White Paper

Communication Service Providers (CoSPs), Cloud Service Providers (CSPs)

Intel-HPE Verified Reference Configuration for NFVI & SASE on Red Hat OpenShift Container Platform

Simplify and Accelerate NFVI and SASE Deployments with Intel-HPE Verified Reference Configurations on 5th Gen Intel® Xeon® Scalable Processors

Authors

Jonathan Tsai

Software Enabling and Optimization Engineer, Intel

Ai Bee Lim

Solution Architect, Intel

Timothy Miskell

Cloud Software Architect, Intel

Vaishnavi Saravanan

Technical Enablement Engineer, Intel

Max Coates

Community Manager, Intel

Table of Contents

Intel Verified Reference Configurations (VRCs) offer workload-optimized infrastructure solutions designed to address today's complex demands using the Intel® Xeon® Scalable processor family. This document presents a reference implementation for the 5th Gen Intel® Xeon® Scalable Processor, featuring the Hewlett Packard Enterprise (HPE) ProLiant DL380 Gen11 server and the Red Hat® OpenShift® Container Platform.

By adopting this Intel VRC for Network Function Virtualization Infrastructure (NFVI) and Secure Access Service Edge (SASE) deployments, network operators, service providers, and enterprises can deploy virtualized network functions with enhanced security, simplicity, and efficiency.

This implementation helps to streamline the deployment process by integrating hardware and software components into a cohesive tried and tested solution. It enables operators to make straightforward design choices while achieving predictable high-performance, reducing the time and costs associated with evaluating the large choice of options. This VRC architecture aims to deliver consistent performance across NFVI and SASE workloads by leveraging in-house configuration, benchmarking and qualification later detailed in this whitepaper.

Intel VRC for NFVI and SASE

Built on a multi-node architecture—including a controller, a cloud node, and storage—the Intel VRC for NFVI and SASE provides all the necessary resources to deploy software-defined infrastructure within each cloud server instance which is managed by a hypervisor.

The Controller Node acts as the central hub for severeal functions; control, signaling, and management functions, implementing Virtual/Container Network Function Management (VNFM) and Virtualization Infrastructure Management (VIM). Due to its role, it typically does not require additional local storage or hardware acceleration, simplifying its function within the overall architecture.

This integrated approach offers a robust and flexible foundation for deploying NFVI and SASE solutions, allowing users to effectively manage and optimize their network and security functions with the reliability and performance of Intel architecture.

Note: For detailed solution specifications, please refer to the Intel VRC for NFVI and SASE on Red Hat OpenShift Container Platform document.

HPE ProLiant DL380 Gen11 Server

The HPE ProLiant DL380 Gen11 Server, powered by 5th Gen Intel® Xeon® Scalable processors, is designed to maximize data value and drive innovation from edge to cloud. Built for today's hybrid computing environments, this versatile and future-ready server delivers exceptional performance and reliability.

Figure 1. HPE ProLiant DL380 Gen11 Server

Key Features:

- **Optimized Workload Performance:** Tailored to handle demanding workloads with ease, ensuring consistent, high-speed processing power.
- **Seamless Cloud Integration:** Offers an intuitive cloud operating experience that simplifies deployment and management across diverse environments.
- **Built-In Security:** Designed with trusted security protocols from the ground up, including Security Protocols and Data Models (SPDM) for component authentication.
- **Enhanced Protection with HPE iLO 6:** Benefit from extended protection through the HPE partner ecosystem, leveraging HPE iLO 6 for robust security verification and enhanced system resilience.

HPE Hardware and Firmware Details

The NFVI solution features a cutting-edge hardware configuration that combines both Intel processor technology and Intel platform innovations, including Ethernet and acceleration technologies. Integrated on the motherboard, these technologies deliver best-in-class, low-latency performance optimized for demanding workloads with the Data Plane Development Kit (DPDK).

Note: For HPE ProLiant DL380 Gen11 drivers and firmware updates, visit HPE Support.

