

napa:tech;
SMARTNIC+IPU SOLUTIONS

SOLUTION DESCRIPTION

Virtual switch offload solution for edge and cloud data centers

NAPATECH LINK-VIRTUALIZATION™ SOFTWARE

napa:tech;

SOLUTIONS

Telecom
Cloud and Edge
Cybersecurity
Monitoring
Infrastructure
Financial

PLATFORMS

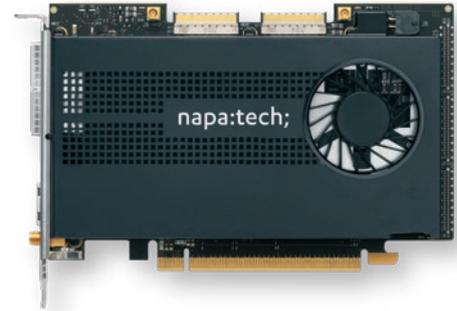
Link-Capture™ Software
Link-Inline™ Software
Link-Virtualization™ Software
Link-Storage™ Software
Link-Security™ Software
Link-Programmable™
Smart Network Interface Cards (SmartNICs)
Infrastructure Processing Units (IPUs)

APPLICATIONS

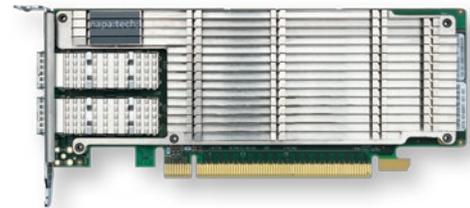
5G UPF Offload
vSwitch Offload
Network Offload
Security Offload
Storage Offload
Data Capture
+ More

SERVICES

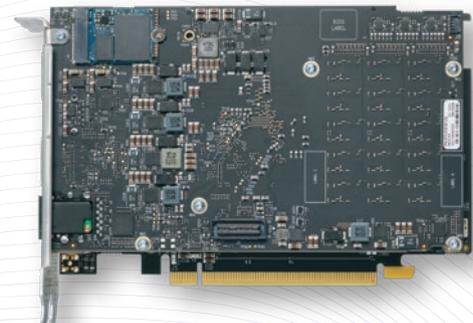
Professional Services
Custom Development



NAPATECH NT200A02 SMARTNIC



NAPATECH NT50B01 SMARTNIC



NAPATECH F2070X IPU

Napatech's virtual switch offload solution maximizes compute resource utilization and energy efficiency within edge and cloud data centers

Offloading virtual switching to a Smart Network Interface Card (SmartNIC) or Infrastructure Processing Unit (IPU) maximizes the Return on Investment (RoI) for edge and cloud data centers.

- Optimizes the allocation of host CPU cores for running applications and services instead of networking workloads;
- Full support for both FPGA-based SmartNICs and FPGA-based IPUs;
- Delivers best-in-class vSwitch performance to accelerate enterprise, cloud and telecom workloads.

High-bandwidth data center workloads drive the need for high-performance virtual switching

Within both edge and cloud data centers, virtualized workloads are becoming ever more complex, driven by **enterprise use cases** including Artificial Intelligence (AI), Machine Learning (ML), surveillance, Augmented Reality (AR) and industrial control, **consumer applications** like Massively Multiplayer Online Gaming (MMOG) and video streaming, and **telecom services** such as virtual RAN (vRAN) and 5G Core. At the same time, network throughput rates continue to increase thanks to the proliferation of fiber and 5G connectivity, as well as the growth in peer-to-peer traffic between data centers.

Both these trends drive an increasing requirement for a high-performance virtual switch (vSwitch), which connects Virtual Machines (VMs) or containers with both virtual and physical networks, while also connecting "East-West" traffic between the VMs or containers themselves.

