Solution Brief

Intel® Network Builders
NFVI Forwarding Platform
Intel Ethernet 800 Series Update

Intel Select Solutions for NFVI Forwarding Platform

The Intel® Select Solutions for NFVI Forwarding Platform are verified systems of compute, network, storage, and middleware elements for packet-processing-intensive workloads that can significantly accelerate the introduction of next-generation virtualized solutions into the network, reducing both time to production and TCO.

Introduction

Communication service providers (CommSPs) continuously transform and introduce new technology into their networks. Today the challenge is to accelerate the introduction of standards-based, high-capacity servers capable of meeting the unprecedented level of data growth in operational networks due to the introduction of 5G, which brings an order of magnitude increase in capacity of network traffic without a corresponding increase of an order of magnitude in cost.

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 for added capacity growth, as well as potentially for new services within CommSPs' operational networks. CommSPs are expected to continue the transformation from purpose-built solutions of the past, beyond virtualized appliances, to a model of fully virtualized and cloud native 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. CommSP 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.

The community effort necessary for the success of the NFVI transformation includes ISVs, OEMs, and OSVs, in addition to the CommSPs themselves. Intel has built an ecosystem of partners from each of these segments within the Intel Network Builders program, and working closely with industry leaders, Intel has co-created the Intel Select Solutions for NFVI. Within this segment, 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 CommSP. 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 400 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 CommSPs' operational network workloads.

What Are Intel Select Solutions?

Intel Select Solutions are pre-defined, workload-optimized solutions designed to minimize the challenges of infrastructure evaluation and deployment. Solutions are validated by OEMs/original design manufacturers (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 Intel Select Solutions, 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. Provide a product website URL with the Intel Select Solutions logo pictured and statement that this solution is now verified as Intel Select Solutions for xxx.
4. Publish a solution brief and an optional detailed implementation guide 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 CommSP Network Infrastructure are shown in Table 1.

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)

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.

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 2nd generation Intel Xeon® Scalable® processors, Intel Ethernet Server 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.

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

2nd Generation Intel Xeon Gold Processors

All node configurations feature 2nd generation Intel Xeon Gold processors, with at least 16 physical execution cores for the controller node and a minimum of 20 physical cores for the user-plane nodes. Solutions can use a range of processor SKUs to tailor performance to the specific needs of the implementation, deployed with two-socket systems that combine rigorous requirements for high performance and efficiency, coupled with low TCO.

To boost processor-to-processor data flow, Intel Ultra Path Interconnect (Intel UPI) provides transfer speeds between sockets of up to 10.4 GT/s,[1] while delivering high energy efficiency. The platforms also feature enhanced performance and bandwidth across six memory channels.

Intel Advanced Vector Extensions 512 (Intel AVX-512) doubles the amount of data handled per instruction compared to predecessor Intel AVX2 technology.[2] Intel Speed Select Technology enables the platform to adjust operating frequency and voltage in response to workloads, optimizing the balance between performance and energy efficiency. In addition, Intel Xeon Gold processors also feature important platform technologies that are required or recommended for Intel Select Solutions for NFVI Forwarding Platform:

- **Intel Virtualization Technology (Intel VT)** provides hardware abstraction so multiple workloads can share resources; workloads are isolated in hardware. Use of Intel VT in the solutions is required.
- **Intel Boot Guard** provides hardware-based integrity protection for the system boot blocks against compromise. Use of Intel Boot Guard in the solutions is required.
- **Intel Trusted Execution Technology (Intel TXT)** tests the integrity of the software environment at system start-up by comparing it to a known good copy. Use of Intel TXT in the solutions is optional.

2nd Generation Intel Xeon Scalable processors:

- Offer high scalability that is cost-efficient and flexible, from the multicloud to the intelligent edge
- Establish a seamless performance foundation to help accelerate data's transformative impact
- Support breakthrough Intel Optane persistent memory technology
- Accelerate AI performance and help deliver AI readiness across the data center
- Provide hardware-enhanced platform protection and threat monitoring

Intel Ethernet 700 and 800 Series

Standards-based networking performance across NFVI workloads is provided by the Intel Ethernet Network Adapter XXV710 and Intel Ethernet Network Adapter E810 CAM-2 through a combination of sophisticated packet-processing, intelligent offloads and accelerators, and high-quality open-source drivers. 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, 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 700 and 800 Series 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 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

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.

Software and Firmware Stack

The Intel Select Solutions for NFVI Forwarding Platform includes a comprehensive, workload optimized software and firmware stack, as shown in Table 4. While the solution initially focuses on Red Hat Enterprise Linux and Red Hat OpenStack, it can accommodate other OSs and VIMs as well.

