

SOLUTION DESCRIPTION

Network Offload

FOR INTEL-BASED INFRASTRUCTURE
PROCESSING UNITS (IPUs)

SOLUTIONS

Telecom
Cloud and Edge
Cybersecurity
Monitoring
Infrastructure
Financial

PLATFORMS


Link-Capture™ Software
Link-Inline™ Software
Link-Virtualization™ Software
Link-Storage™ Software
Link-Security™ Software
Link-Programmable™
Smart Network Interface Cards (SmartNICs)
Infrastructure Processing Units (IPUs)

APPLICATIONS

5G UPF Offload
vSwitch Offload
Network Offload
Security Offload
Storage Offload
Data Capture
+ More

SERVICES

Professional Services
Custom Development



NAPATECH F2070X IPU

Napatech Infrastructure Processing Unit (IPU) solution maximizes the performance of data center networking for Infrastructure-as-a-Service platforms

Operators of enterprise and cloud data centers are continually challenged to maximize the compute performance and data security available to tenant applications, while at the same time minimizing the overall CAPEX, OPEX and energy consumption of their Infrastructure-as-a-Service (IaaS) platforms. However traditional data center networking infrastructure based around standard or “foundational” Network Interface Cards (NICs) imposes constraints on both performance and security by running the networking stack on the host server CPU, as well as related services like the hypervisor.

This solution brief explains how an integrated hardware-plus-software solution from Napatech addresses this problem by offloading the networking stack from the host CPU to an Infrastructure Processing Unit (IPU) while maintaining full software compatibility at the application level.


The solution not only frees up host CPU cores which would otherwise be consumed by networking functions but also delivers significantly higher data plane performance than software-based networking, achieving a level of performance that would otherwise require more expensive servers with higher-end CPUs. This significantly reduces data center CAPEX, OPEX and energy consumption.

The IPU-based architecture also introduces a security layer into the system, isolating system-level infrastructure processing from applications and revenue-producing workloads. This increases protection against cyber-attacks, which reduces the likelihood of the data center suffering security breaches and high-value customer data being compromised, while allowing customers full control of host system resources.

Finally, offloading the infrastructure services such as the networking stack and the hypervisor to an IPU allows data center operators to achieve the deployment agility and scalability normally associated with Virtual Server Instances (VSIs), while also ensuring the cost, energy and security benefits mentioned above.

Napatech F2070X IPU

The Napatech F2070X Infrastructure Processing Unit (IPU) is a 2x100G PCI Express (PCIe) card with an Intel® Agilex® F-Series FPGA and an Intel® Xeon® D processor, in a Full Height, Half Length (FHHL), dual-slot form factor.



The standard configuration of the F2070X IPU comprises an Agilex AGF023 FPGA with 4 banks of 4GB DDR4 memory, together with a 2.3GHz Xeon® D-1736 SoC with 2 banks of 8GB DDR4 memory, while other configuration options can be delivered to support specific workloads.

The F2070X IPU connects to the host via a PCIe 4.0 x16 (16 GT/s) interface, with an identical PCIe 4.0 x16 (16 GT/s) interface between the FPGA and the processor.

Two front-panel QSFP28/56 network interfaces support network configurations of

- 2x 100G;
- 8x 10G or 8x 25G (using breakout cables).

Optional time synchronization is provided by a dedicated PTP RJ45 port, with an external SMA-F and internal MCX-F connector. IEEE 1588v2 time-stamping is supported.

Board management is provided by a dedicated RJ45 Ethernet connector. Secure FPGA image updates enable new functions to be added, or existing features updated, after the IPU has been deployed.

The processor runs Fedora Linux, with a UEFI BIOS, PXE boot support and full shell access via SSH and a UART.

Foundational NICs constrain the capabilities of Infrastructure-as-a-Service platforms

Operators of enterprise and cloud data centers seeking to deliver Infrastructure-as-a-Service (IaaS) platforms to their customers or “tenants” have traditionally been forced to select between two architecture approaches.

By dedicating an entire physical server to a single tenant, they can guarantee the maximum compute performance available from that server’s CPU, together with fully-deterministic operation and full isolation from all other data center tenants. However it can take the operator several hours to provision a server for this true “bare metal” experience, during which time the server is unused and generating no revenue. The physical provisioning process also limits the operator’s ability to respond quickly to new customer orders or dynamically-changing requirements, while the scalability of the customer’s configuration is inherently limited. All-in-all, using physical servers to deliver IaaS platforms results in a less-than-ideal OPEX model for the operator.