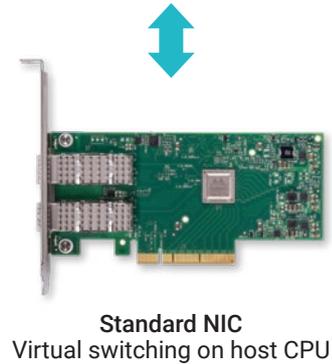
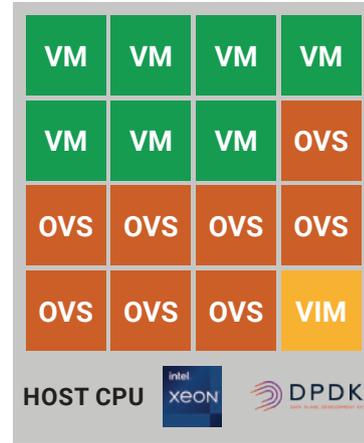
Software-based vSwitches are available, with OVS-DPDK being the most widely used due to the optimized performance that it achieves through the use of Data Plane Development Kit (DPDK) software libraries as well as its availability as a standard open-source project.



Open vSwitch constrains VM density

In a high bandwidth use case, however, OVS-DPDK consumes a significant number of CPU cores running on the host CPU. Additionally, at least one core is typically required for a Virtual Infrastructure Manager (VIM) such as OpenStack, Kubernetes or NETCONF.

This limits the number of cores that can be assigned to revenue-generating applications and services. This in turn impacts the effectiveness, profitability and energy efficiency of the data center, so there's a strong business incentive to minimize the number of cores consumed by the vSwitch in order to maximize the number of cores available for VMs, commonly referred to as the "VM density".

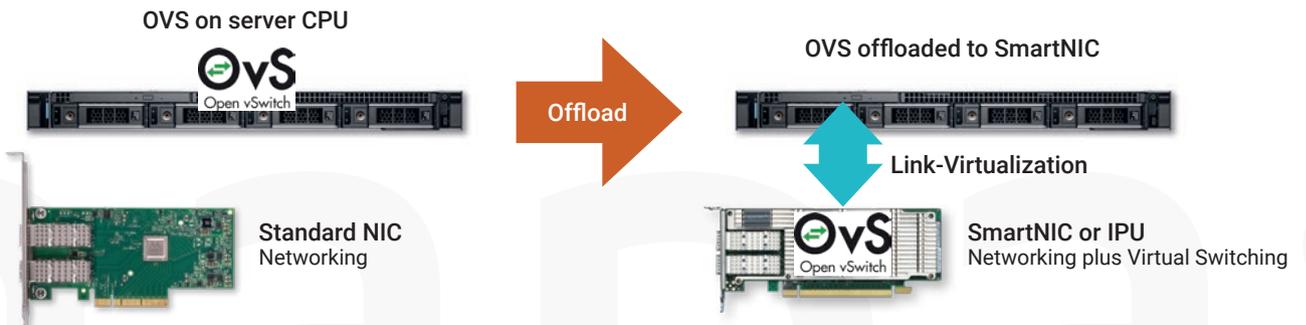


vSwitch offload frees up CPU resources

vSwitch offload solutions from Napatech address this challenge by offloading the vSwitch data path from the host CPU to a Field-Programmable Gate Array (FPGA) on a programmable Smart Network Interface Card (SmartNIC) with a standard PCI-Express (PCIe) card form-factor.

Napatech's Link-Virtualization software abstracts the FPGA "fast path" from the vSwitch API, while executing on the card all the high-performance networking functions required to implement a vSwitch that is 100% compatible with the relevant software APIs such as OVS-DPDK. Applications or services originally written to use OVS-DPDK can leverage the performance advantages of the Napatech vSwitch offload without any software changes.

Reclaiming host CPU cores previously required to run OVS and making them available to run applications and services leads to a significant reduction in the number of servers required to support a given workload or user base. This in turn drives significant reductions in overall server CAPEX and OPEX. It also results in lower system-level power consumption and improved energy efficiency for the edge or cloud data center. To aid in the estimation of cost and energy savings for specific use cases, Napatech provides an [online ROI calculator](#), which data center operators can use to analyze their projected savings.