Hardware	Description
Processors	2 x Intel® Xeon® Gold 6538N Processor
Memory	16 x 32GB Dual Rank DDR5, 4800MHz, 1 DIMM per channel, Total Memory 512GB
Network Interface Card (NIC)	2 x Dual Port 100GbE Intel® Ethernet Network Adapter E810-2CQDA2 (Gen 4 x16)
Storage	2 x 960GB SSD NVMe solution as boot device
Platform Firmware	BIOS Version: 2.2, iLO6 Version: 1.54

Table 1. Hardware and Firmware details

Intel BIOS Reccomendation

Intel recommends configuring BIOS settings for deterministic performance with Turbo enabled to meet the stringent performance requirements of packet processing throughput workloads. For detailed BIOS settings recommendations, refer to Chapter 3.0 of "**BIOS Settings for Intel Wireline, Cable, Wireless, and Converged Access Platform**" - Document ID #747130. You can access this document through the Intel® Resource & Documentation Center.

Note: Please contact your Intel Field Representative to obtain access to this documentation.

HPE BIOS configuration

This section outlines the HPE BIOS configuration necessary to achieve optimal performance with an energy-balanced approach and Turbo mode enabled for high-throughput packet processing. HPE recommends using the "Network Function Virtualization Infrastructure - Secure Access Service Edge" profile for best results.

To apply this profile, navigate to **System Configuration > BIOS/Platform Configuration (RBSU) > Workload Profile** and select the "Network Function Virtualization Infrastructure - Secure Access Service Edge" profile. After applying the profile, reboot the system to activate the settings. This profile automatically adjusts other BIOS settings to ensure energy-efficient Turbo performance while maintaining high throughput for packet processing.

Note: After applying the above settings, a reboot is required for the changes to take effect.

If the BIOS is upgraded to a newer version after the profile is set, you should first load the BIOS defaults, and then re-apply the "Network Function Virtualization Infrastructure - Secure Access Service Edge" profile to ensure that the updates in the Telco Optimized Profile are correctly implemented.

Solution Software

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

Ingredient	Software Version
Red Hat* OCP	4.14
RHCOS Kernel	Kernel 5.14
Async mode NGINX*	0.5.1
VPP IPSec	23.10
Intel®-Tensorflow	2.14.0 with oneDNN 3.4.1

Table 2. Solution Software Versioning

Performance

This section verifies the performance metrics of the NFVI reference configuration, ensuring it operates as expected without anomalies. Review the information provided here to confirm that the platform meets its performance baseline targets. For a detailed explanation of the testing methodologies and workload evaluations, please refer to the whitepaper titled "**Intel Verified Reference Configuration for NFVI v5 and SASE on Red Hat OpenShift Container Platform (RHOCP).**" You can access this document through the Intel® Resource & Documentation Center.

Performance Baseline

To establish the performance baseline, several key applications must be executed after configuring the platform according to the Bill of Materials (BOM), BIOS settings, and Software Stack as outlined earlier. These applications measure critical performance metrics, including latency, memory bandwidth, and jitter, providing a clear understanding of the system's capabilities and ensuring it meets the expected performance criteria.

Memory Latency Checker (MLC)

The first application to run is the Memory Latency Checker (MLC), which is used to measure memory latency and bandwidth performance.

To download the latest version of MLC, visit the Intel® Memory Latency Checker download page. After downloading, unzip the tarball package, navigate to the Linux folder, and execute the MLC application using the command: `./mlc`

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

Table 3. MLC Data - Idle Latencies for Sequential Access

Read-Write Ratio	Peak Injection Memory bandwidth (MB/s)
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

Table 4. MLC Data - Read-Write Ratio Performance

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

Table 5. MLC Data - Idle Latencies for Sequential Access

Inject Delay	Latency (ms)	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

Table 6. MLC Data - Inject delay Latency Performance

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.

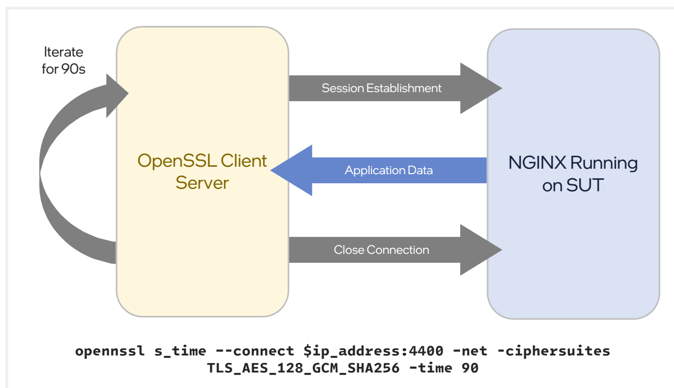

Figure 2. Test Methodology for SSL with NGINX

Figure 3 shows the performance results obtained for each configuration of the workload testing. The cipher ECDHEX448-RSA4K with TLS 1.3 was used for the testing. Intel® QuickAssist Technology (Intel® QAT) hardware acceleration helps to offload public key exchange for SSL layer 7 applications.

