 Server

HPE Hardware and Firmware Details

Details for the NFVI solution are based on the hardware configuration, which showcases the best combination of the latest Intel® CPU technology coupled with Intel® platform technologies, Intel® Ethernet, and Intel® acceleration technologies. These technologies are integrated into the motherboard to deliver best-in-class performance with low latency requirements using the Data Plane Development Kit (DPDK).

Table 1 shows the Platform Hardware Configuration and Table 2 provides the Platform Firmware details.

Table 1: Platform Hardware Configuration

Hardware	Description
Processors	5 th Gen Intel® Xeon® Scalable Processor Two Sockets Intel® 6538N Processors
DRAM	16x 32GB Dual Rank DDR5 4800MHz 1 DIMM per channel Total Memory 512GB
Network Interface Card	2x Dual Port 100GbE Intel® Ethernet Network Adapter E810-2CQDA2 (Gen 4 x16)

Hardware	Description
Storage	2x 960GB SSD NVMe solution as boot device

Table 2: Platform Firmware

System	Components	Version
ProLiant DL380 Gen11	BIOS	2.2
	iLO 6	1.54

Notes:

1. For HPE ProLiant Compute DL380 Gen11 drivers and firmware updates go to:
https://support.hpe.com/connect/s/product?language=en_US&kmoid=1014696069&tab=driversAndSoftware

Intel® BIOS Recommendation

Intel® recommends using the BIOS Settings for deterministic performance with turbo enabled to meet performance requirements for the Packet Processing throughput workload. For detailed information on the BIOS settings recommendation, refer to the document BIOS Settings for Intel® Wireline, Cable, Wireless and Converged Access Platform (Document ID #747130) Chapter 3.0.

NOTE: Please contact your Intel® Field Representative for access to documentation.

HPE BIOS Config

This section describes the HPE BIOS Configuration required to achieve optimal performance with energy-balance and turbo enabled for high throughput packet processing. HPE recommends using the “Network Function Virtualization Infrastructure - Secure Access Service Edge” profile.

To set the “Network Function Virtualization Infrastructure - Secure Access Service Edge” profile, navigate to System Configuration->BIOS/Platform Configuration (RBSU)->Workload Profile and select the profile. Reboot the system after applying the workload profile. “Network Function Virtualization Infrastructure - Secure Access Service Edge” sets other BIOS settings necessary to meet energy-efficient turbo performance while maintaining high throughput packet processing performance.

Notes:

1. Reboot is required after applying the above settings for changes to take effect.
2. If it is set and the BIOS is upgraded to a newer version, follow the steps below for updates in the Telco Optimized Profile to take effect.
 - a. Load BIOS defaults, and
 - b. Re-apply “Network Function Virtualization Infrastructure - Secure Access Service Edge” profile

Solution Software

Solution deployment is on Red Hat* OpenShift 4.14 with the following Software versions:

Table 3: Solution Software Version

Ingredient	SW Version Details
Red Hat* OCP	4.14
RHCOS kernel	Kernel 5.14
Async mode NGINX*	v0.5.1
VPP IPSec	23.10
Intel®-Tensorflow	2.14.0 with oneDNN 3.4.1

Performance

This chapter aims to verify the performance metrics for the reference configuration for NFVI, ensuring no anomaly is seen. Refer to the information in this section to ensure that the performance baseline for the platform is as expected. For the full details of the testing methodologies of the workloads, refer to Intel Verified Reference Configuration for Network Function Virtualization Infrastructure (NFVI) v5 and Secure Access Service Edge (SASE) on Red Hat* OpenShift Container Platform (RHOCPS*) (Intel® Resource & Documentation Center (<https://www.intel.com/content/www/us/en/resources-documentation/developer.html#gs.afp7gr>) document #TBD)

Performance Baseline

This section includes information on a few applications that must be executed after the Platform is configured according to the BOM, BIOS configuration, and Software Stack, as described in the earlier sections.

The output provides a performance baseline on expected latency performance, memory bandwidth and jitter seen on the system.

Memory Latency Checker (MLC)

The first application is the Memory Latency Checker which can be downloaded from <https://software.intel.com/en-us/articles/intelr-memory-latency-checker>. Download the latest version and execute this application, unzip the tarball package and go into Linux* folder and execute /mlc.

Table 4: MLC Data

Idle Latencies for Sequential Access	NUMA Node 0
NUMA Node 0	93.4 ns

Read-Write Ratio	Peak Injection Memory Bandwidth (MBPS)
ALL Reads	411977.1
3:1 Reads-Writes	374517.7
2:1 Reads-Writes	377096.7
1:1 Reads-Writes	367674.9
Stream-triad like	355539.4

Inject Delay	Latency (ns)	Bandwidth (MBPS)
0	267.08	415736.3
2	269.44	416115.1
8	250.31	413530.1
15	251.61	409428.9
50	248.91	398511.2
100	173.73	319072.3
200	128.78	168139.9
300	131.89	114875.5
400	120.79	88612.6
500	109.17	71957.9
700	110.09	51940.4
1000	108.39	36961.2
1300	105.86	28577.3
1700	107.76	22039.3
2500	104.90	15222.5
3500	105.57	11138.0
5000	102.56	7962.0
9000	101.97	4784.1
20000	101.41	2521.4

Cache-to-cache Transfer Latency	Latency (ns)
Local Socket L2→L2 HIT Latency	59.1
Local Socket L2→L2 HITM Latency	60.4

NGINX* Software Benchmarks

Figure 2 shows the testing methodology using NGINX* to measure the connections per second that a server can sustain when handling 0-byte requests to the server.

