

Solution Brief

**Intel® Select Solution for NFVI Forwarding Platform
2nd Generation Intel® Xeon® Processor Scalable Family**

Accelerating Forwarding Plane Performance

Overview

As Communication Service Providers transition from proprietary single-function networking elements to Network Function Virtualization Infrastructure (NFVI) they will begin to realize the transformational benefits of Software Defined Networking (SDN) where services become increasingly migrated into cloud-native workloads. Both new and legacy networking services and business policies can then be instantiated on-demand at the network edge in the form of software instances that are fully managed and orchestrated in the cloud.

In preparation for the widespread use of 5G and the broad array of services that will be sliced over the network, it is vital that newly deployed platforms based on Intel® Architecture (IA) take full advantage of Intel® Select Solutions for NFVI configurations in order to provide predictable performance and offer the throughput headroom needed by new and future services. More specifically, the key to network transformation success lies in the use of servers with NUMA-balanced I/O designs capable of accelerating data plane performance beyond that of traditional IT servers.

With the right server architecture and configuration, further CAPEX and OPEX savings can be achieved with enhanced commercial off the shelf solutions destined for next generation use cases: 4G or 5G core User Plane Functions (UPF), broadband use cases such as virtual Broadband Network Gateways (vBNG), network services such as virtual Evolved Packet Core (vEPC) and IPsec Gateways, as well as cable use cases such as virtual Cable Modem Termination System (vCMTS) -- all demand scalability, high performance and high throughput.

When CommSPs are ready to turn to an NFVI forwarding plane solution, they must be able to quickly and efficiently deploy it, realizing reliable, secure and workload-optimized deterministic performance on a balanced platform. To achieve this, they need a framework that makes it easier for them to deploy NFVI in a timelier fashion. Intel® Select Solutions for NFVI Forwarding Platform (NFVI FP) provides this new framework, and helps CommSPs to meet the challenges of next generation networking and compute more efficiently.

This solution brief describes how Intel® Select Solutions for NFVI FP address the complexity that service providers face in choosing the right infrastructure for NFV. It discusses Advantech's SKY-8101D configurations in detail and compares it with the base and plus specifications required to meet the Intel® Select Solutions for NFVI FP specifications based on 2nd generation Intel® Xeon® Scalable processors. Ready to ship as pre-configured platforms developers can now gain faster access to optimized and stable platform configurations to conduct testing and modeling of solutions that will help them define next-generation services.

Solution Highlights

- ✓ *NFV-Ready 1RU high-performance, dual-socket, NUMA balanced server*
- ✓ *Fully certified with Red Hat Enterprise Linux and Red Hat OpenStack Platform*
- ✓ *Maximizes network I/O per NUMA node for increased subscriber density, higher throughput and greater cost efficiency per rack unit*
- ✓ *400Gbps throughput per system, 200Gbps full-duplex per CPU socket*

Advantech's SKY-8101D enables enhanced solutions that demand high performance and throughput:

- 4G or 5G core User Plane Functions (UPF)
- Broadband use cases, such as virtual Broadband Network Gateway (vBNG)
- Network services such as virtual Evolved Packet Core (vEPC) and IPsec Gateways (IP-SecGW)
- Cable use cases such as virtual Cable Modem Termination System (vCMTS)

The SKY-8101D is a verified Intel® Select Solution for NFVI Forwarding Platform (NFVI FP) providing operators with faster access to optimized and stable server configurations that accelerate their NFV deployments.



Virtualize, Automate & Deploy with Advantech White Boxes

ADVANTECH


Introduction

Advantech works closely with Intel on a number of network transformation initiatives including Intel® Select Solutions, a program created by Intel to provide reference designs unique to their targeted workloads. An Intel Select Solution reference design offers a checklist of required and recommended components including CPU, chipset, accelerators, NICs, storage, BIOS settings etc. In addition Advantech collaborates closely with independent software vendors (ISVs) and operating system vendors (OSVs), as active members of the Intel Network Builder ecosystem, to ensure system-level optimization of hardware, software and firmware, including the bios, and other settings.

