Solution Brief

Intel® Select Solutions | Version 2
NFVI Forwarding Platform
3rd Generation Intel® Xeon® Scalable Processors

Intel Select Solutions for NFVI Forwarding Platform

Intel Select Solutions for NFVI Forwarding Platform are verified configurations of compute, network, storage, and middleware elements for packet-processing-intensive network workloads that can significantly accelerate time to production and reduce TCO.

Introduction

As communications service providers transform their networks for 5G, they are driving a rapid transition to virtualization and a dependency on network technology innovations that meet massive new data-processing requirements.

The primary traffic load in 5G is in the user-plane functions (the data, as opposed to the control or signaling associated with establishing traffic flow). Network implementations based on NFVI can increase speed and agility, improving time to market, scale for future capacity, and new services within CoSPs' operational networks. CoSPs are expected to continue the transformation from purpose-built solutions, beyond virtualized appliances, to a model of fully virtualized and cloud-native based networks. To accomplish this transformation, the journey to a more fully disaggregated deployment model must be realized.

NFVI standards include the consideration that solutions at the application layer may include both virtual machine (VM) or containerized application functions; the disaggregation of the lower elements of compute, network, and storage remain critical. CoSP workloads will demand that the underlying systems can scale network capacity rapidly and yield resources when capacity requirements abate during the normal course of a day and throughout the lifetime of the installed systems.

Strong ecosystem collaboration is essential, and Intel has worked closely with members of the Intel Network Builders program to create the Intel Select Solutions for NFVI. The Intel Select Solutions for NFVI Forwarding Platform delivers a verified hardware and software stack capable of sustained optimized data- and control-plane processing within the operational networks of the CoSP. These workloads span the 4G and 5G wireless functions found in vEPC and vUPFs, as well as in the wireline network gateway functions vBRAS/vBNG and vCMTS within the DSL and MSO networks.

In certain configurations, the Intel Select Solutions for NFVI Forwarding Platform shows a theoretical throughput of up to 800 Gbps per system. The disaggregated model of separating the network, compute, and storage from the upper-layer application function has led in some cases to less than desired processing capability, mainly due to architectural shortcomings at various layers of the NFVI implementation. Examples of the causes for these shortcomings include resources being misaligned on the server, a lack of access to interfaces, or incompatible versions of drivers or accelerators.

Working together, Intel and other leaders in the NFVI transformation have recognized this challenge and designed the Intel Select Solutions for NFVI Forwarding Platform to address many of these concerns. The reference architecture builds from a foundation of a balanced server design, where each NUMA node has identical capacity resources. The solution is then further specified with a well-known vertical stack of all the critical software elements, including the BIOS configuration, host operating system, and network and acceleration drivers. All components are verified together to provide a system with known performance and optimized capabilities to meet the demands of CoSPs' operational network workloads.

What Are Intel Select Solutions?

Intel Select Solutions are predefined, workload-optimized solutions designed to minimize the challenges of infrastructure evaluation and deployment. Solutions are validated by OEMs/ODMs, certified by ISVs, and verified by Intel. Intel develops these solutions in extensive collaboration with hardware, software, and operating system vendor partners and with the world's leading data center and service providers. Every Intel Select Solution is a tailored combination of Intel data center compute, memory, storage, and network technologies that delivers predictable, trusted, and compelling performance.

To refer to a solution as an Intel Select Solution, a vendor must:

1. Meet the software and hardware stack requirements outlined by the solution's reference-design specifications
2. Replicate or exceed established reference-benchmark test results
3. Publish solution content to facilitate customer deployment

Solution providers can also develop their own optimizations in order to give end customers a simpler, more consistent deployment experience.

Common High-Throughput Node Architecture Across Sites

The Intel Select Solutions for NFVI Forwarding Platform are designed to maximize network I/O capacity and packet-processing throughput per node with a scalable architecture for deployment across various types of network sites. The specification includes nodes with various theoretical throughputs to support sites with different requirements, from multiple terabits per second at the network core to lower throughputs as the network branches out to progressively smaller sites.

The multi-node architecture of the Intel Select Solutions for NFVI Forwarding Platform supports the Control and User Plane Separation (CUPS) strategy. Specifically, that separation allows user-plane functions to be scaled out across multiple systems. The resulting topology allows for a many-to-one relationship between the node types that enables more efficient use of hardware resources, for lower TCO. Examples of typical user-plane services that might be deployed at different levels of a theoretical CoSP Network Infrastructure are shown in Table 1.

The solution has been tuned and pre-tested to ensure high throughput across different types of VNFs, using the reference architecture's hardened stack in a controlled environment. In production, this assurance helps accelerate time to market and mitigates implementation risk.

Table 1. Typical VNF workloads deployed at specific network locations.