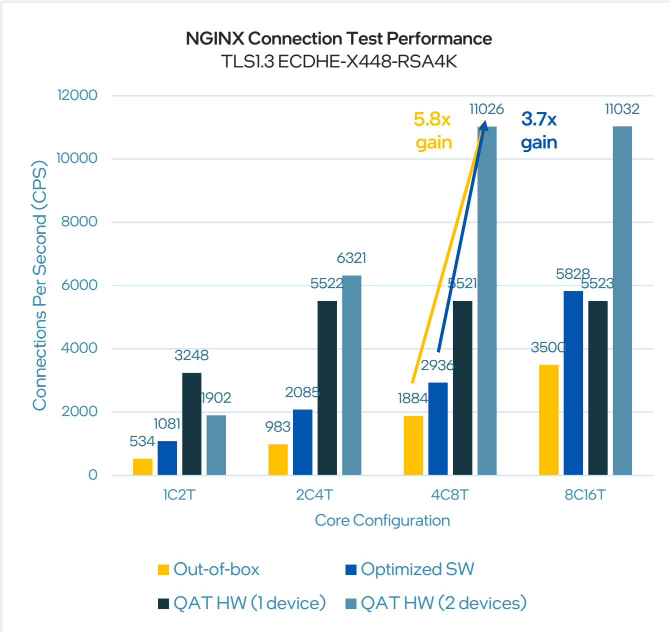

Figure 3. NGINX* Software Results

In a 4C8T configuration, the Intel Verified Reference Configuration for NFVI demonstrated approximately 11,026 Connections Per Second (CPS) when running the full software stack of the NGINX* application with Intel® QAT hardware (2 devices). In comparison, without Intel® QAT, the system achieved about 1,884 CPS using the out-of-box software stack of NGINX*. This results in a significant performance boost: the full software stack with Intel® QAT hardware (2 devices) delivers approximately a 5.8x increase in CPS over the out-of-box software stack, and about a 3.7x gain over the optimized software configuration.

VPP-IPsec Software Benchmarks

The VPP IPsec performance testing was conducted using a system configured with four 50 Gbps NIC ports, connected to corresponding 50 Gbps NIC ports on a server acting as the traffic generator. The system under test hosted the VPP IPsec container endpoints, while the traffic generating server utilized the DPDK testpmd application as a simple switch.

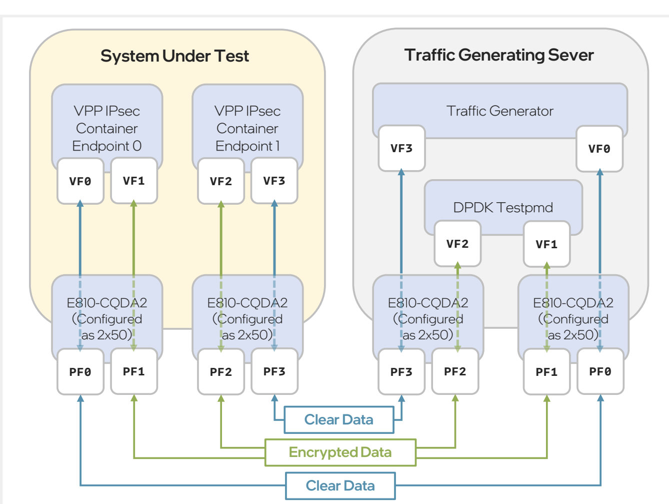

Figure 4. Test Methodology for VPP-IPsec

Traffic was generated and flowed bidirectionally across the setup: the server's ports receiving clear data handled traffic generation, while the ports receiving encrypted data used the DPDK testpmd application to forward traffic between the VPP IPsec endpoints.