Standard server running Virtual Server Instance


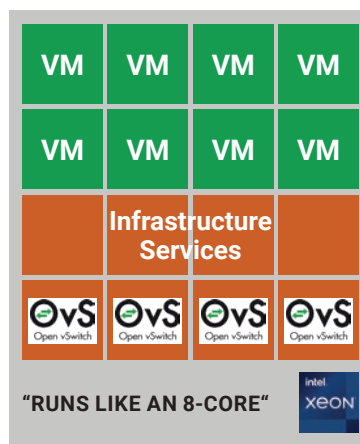
The traditional alternative is to provision the platform as a “Virtual Server Instance” or “VSI” running on a standard server configured with a foundational Network Interface Card (NIC), so that customers access Virtual Machines (VMs) rather than physical CPU cores, under the control of a hypervisor. (Typically, a “Type 1” bare metal hypervisor is used for IaaS applications, rather than a “Type 2” hypervisor which runs on top of an operating system.)

The advantages of a VSI deployment are that provisioning is quick (typically only minutes or even seconds) which minimizes server downtime, while the platform is highly scalable and customizable as customers’ needs change.

There are however some important downsides to the VSI approach. A significant fraction of the CPU cores are required for running infrastructure services such as the networking stack, the virtual switch (vSwitch) and other functions such as the hypervisor, which limits the number of CPU cores can be monetized for VMs. A 16-core CPU, for example, might only deliver the performance of an 8-core.

A vSwitch connects VMs or containers with both virtual and physical networks, while also connecting “East-West” traffic between the VMs or containers themselves. Software-based vSwitches are available, with Open vSwitch - Data Plane Development Kit (OVS-DPDK) being the most widely used due to the optimized performance that it achieves through the use of DPDK software libraries as well as its availability as a standard open-source project.

Example 16-core CPU



Foundational NIC
Only performing networking

Also, since tenants are unable to access the whole CPU, this architecture does not constitute a true bare metal cloud. There is a security risk if either the hypervisor or vSwitch is compromised in a cyber-attack, while there is no way to ensure full isolation between tenant workloads.



Finally, this approach delivers limited networking performance since a foundational NIC lacks the capability to accelerate the stack by offloading it from the CPU.

IPU offload for high-performance bare metal clouds

The best approach for delivering IaaS is to use an Infrastructure Processing Unit (IPU) instead of a foundational NIC. An IPU comprises an FPGA to execute data plane functions like Layer 2 and 3 of the networking stack, as well as a general-purpose CPU for control plane workloads as well as Layer 4 through 7 of the stack and the hypervisor.

Infrastructure services such as the hypervisor, vSwitch and networking protocols can be completely offloaded to the SmartNIC, so that all the CPU cores are available (and monetizable) for running tenant VMs, guaranteeing maximum compute performance. In this example, the 16-core CPU truly runs like a 16-core.

Example 16-core CPU



IPU plus Napatech Software
IPU runs the infrastructure services
plus high-performing networking

This architecture enables the deployment of virtualized, multi-tenant bare metal clouds in support of an IaaS business model, with fully-deterministic performance for each tenant and no risk of jitter due to "noisy neighbors".

In a multi-tenant scenario, compute resources can be split and allocated dynamically based on the specific needs of each tenant, with a finer level of granularity and faster provisioning compared to VSIs.

At the same time, the IPU ensures maximum networking performance, by offloading all the necessary, protocols and security functions that were performed on the server CPU in the VSI scenario.


This IPU platform architecture supports a highly-agile deployment model that enables data center operators to maximize the ROI from their IaaS offerings. As in the VSI scenario, provisioning is quick (typically only minutes or even seconds) which minimizes server downtime, while the platform is highly scalable and customizable as customers' needs change. In addition, the infrastructure services are fully upgradable in software, so there is no risk of hardware lock-in, which maximizes the effective lifetime of both the server and the IPU.

Napatech network offload solution

Napatech provides an integrated, system-level solution for data center network offload, comprising the high-performance Link-Virtualization™ virtualized data plane software stack running on the F2070X IPU.

Link-Virtualization is fully compatible with the virtio-net networking interface and applicable industry standard APIs like DPDK, so no changes are required to guest applications in order to leverage the performance delivered by the Napatech solution.