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 6252 processor @ 2.1 GHz or Intel Xeon Gold 6252N processor @ 2.3 GHz, 24C/48T or higher	2x Intel Xeon Gold 6230 processor @ 2.1 GHz or Intel Xeon Gold 6230N processor @ 2.3 GHz, 20C/40T or higher	2x Intel Xeon Gold 5218 processor @ 2.3 GHz or Intel Xeon Gold 5218N processor @ 2.3 GHz, 16C/32T or higher
Memory	384 GB DDR4-2666 or 192 GB DDR4-2666 plus Intel Optane persistent memory		192 GB DDR4-2666 or 192 GB DDR4-2666 plus Intel Optane persistent memory
Intel Optane Persistent Memory	1-1.5 TB Intel Optane persistent memory (recommended)	512 GB Intel Optane persistent memory (recommended)	
Discrete Network Adapters	4x Intel Ethernet Network Adapter E810-CAM2 dual-port @ 100 Gbps or 4x Intel Ethernet Network Adapter XXV710 SPF 28+ quad-port @ 25 Gbps	2x Intel Ethernet Network Adapter XXV710 SPF 28+ quad-port @ 25 Gbps or 4x Intel Ethernet Network Adapter XXV710 SPF 28+ dual-port @ 25 Gbps	2x Intel Ethernet Network Adapter XXV710 SPF 28+ dual-port @ 25 Gbps
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		

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[3,5]	Target for Base Configuration[3,4]
Overall Throughput	310 Gbps per server	150 Gbps per server
Latency	50 µs at 256 B packets	50 µs at 256 B packets
Overall Server Power	<600 W if using Intel Ethernet 800 Series <650 W if using Intel Ethernet 700 Series	<430 W
Per-CPU Power	<160 W per socket	<100 W per socket

Table 4. Intel Select Solutions for NFVI Forwarding Platform software configurations (minimum).

	Ingredient	Software Version Details	
Firmware	BIOS MCU	SE5C620.86B.0X.02.0245 Release Date: Sept 28, 2019 0x5002f00	
	XXV710 NIC FW	v7.30 or later	
	E800 NIC FW	v2.0.0	
	Intel Ethernet Converged Network Adapter E810	800038FC_signed	
	Intel Optane persistent memory FW, DIMM FW	NVMDIMMDriver: v01.00.00.3371 NVMDIMMHii: v01.00.00.3371	
Host	OS	Red Hat Enterprise Linux	RHEL8.2-kernel-4.18.0-193.el8.x86_64
	Hypervisor	KVM/QEMU	2.12.0
	Libvirt	Libvirt	4.5.0
	Docker	docker	Version 18.09.7, build 2d0083d
	APPs	DPDK	20.05
		CollectD	5.8.0
	Drivers	I40e	2.8.10-k
		ice	1.0.4
		iavf	3.11.0_rc6
Guest	APPs	DPDK	18.11, 19.02
	OS	Ubuntu	20.04 LTS
		Red Hat Enterprise Linux	8.2
		CentOS	8.1
	Drivers	I40evf	3.2.3-k

Conclusion

The Intel Select Solutions for NFVI Forwarding Platform are based on a workload-optimized reference architecture that is purpose-built for high throughput and to provide a single coherent infrastructure for NFV that will enable the 5G rollout over the next several years, by means of network transformation at the access, next-gen central office (NGCO) and core. 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 CommSPs.

Learn More

Intel Select Solutions:
intel.com/selectsolutions

2nd-generation Intel Xeon Scalable processors:
intel.com/xeonscalable

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

[1] "Second Generation Intel Xeon Processor Scalable Family Technical Overview." https://software.intel.com/en-us/articles/second-generation-intel-xeon-processor-scalable-family-technical-overview.

[2] "Intel Advanced Vector Extensions 512 (Intel AVX-512) Overview." https://www.intel.com/content/www/us/en/architecture-and-technology/avx-512-overview.html.

[3] Testing conducted by Intel on July 22, 2019 and may not reflect all publicly available security updates. No product or component can be absolutely secure. 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 complete information, visit www.intel.com/benchmarks.

[4] Servers for the base configuration used were based on the Intel Xeon Gold 6230N CPU with 20 cores @ 2.30 GHz, 384 GB of RAM, 2x Quad-port Intel Ethernet Controller XXV710 (8 x 25 GbE ports) from Silicom. Key software tested: Red Hat Enterprise Linux Server release 8.2 (Ootpa), VBNG.L.v19.03.

[5] Servers for the plus configuration used were based on the Intel Xeon Gold 6252N CPU with 24 cores @ 2.30 GHz, 384 GB of RAM, 4x Intel E810 CAM2 (dual-port 100 GbpE). Key software tested: Red Hat Enterprise Linux Server release 8.2 (Ootpa), VBNG.L.v20.7.0.