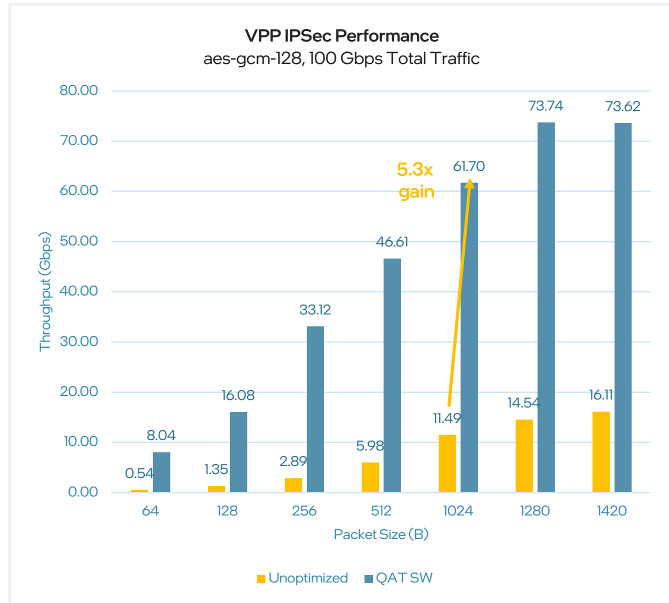

Figure 5. VPP IPsec AES-GCM-128 Test Results

Figures 5-7 illustrate the performance results for the AES-GCM-128, AES-GCM-256, and AES-CBC-256 cryptographic algorithms across various packet sizes. The testing was conducted at a connectivity speed of 50 Gbps, with traffic generated at 100% line rate bidirectionally through the configured ports.