Intel Select Solutions provides workload optimized configurations across the data center, and for network specific workloads, Intel has launched Select Solutions for NFVI Forwarding Platform (NFVI FP). Specifically targeted at vEPC, vBNG & vCMTS workloads, the Intel Select Solution for NFVI FP helps developers by giving them an ultra-performant

commercially viable and tested platform.

As CommSPs gain more experience with NFV solutions based on Intel Select Solutions, they will gain from the flexibility enabled by a software-centric approach where a single edge NFVI platform can serve the evolving service needs of consumer, residential, and enterprise customer segments in both wireless and fixed networks.

The end result is faster time to deployment using off-the-shelf building blocks and improved Quality of Experience (QoE) for customers.

Advantech verified Intel® Select Solutions

Advantech's SKY-8000 Server product range is designed with a NUMA balanced PCIe approach in mind. PCIe lanes are distributed equally between each CPU socket rather than

connected to a single CPU or in the unbalanced manner found in many commercial servers. This avoids unnecessary latency and makes it ideal for networking applications as each socket can scale the application linearly.

Intel Select Solutions for NFVI defines the minimum NFVI hardware and software for specific application use-cases and sets the baseline for performance.

For example, the SKY-8101D 1RU dual-socket server based on Intel® Xeon® Scalable processors offers 200Gbps of PCIe bandwidth per socket, for up to 400Gbps per 1RU. This is ideal for applications that are inline (for example BNG) where 100Gbps upstream and 100Gbps downstream capacity per socket maximizes the number of serviceable user sessions.

While testing the platform with the Intel Select Solutions for NFVI FP scripts, Advantech recorded over 310 Gbps of PPPoE traffic throughput on the SKY-8101D with no packet loss, while overall server power consumption was less than 650W. When testing the

system with simple Layer 3 forwarding, a platform throughput of 360 Gbps was achieved with 256-byte packets.

Tests were also performed which make use of Dynamic Device Personalization (DDP) on the Intel 700 series Ethernet controllers which allows reconfiguration of the NICs internal packet filtering and processing pipeline on-demand to meet specific use case needs. This is achieved by adding new configuration profiles to the Ethernet controllers at runtime without resetting or rebooting the server using the Data Plane Development Kit (DPDK).


These configuration profiles provide the ability to classify packet types and protocols inline and distribute the packets to specific queues on the device host interface. These Queues can be assigned to specific cores on the CPU. This enables application performance scale with the number of cores in the system improving overall performance and utilization. It also improves latency because the packets end up automatically on the right worker core. For example in a BNG application it is vital for performance that all the packet flows from the same CPE are handled by the same core to effectively enable technologies like billing & HQoS.

Server Hardware Requirements


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™ DC 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:



Base node
Value/performance-optimized configuration suited to deployments further from the network core.



Plus node
Tailored for the highest performance and highest density for maximum networking density and packet throughput per NUMA node.

In addition, the reference architecture specifies system configuration parameters for the Intel Select Solutions for NFVI Forwarding Platform controller node.



Controller node
Lower performance CPU and NICs able to manage and communicate with all cloud nodes.

The Advantech SKY-8101D is a verified Intel® Select Solutions for NFVI FP and is designed for high throughput across the CommSP Network Infrastructure. It is built around the architectural and deterministic performance features of 2nd Generation Intel® Xeon® Gold processors.

The system is made up of the following key components and technologies:

- **Balanced non-uniform memory access (NUMA) connectivity:** I/O is evenly distributed across both CPU sockets. Workloads running on both sockets can benefit from

direct access to network, storage (Non-Volatile Memory Express*, or NVMe* drives), and Intel® QuickAssist Technology (Intel® QAT) acceleration.

- **High throughput network interface cards (NICs):** The SKY-8101D offers four NUMA balanced PCIe x16 Gen3 slots. For Plus node validation four Intel® Ethernet Network Adapters with quad 25GbE ports were installed in the system. (XXV710 SPF 28+ quad-port @ 25 Gbps Intel® Ethernet Network Adapter XXV710).