See the sidebar for details of the F2070X IPU hardware.



The Link-Virtualization software incorporates a rich set of functions, including:

- Full offload of Open vSwitch (OVS) from the host to the IPU, with the OVS data plane running in the FPGA and the control plane on the Intel® Xeon® D SoC.
- VLAN and VxLAN encapsulation and decapsulation:
- L2 tunneling: IEEE 802.1Q VLAN and Q-in-Q (IEEE 802.1ad);
- L3 tunneling: VXLAN.
- Virtual data path acceleration (vDPA), enabling the live migration of guest (tenant) VMs while retaining the high performance and low latency of Single Root I/O Virtualization (SR-IOV).
- Port mirroring.
- RSS load balancing.
- Link aggregation.
- Port rate limiting for Quality of Service (QoS):
- Transmit rate limiting on a per Virtual Function VF or Physical Function (PF) basis;
- Receive rate limiting on a per VF basis.
- Security isolation layer between the host CPU and the IPU.


With no host CPU cores being utilized for the virtualized data plane, the platform enables the deployment of both single-tenant and multi-tenant clouds with full isolation, protecting and preserving the network management model and topology.

Since the F2070X is based on an FPGA and CPU rather than ASICs, the complete functionality of the platform can be updated after deployment, whether to modify an existing service, to add new functions or to fine-tune specific performance parameters. This reprogramming can be performed purely as a software upgrade within the existing server environment, with no need to disconnect, remove or replace any hardware.

Industry-leading performance

The Napatech F2070X-based storage offload solution delivers industry-leading performance on benchmarks relevant to data center use cases, including an aggregate switching capacity of 300 million packets per second (Mpps) in a port-to-port configuration with 64-byte packets (line rate).

Note that the above benchmarks are preliminary pending general availability of hardware and software.



Summary

Operators of enterprise and cloud data centers are continually challenged to maximize the compute performance and data security available to tenant applications, while at the same time minimizing the overall CAPEX, OPEX and energy consumption of their Infrastructure-as-a-Service (IaaS) platforms. However traditional data center networking infrastructure based around standard or “foundational” Network Interface Cards (NICs) imposes constraints on both performance and security by running the networking stack on the host server CPU, as well as related services like the hypervisor.

Napatech’s integrated hardware-plus-software solution, comprising the Link-Virtualization software stack running on the F2070X IPU, addresses this problem by offloading the networking stack from the host CPU to an Infrastructure Processing Unit (IPU) while maintaining full software compatibility at the application level.

The solution not only frees up host CPU cores which would otherwise be consumed by networking functions but also delivers significantly higher data plane performance than software-based networking, achieving a level of performance that would otherwise require more expensive servers with higher-end CPUs. This significantly reduces data center CAPEX, OPEX and energy consumption.

The IPU-based architecture also introduces a security layer into the system, isolating system-level infrastructure processing from applications and revenue-producing workloads. This increases protection against cyber-attacks, which reduces the likelihood of the data center suffering security breaches and high-value customer data being compromised, while allowing customers full control of host system resources.

Finally, offloading the infrastructure services such as the networking stack and the hypervisor to an IPU allows data center operators to achieve the deployment agility and scalability normally associated with Virtual Server Instances (VSIs), while also ensuring the cost, energy and security benefits mentioned above.

For more information, visit: Napatech at <https://www.napatech.com/support/resources/solution-descriptions/network-of-flood-solution/>

Napatech is the leading supplier of Smart Network Interface Card (SmartNIC) and Infrastructure Processing Unit (IPU) solutions used in cloud, enterprise, and telecom datacenters.

Through commercial-grade software suites integrated with high-performance hardware, Napatech accelerates network infrastructure and security workloads to deliver best-in-class system-level performance while maximizing the availability of server compute resources for applications and services.

Additional information is available at:
www.napatech.com

NAPATECH RECONFIGURABLE COMPUTING

EUROPE, MIDDLE EAST AND AFRICA

Napatech A/S
Copenhagen, Denmark
T +45 4596 1500
info@napatech.com
www.napatech.com

NORTH AMERICA

Napatech Inc.
Portsmouth, NH, USA
T +1 888 318 8288
info@napatech.com
www.napatech.com

ASIA-PACIFIC

ntapacsales@napatech.com
www.napatech.com