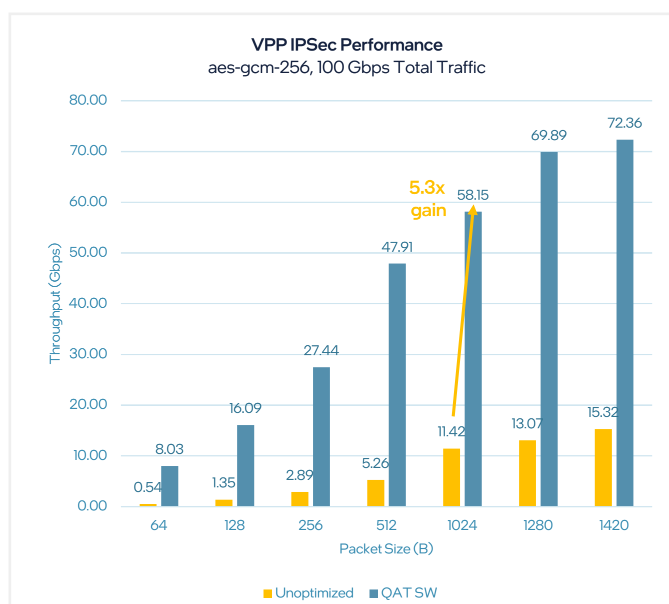

Figure 6. VPP IPsec AES-GCM-256 Test Results

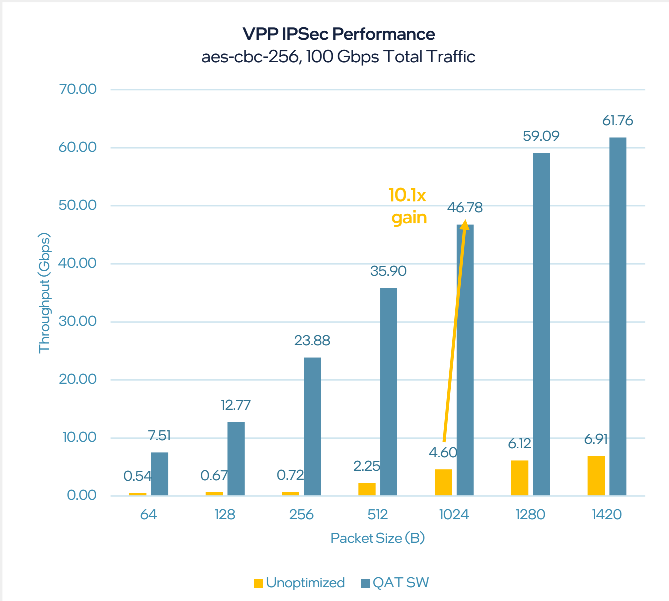

Figure 7. VPP IPsec AES-CBC-256 Test Results

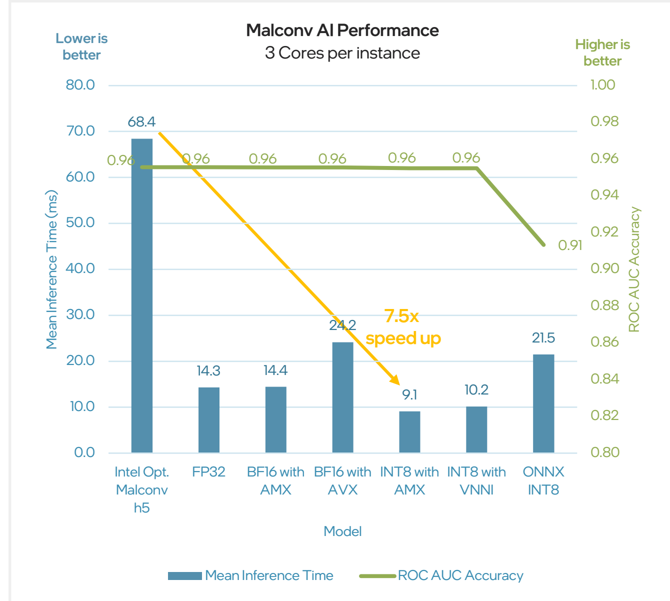

Figure 8. Malconv AI Software Results

Malconv AI Software Benchmarks

AI inference plays a crucial role in network and security applications, helping to prevent advanced cyber-attaks. To enhance latency performance in AI inference, the Intel® Xeon® Scalable Processor leverages advanced technologies such as AVX-512, Advanced Matrix Extensions (AMX), and Vector Neural Network Instructions. The Malconv AI workload utilizes the TensorFlow deep-learning framework, Intel® oneAPI Deep Neural Network Library (oneDNN), AMX, and Intel® Neural Compressor to optimize the AI inference model's performance.

Malconv is an open-source deep-learning model designed for malware detection by analyzing raw execution bytes of files. For this workload, an Intel®-optimized version of the Malconv model was used as the initial configuration. Performance improvements were achieved through various techniques, including converting the model to a floating point frozen version and applying Intel® Neural Compressor for post-training quantization to obtain BF16, INT8, and ONNX INT8 precision models. For more details on the Intel® Optimized Malconv Model, refer to the Intel® Optimized Malconv Model on Docker Hub.

Figure 8 displays the performance results for each model tested, with seven models in total. Multiple instances of the inference test were executed to fully utilize the cores of the socket, with data shown for configurations using three cores per instance.

The best performance was achieved with the INT8 model utilizing AVX512_CORE_AMX, resulting in a mean inference time of approximately 9.1 milliseconds. This represents a significant speedup—about 7.5x faster—compared to the initial Intel® Optimized Malconv model, which had a mean inference time of approximately 68.4 milliseconds.

Conclusion

This Intel®-HPE Verified Reference Configuration for NFVI and SASE is expertly engineered for the HPE ProLiant DL380 Gen11 server, integrating 5th Gen Intel® Xeon® Scalable processors with Intel® QuickAssist Technology (Intel® QAT) and Intel® Ethernet Adapter E810. This configuration leverages cutting-edge architectural enhancements and built-in accelerators to deliver exceptional memory and I/O bandwidth, ensuring superior performance and scalability required for today's NFVI and SASE environments.

Key Performance Improvements Achieved with This Configuration Include:

- **NGINX:** 12.73x increase in Connections Per Second using Intel® QAT hardware (2 devices) compared to the standard software stack for a 4C8T setup.
- **VPP IPSec:** Delivered 5.3x, 5.0x, and 10.1x throughput improvements for the QAT software stack over the unoptimized stack with 1024B packets for aes-gcm-128, aes-gcm-256, and aes-cbc-256 encryption algorithms respectively.
- **Malconv AI:** Achieved a 7.5x speedup in Mean Inference Time for the INT8 model with AVX512_CORE_AMX compared to the initial Intel® Optimized Malconv h5 model.

Designed to support network, cloud-native, wireline, and wireless core-intensive applications, this configuration also utilizes Intel® Ethernet E810 Network Controllers and the Data Plane Development Kit (DPDK). Together, these best in class components provide a robust, scalable, and high-performance infrastructure solution tailored to meet today's complex networking and security demands.

intel.

Notices & Disclaimers

Availability of accelerators varies depending on SKU. Visit the Intel Product Specifications page for additional product details.
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 configuration disclosure for configuration details.
No product or component can be absolutely secure.
Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
Your costs and results may vary.
Intel technologies may require enabled hardware, software or service activation.
You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a nonexclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.
The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.
© 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.

1009/MC/TM/JT/ABL/PDF ♻ Please Recycle 362171-002US