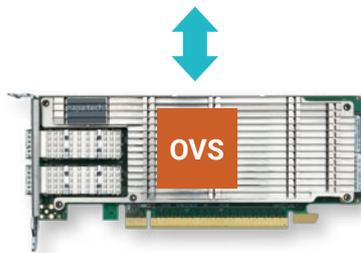


Match the hardware to the use case

By supporting two different types of hardware, the Link-Virtualization software provides users with the flexibility to select the platform that will result in the best match for the business goals of their use case.

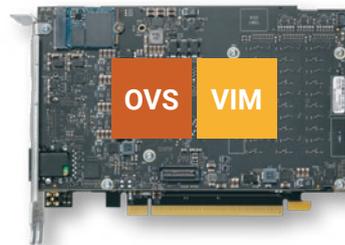
In the case of an FPGA-based programable SmartNIC such as the **NT200 or NT50 cards from Napatech**, the OVS-DPDK data plane runs in the FPGA, providing a fast path implementation that delivers the ultra-high throughput, low latency and deterministic vSwitch performance required to support the most demanding workloads. At the same time, the host server CPU runs the OVS-DPDK “slow path”, which classifies the initial packet in a flow before the remaining packets are processed in the FPGA, as well as running the VIM.

When an FPGA-based IPU is selected as the hardware platform, the OVS-DPDK data plane again runs on an FPGA, while the OVS-DPDK slow path and the VIM run on a general-purpose processor also located on the IPU. This 100% offload of the slow path and VIM ensures that all the host CPU cores are available for running virtualized services and applications, enabling the provisioning of “bare metal” clouds.



SmartNIC
OVS fast path offloaded to SmartNIC

This is the typical scenario for an **embedded telecom or industrial use case**, whether hosted at the network edge or in the core network, where optimized cost-performance is the primary objective.



IPU
OVS and VIM 100% offloaded to IPU

This is the optimum solution for an **IT use case** where compute resources must be allocated for use by multiple tenants, for example in support of an Infrastructure-as-a-Service (IaaS) business model. Tenant VMs have full access to the entire host CPU in order to obtain deterministic performance, optimized security and full isolation for their private workloads.

Bare metal configurations are widely used in applications such as public and private cloud data centers, financial services, healthcare and retail locations.

The above SmartNICs and IPUs are available from Napatech as integrated solutions that include the Link-Virtualization software. The FPGA on the SmartNIC or IPU ensures that 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.

Offloading the virtual switching function from the server's host CPU to a programmable SmartNIC running Napatech's Link-Virtualization software results in significant reductions in CAPEX, OPEX and power consumption for edge and cloud data centers.

napatech

Industry-leading performance

Leveraging Napatech's highly-optimized FPGA firmware, the Link-Virtualization software delivers an aggregate switching capacity of 129 million packets per second (Mpps) per core with 64-byte packets. This represents a 15x performance improvement compared to a pure software implementation of OVS-DPDK, which achieves 8Mpps per core with 64 byte packets. *(These results were measured on a Dell R740 server with 2.30GHz Intel® Xeon® Gold 6230 CPU and 128G DRAM, running CentOS 7.8 and DPDK 18.11.2, configured with a Napatech NT200A02 SmartNIC with dual 25Gbps ports, running Link Virtualization v4.0.6.)* In addition, Link-Virtualization implements the standard virtual Data Path Acceleration (vDPA) kernel framework, which enables full support for the live migration of guest workloads while retaining the high performance and low latency of Single Root I/O Virtualization (SR-IOV).

Summary

Offloading the virtual switching function from the server's host CPU to a programmable SmartNIC or IPU running Napatech's Link-Virtualization software results in significant reductions in CAPEX, OPEX and power consumption for edge and cloud data centers. This offload strategy addresses the dual problems of heightened workload complexity and increased network traffic for enterprise, consumer and telecom use cases.

For more information, visit: Napatech at <https://www.napatech.com/products/link-virtualization-software/>

napatech

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