Level 1: Core Network Site	vEPC (virtual Evolved Packet Core)
	vGiLAN (virtual Gateway Interface Local Area Network) / 5G UPF (5G User Plane Function)
	vIMS (virtual IP Multimedia System)
Level 2: Regional Points of Presence	vEPC (virtual Evolved Packet Core)
	vGiLAN (virtual Gateway Interface Local Area Network)
	vIMS (virtual IP Multimedia System)
	vCGNAT (virtual Carrier-Grade Network Address Translation)
	vCRAN (virtual Cloud Radio Area Network)
Level 3: Remote Central Offices	vBNG (virtual Broadband Network Gateway)
	vBRAS (virtual Broadband Remote Access Server)
	dEPC (distributed Evolved Packet Core) / S/P GW (Secure/Packet Gateway)
	vDPI (virtual Deep Packet Inspection) / vCPE (virtual Customer Premise Equipment)
	vMEC (virtual Multi-Access Edge Compute)
	vCMTS (virtual Cable Modem Termination System)
Level 4: Access Central Offices	vRAN (virtual Radio Access Network)
	vOLT (virtual Optical Line Terminator) / DSL (Digital Subscriber Line)
	vMEC (virtual Multi-Access Edge Compute)
	vCMTS (virtual Cable Modem Termination System)

Intel Select Solutions for NFVI Forwarding Platform: Hardware Configurations

The Intel Select Solutions for NFVI Forwarding Platform defines a hyperconverged infrastructure in a 1U or 2U rack-mounted configuration, with solution components and configurations selected to ensure maximum I/O throughput. The hardware topology incorporates 3rd generation Intel Xeon Scalable processors, Intel Ethernet Network Adapters for DPDK-accelerated networking, and Intel Solid State Drives (Intel SSDs).

This platform addresses general use cases for NFVI, focusing resources on I/O to provide the widest data path possible into each NUMA node. Optional components can be added to meet the requirements of specific use cases. Intel® Optane™ persistent memory can be added to the configurations to provide massive memory resources that enlarge the pool of warm data that can be held in close proximity to the processor. User-plane nodes are available in two primary configurations, with configurability to fine-tune the stack for specific solution needs:

- **Plus node:** This configuration is tailored for the highest performance and highest density for maximum I/O packet-processing.
- **Base node:** This configuration is a value/performance- optimized offering suited to deployments further from the network core.

In addition, the reference architecture specifies system configuration parameters for the Intel Select Solutions for NFVI Forwarding Platform controller node. Configuration guidelines for all three types of nodes are given in Table 2. All components are required unless otherwise noted.

Table 2. Intel Select Solutions for NFVI Forwarding Platform hardware configurations (required unless otherwise noted).

Ingredient	Plus Configuration	Base Configuration	Controller Node Configuration
Processors	2x Intel Xeon Gold 6338N processor	2x Intel Xeon Gold 5318N processor	2x Intel Xeon Gold 5318N processor
Memory	512 GB DDR4	256 GB DDR4	
Intel® Optane™ PMem	Recommended		
Discrete Network Adapters	4x Intel Ethernet Network Adapter E810 2CQDA2	2x Intel Ethernet Network Adapter E810 2CQDA2 or 4x Intel Ethernet Network Adapter E810 CQDA2	2x Intel Ethernet Network Adapter E810 CQDA2
Local Storage	2x Intel SSD D3-S4510 Series or higher @ 480 GB or larger		
LAN on Motherboard	10 Gbps or 25 Gbps port for Pre-boot Execution Environment (PXE) and Operation, Administration and Management (OAM)		
	1/10 Gbps port for management		

The Intel Select Solutions for NFVI Forwarding Platform reference architecture is designed for high throughput across the CoSP Network Infrastructure. Intel technologies used in the configuration specifications are described below.

3rd Generation Intel Xeon Scalable processors are the foundation for all hardware configurations used in the solutions. These CPUs are available in configurations optimized for diverse network environments, in a wide range of frequency, feature and power levels. The following platform characteristics and technologies are particularly valuable to the Intel Select Solutions for NFVI Forwarding Platform:

- **New balanced, scalable architecture** increases per-core performance, memory and I/O bandwidth to accelerate diverse workloads from the edge to the data center.
- **Increased core count and cache**, available in a flexible range of SKUs with from 8 to 40 powerful cores and L1 cache up to 48 KB (50 percent larger than the previous generation), helping drive up the number of subscribers that can be handled per node.
- **New extensions to Intel AVX-512 instructions** accelerate bit-processing kernels that move and reorder blocks of data within the wireless signal processing pipeline.
- **Intel Software Guard Extensions (Intel SGX)** provides protected execution enclaves that isolate and help protect application code and data while in use and enable new ways to collaborate using shared data without compromising privacy.