Standards-based networking performance across NFVI workloads is provided by the adaptors 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:

- **Data Plane Development Kit (DPDK)** to accelerate packet processing in the data path. It also facilitates building packet forwarders designed to operate on general-purpose, standards-based servers.
- **Dynamic Device Personalization (DDP)** is a programmable packet-processing pipeline provided by the Intel® Ethernet 700 Series that supports on-demand reconfiguration of network controllers at runtime, enabling workload-specific optimizations to increase throughput and decrease latency.



Figure 1: Advantech SKY-8101D front and rear views

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).

- **Storage:** 2 x 480 GB Intel® SATA Solid State Drives are configured as boot drives.

Advantech SKY-8101D - Intel® Select Solution for NFVI FP Plus Configuration

The Intel® Select Solution for NFVI FP plus node is based on the Intel® Xeon® Gold 6252 CPU @ 2.1 GHz. This hardware configuration is tailored for the highest performance and highest density for maximum networking density and packet throughput per NUMA node.

The Advantech SKY-8101D passed Plus configuration tests and was validated with dual Intel® Xeon® Gold 6252 processors (24 cores each, 2.1GHz). The SKY-8101D has been designed for maximum performance, scalability and functionality in a 1U rack mount footprint. In addition to the configuration verified to meet Intel's reference benchmark-performance

threshold the server is also available with a broader choice of processors from the 2nd generation Intel® Xeon® Scalable family.

The SKY-8101D is a high-end server optimized for computing power, accelerated workloads and high speed, high density I/O with optimum energy efficiency.

Two Intel® Xeon® Scalable processors provide the latest architectural enhancements, including rebalanced cache hierarchy, and Intel® Ultra Path Interconnect (Intel® UPI) for increased bandwidth and transfer rates between sockets at up to 10.4GT/s.

In addition, the new Intel® Advanced Vector Extensions 512 (Intel® AVX-512) Vector Neural Network Instruction (VNNI) extension increases the throughput of tight inner convolutional loop operations, reduces the memory bandwidth required to perform deep learning operations and will improve the performance of image matching algorithms on Advantech white box servers.

Each socket supports 6 memory channels and up to 12 DDR4 RDIMMs at 2666 MHz for up to 1536GB of ECC memory using the latest technology. Advanced RAS modes such as mirroring and sparing increase platform reliability.

The SKY-8101D's thermal system design enables support for processors with up to 165W TDP. This allows the appliance to scale from 8 core CPUs

to the highest performance 28 core processors available today.

With an abundance of PCI Express lanes, the SKY-8101D can support up to four full height 3/4 length (10.5") PCIe x16 adapters for modular, configurable networking I/O and acceleration. PCIe Gen3 technology on all slots provides sufficient bandwidth to support multiple 40GbE 25GbE and 10GbE NICs as well as the latest adapters offering 100GbE connectivity. With integrated security and compression acceleration based on Intel® QuickAssist Technology and two 10GbE ports with SR-IOV and RDMA support, the system offers best-in-class integration in a 1RU form factor.

Advanced Lights Out Management based on Advantech code base BMC and IPMI suite improves system manageability and reliability, providing platform thermal management, H/W monitoring and supervision. Remote firmware upgrade capability and hardware-based BIOS redundancy make the SKY-8101D an ideal platform for mission-critical and highly available networks.

Redundant power supplies, the ability to withstand single fan failures, redundant firmware images with failsafe upgrades and hot swappable FRUs make the SKY-8101D the platform of choice for applications requiring zero downtime.

The SKY-8101D is CE, FCC, UL, CB, CCC, and RoHS compliant.