Figure 2: Test Methodology for SSL with NGINX*





Figure 3 shows the performance results obtained for each configuration of the workload testing. The cipher ECDHE-X448-RSA4K with TLS 1.3 was used for the testing. Intel® QuickAssist Technology (QAT) hardware acceleration helps to offload public key exchange for SSL layer 7 applications.

Figure 3: NGINX* Software Results




With 4C8T, the Intel Verified Reference Configuration for NFVI should be able to demonstrate approximately 11026 Connections Per Second (CPS) with the full software stack of the NGINX* application with Intel® QAT HW (2 devices). Without Intel® QAT, the system should demonstrate approximately 1884 CPS with the out-of-box software stack of the NGINX* application. Overall, the full software stack with **Intel® QAT HW (2 devices)** shows approximately a **5.8x** gain over the out-of-box software stack and approximately a **3.7x** gain over the Optimized SW configuration.

VPP IPsec Software Benchmarks

Figure 4 shows the testing methodology using VPP IPsec. The system under test contains four 50Gbps NIC ports connected to 50 Gbps NIC ports in a server acting as the traffic generator. The system under test hosts the VPP IPsec container endpoints, while the traffic generating server hosts the traffic generator and also uses the DPDK testpmd application as a simple switch. The generated traffic flows bidirectionally throughout the setup. In the traffic generating server, the two ports receiving the clear data are used for the traffic generator, while the two ports receiving the encrypted data run the DPDK testpmd application to forward the traffic between the VPP IPsec endpoints.

Figure 4: Test Methodology for VPP IPsec



Figures 5-7 show the performance results obtained for aes-gcm-128, aes-gcm-256, and aes-cbc-256 crypto algorithms with different packet sizes. The connectivity speed of the ports used for the testing was 50 Gbps, and the traffic was generated at 100% line rate bidirectionally using these ports.

Figure 5: VPP IPsec aes-gcm-128 Results




Figure 6: VPP IPsec aes-gcm-256 Results





Figure 7: VPP IPsec aes-cbc-256 Results



The testing used 2C4T for each VPP IPsec endpoint, and in each configuration, the QAT SW stack using the IPsec Multi Buffer crypto handler provided a large throughput improvement over the unoptimized stack using the OpenSSL crypto handler. For 1024B packets, the **QAT SW** stack reached a throughput of **61.70 Gbps**, **58.15 Gbps**, and **46.78 Gbps** in the aes-gcm-128, aes-gcm-256, and aes-cbc-256 cases respectively. This corresponded to a performance gain of **5.3x**, **5.0x**, and **10.1x** in the respective crypto algorithm cases over the Unoptimized stack.

Malconv AI Software Benchmarks


AI inference is used in network/security to help prevent advanced cyber-attacks. In order to improve the latency associated with this application, the Intel® Xeon® Scalable Processor contains technologies to accelerate AI inference such as AVX-512, AMX, and Vector Neural Network Instructions. The Malconv AI workload utilizes the TensorFlow deep-learning framework, Intel® oneAPI Deep Neural Network Library (oneDNN), Advanced Matrix Extensions (AMX), and Intel® Neural Compressor to improve the performance of the AI inference model.

Malconv is an open-source deep-learning model used to detect malware by reading the raw execution bytes of files. An Intel® optimized version of the Malconv model is used as the initial configuration for this workload. The performance of the model can be improved by various procedures including conversion to a floating point frozen model and using the Intel® Neural Compressor for post-training quantization to acquire BF16, INT8, and ONNX INT8 precision models.

NOTE: Refer to <https://hub.docker.com/r/intel/malconv-model-base> for the Intel® Optimized Malconv Model.

Figure 8 shows the performance results obtained for each model of the workload testing. Seven models were tested and compared in total, and multiple instances of the inference test were run to fully saturate the socket's cores. The data in the figure uses three cores per instance.

Figure 8: Malconv AI Software Results



The best performance was obtained in the **INT8 model with**

AVX512_CORE_AMX enabled with a mean inference time of approximately **9.1ms**. Compared to the initial Intel® Optimized Malconv model which had a mean inference time of approximately 68.4 ms, the INT8 model with AVX512_CORE_AMX showed a speedup of about **7.5x**.

Summary

The Intel®-HPE Verified Reference Configuration for NFVI and SASE is an Intel® Accelerated Solution defined on 5th Gen Intel® Xeon® Scalable processors with Integrated Intel® QAT and Intel® Ethernet Adapter E810 for the solution. This solution, optimized for the HPE ProLiant Compute DL380 Gen11 system, combines architectural improvements, feature enhancements, and integrated accelerators with high memory and IO bandwidth, and provides a significant performance and scalability advantage in today's NFVI and SASE environments. With the advantages offered by the 5th Gen Intel® Xeon® Scalable processors, the following gains were achieved for each of the workloads:

- NGINX*: **12.73x** gain in Connections Per Second for Intel® QAT HW (2 devices) over the out-of-box software stack for 4C8T
- VPP IPSec: **5.3x**, **5.0x**, and **10.1x** gain in Throughput for the QAT SW stack over the Unoptimized stack for 1024B packets with the aes-gcm-128, aes-gcm-256, and aes-cbc-256
- Malconv AI: **7.5x** speedup in Mean Inference Time for the INT8 model with AVX512_CORE_AMX over the initial Intel® Optimized Malconv h5 model

The Intel®-HPE Verified Reference Configuration for NFVI and SASE defined for the ProLiant DL380 Gen11 utilizing these 5th Gen Intel® Xeon® Scalable processors with integrated Intel® QAT accelerator is optimized for network, cloud native, wireline, and wireless core-intensive workloads and is especially suited for NFVI and SASE workloads coupled with Intel® Ethernet E810-Network Controllers and Data Plane Development Kit.



Notices & Disclaimers

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein. 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.

2024© 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. 0607/ABL/Wipro/PDF  Please Recycle 787009-001US