Intel Ethernet 800 Series Network Adapters

Standards-based networking performance across NFVI workloads is provided by the Intel Ethernet 800 Series Network Adapters through a combination of sophisticated packet-processing, intelligent offloads and accelerators, and high-quality open-source drivers for data plane processing. In addition to optimizing throughput, the adapters are designed to enable broad interoperability and agility. Key features and capabilities associated with the adapters include the following:

- **Application Device Queues (ADQ)**, a capability unique to the Intel Ethernet 800 Series Network Adapters, provides dedicated queues to key workloads, enabling application-specific data steering, signaling, and rate limiting using an optimized application thread-to-device data path. ADQ increases predictability, reduces latency and jitter, and improves throughput.
- **Dynamic Device Personalization (DDP)** is a programmable packet-processing pipeline provided by the Intel Ethernet 800 Series Network Adapters that supports on-demand reconfiguration of network controllers at runtime, enabling workload-specific optimizations to increase throughput and decrease latency. DDP is enhanced in the Intel Ethernet 800 Series Network Adapters with greater programmability than its predecessor, as well as workload-specific protocols for added flexibility.
- **Data Plane Development Kit (DPDK)** is an open-source set of libraries and drivers that accelerates packet- processing in the data path. It also facilitates building packet forwarders designed to operate on general-purpose, standards-based servers.

The adapters deliver excellent small-packet performance that is well suited to the requirements of NFVI, together with advanced I/O virtualization that helps drive up throughput on virtualized servers. In addition, they offer network virtualization optimizations including VXLAN, GENEVE, NVGRE, MPLS, and VXLAN-GPE with Network Service Headers (NSH).

Intel Optane Persistent Memory 200 Series

A redefined memory tier based on Intel Optane persistent memory (recommended) improves overall system performance and reduces latency by putting more data close by the processor on non-volatile memory, reducing the need for disk accesses. It combines the byte-addressability of DRAM with the persistence of storage, with idle read latency that's an order of magnitude faster than SSDs or other storage types. In a form factor that's socket-compatible with DDR4, Intel Optane persistent memory is available in capacities from 128 GB to 512 GB.

Verified Performance Through Benchmark Testing

All Intel Select Solutions are verified by Intel to meet a specified minimum level of workload-optimized performance capability. Verified Intel Select Solutions for NFVI Forwarding Platform meet or exceed vBNG design and testing standards, as shown in Table 3.

Table 3. Minimum vBNG performance standards for Intel Select Solutions for NFVI Forwarding Platform. System builders, system integrators, and solution and service providers can further optimize their systems to achieve higher performance and capability.

Benchmark	Target for Plus Configuration[1]
Overall Throughput	620 Gbps per server
Latency	50 µs at 256 B packets
Overall Server Power	<800 W
Per-CPU Power	<185 W per socket

Software and Firmware Stack

The Intel Select Solutions for NFVI Forwarding Platform includes a comprehensive, workload optimized software and firmware stack based on Red Hat Enterprise Linux and Red Hat OpenStack, though it can accommodate other operating systems and VIMs as well. Contact your Intel representative for access to a detailed design specification for each configuration. These specifications and associated test plans are posted to the Intel Resource and Design Center.

The solution details each software stack component in order to reduce chances for incompatibilities, eliminate errors, and provide for reliable and high-performance operation. Variables specified include minimum supported versions of the following key items:

- Intel processor microcode update versions
- UEFI firmware (BIOS)
- Intel Ethernet Controller E810 firmware and drivers
- Host and guest operating system versions
- Docker container version
- Data Plane Development Kit version

Conclusion

The Intel Select Solutions for NFVI Forwarding Platform are based on a workload-optimized reference architecture that is purpose-built for high-throughput workloads across the network infrastructure. With pre-validated, pre-optimized solution stacks from a choice of OEMs, the solution can dramatically accelerate deployment and time to new services, while reducing implementation risk for CoSPs.

Learn More

Intel Select Solutions: **intel.com/selectsolutions**

3rd Generation Intel Xeon Scalable processors:
intel.com/content/www/us/en/products/processors/xeon/scalable.html

Intel Ethernet Technology: **intel.com/ethernet**

Intel Select Solutions are supported by the Intel Network Builders Program:
networkbuilders.intel.com/intelselectsolutions/network

[1] Testing conducted by Intel on March 17, 2021 and may not reflect all publicly available updates. One node, 2x Intel® Xeon® Gold 6338N processors (2.2 GHz, 32 cores/64 threads); Intel Hyper-Threading Technology enabled; Intel Turbo Boost Technology disabled; DRAM 512 GB (16 slots/32 GB/2666 MT/s); 1x Intel SSD (960 GB); 4x Intel® Ethernet Controller E810 2CQDA-2; BIOS: SE5C6200.86B.0020.P09.2012290309, 12/29/2020; Red Hat Enterprise Linux 8.2 (Ootpa); kernel: 4.18.0-193.28.1.el8_2.x86_64; Workload & version: vBNG 20.11; Compiler: 8.3.1 20191121 (Red Hat 8.3.1-5) (GCC); DPDK 20.11; libvirt/qemu: 6.0.0/4.2.0; XXV710 driver/FW: 1.3.2 / 2.40 0x800063ef 1.2852.0 / 4.0.2; E810 driver/FW/iavf: 1.3.24; E810 DDP Comms package: 30,024x 2MB Hugepages configured.

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 non-exclusive, 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.

0321/MMQ/MESH/338348-004US