Table 1: Checklist for Intel® Select Solution for NFVI FP Cloud Node Plus HW Configuration VS SKY-8101D

	Requirement	Required/ Recommended	Qty/ node	Advantech SKY-8101D
CPU	Intel® Xeon® Gold 6252 processor @ 2.1 GHz or Intel® Xeon® 6252N processor @ 2.3 GHz 24C/48T, 150 W or higher number Intel® Xeon® Gold/Platinum processor SKU	Required	2	2x Intel® Xeon® Gold 6252 processor @ 2.1GHz 24C/48T
Memory	Option 1: DRAM only configuration: 384GB (12 x 32 GB DDR4 2666 MHz) Option 2: DRAM only configuration: 384GB (24 x 16 GB DDR4 2666 MHz) Option 3: DRAM + Intel® Optane™ DC Persistent Memory 192GB (12 x 16 GB DDR4 2666 MHz)	Required	12 24 12 + Intel® Optane™	384 GB DRAM (12 x 32 GB 2666 MHz 288-pin DDR4 RDIMM)
NIC	4 x 25GbE quad port Intel® Ethernet Network Adapter XXV710 SFP28+. 4 (2 per NUMA node)	Required	4	4 x 25GbE quad port Intel® Ethernet Network Adapter XXV710 SFP28+. 4 (2 per NUMA node)
Storage	2 x 480 GB Intel® SATA Solid State Drive or Equivalent boot drive	Required	2	2 x 480 GB Intel® SATA Solid State Drives
LAN on Motherboard (LOM)	10Gbps or 25Gbps port for Pre-boot Execution Environment (PXE) and Operation, Administration and Management (OAM)	Required	2	2x 10Gbps with Intel® Ethernet Network Adapter X722 (Integrated into PCH)
	1/10Gbps port for Management NIC	Required	1	2x 1Gbps with Intel® Ethernet Controller I210-AT

Table 1 shows the exact configuration of the SKY-8101D and compares it to the reference specifications.

Enhanced Features

Both of Advantech's SKY servers come with an enhanced feature set to improve availability, serviceability and usability:

- Remote Intelligent Platform Monitoring & Control
 - Integrated IPMI Based Management Controller
 - Development, Customization, Validation and Life Cycle Management
 - Standard and Advanced IPMI Features
- Redundant BIOS
 - Physical Redundant Flashes for Current/Backup BIOS
 - Watchdog Mechanism to Detect Failing / Corrupted BIOS
 - Rollback Mechanism for System Recovery if BIOS Upgrade Fails
 - Dedicated Update Utility (ABU)
- Remote BMC/BIOS Upgrade

- x86 BIOS Upgradable By BMC and ABU (Advantech BIOS Utility)
- Industry Standard HPM.1 Protocol

Remote Evaluation Service

Advantech's unique Remote Evaluation Service (RES) offers developers easy and secure access to an entire range of platforms upon which they can rapidly evaluate Advantech value-add and test new services. In concert with other Intel® Network Builders ecosystem members, Advantech enables developers with early access to the latest technology, which accelerates their next generation product designs. As a result, they can apply innovative new technology sooner to reduce operating expense and grow new revenue faster. RES offers an evaluation framework that brings together members of the Intel® Network Builders community who share similar philosophies about telecom and edge cloud architecture and where they can openly collaborate together on a range of platforms from two Intel Atom®

processor cores to several hundred Intel® Xeon® processor cores.

With RES, developers can get ahead of the curve and begin to test different NFV infrastructures on platforms destined for deployment closer to the subscriber in the access network, mobile edge and in customer premises (uCPE) as well as the network core and telecom data center.

For more information on how to access RES for an evaluation of the Advantech verified Intel® Select Solutions for NFVI FP, or to order a platform:

Email: Cloud.IoT@advantech.com

Advantech's 2nd Generation Intel® Xeon® Scalable landing page including Forwarding Platform information:
www.advantech.com/nc/Spotlight/Intel-NFVI-FP

For more information on Advantech solutions please contact us at:

Email:

Cloud.IoT@advantech.com

Or visit

www.advantech.com/nc

Advantech Contact Information

Hotline Europe: 00-800-248-080 | Hotline USA: 1-800-866-6008

Email: NCG@advantech.com

Regional phone numbers can be found on our website at <http://www.advantech.com/contact/>

www.advantech.com/nc

Intel, the Intel logo, Intel Atom, Intel Optane, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

All other trademarks are property of their respective owners

© 2019 Advantech Co Ltd 0521F