

Enhancing Networks with SmartNICs

Defining Intelligent Interconnects and their Role in Today's Data Centers

OVERVIEW

In today's world of faster and more virtualized servers, storage, and network connections, CPUs cannot keep up with the increasing network processing demands. Legacy or foundational NICs may deliver efficient networking however when running demanding workloads, they cause overhead that burdens CPUs, chewing into available processing power. To deploy more advanced networking capabilities a new generation of intelligent or smart NICs are required to deliver accelerations and additional processing power to offload CPUs. SmartNICs can be based on an ASIC, FPGA or System on a Chip (SoC), and Mellanox offers all three: BlueField SmartNICs (SoC-based), ConnectX intelligent NICs (ASIC-based), and Innova (FPGA-based) SmartNICs. Learn the benefits of each and how to decide which is the best fit for your environment.

WHAT ARE SMARTNICs?

A traditional NIC, let's call it a "foundational NIC," performs basic network connectivity functions while performing rudimentary offloads like checksum and segmentation that would ordinarily be performed by the CPU. However, a trend towards cloud and software defined everything is straining current network bandwidth limits, driving the need for higher speeds. New complexities such as overlay tunneling and virtual switching all add networking and overhead that burden the CPU. Performing these functions in software creates a substantial load on CPU resources, leaving little CPU to run applications. While this software-based approach promises to deliver improved efficiencies, it also requires a more elaborate, feature rich and powerful NIC called a "Smart" or "Intelligent" NIC. These Intelligent or SmartNICs accelerate data path, virtualization, storage and security functions by completing the tasks in hardware, freeing CPU cycles for application workloads.

WHY ARE SMARTNICs NEEDED?

The complexities of server and storage-based networking are increasing, and the adage of solving problems by throwing CPU cores at them is no longer viable. A SmartNIC can be implemented in many flexible ways to perform complex actions that save CPU cycles.

HIGHLIGHTS

- Eliminates inefficiencies in the OS network stack and optimizes I/O processing
- Offload networking, security and storage tasks with hardware acceleration to realize huge efficiency gains
- Combine high-performance connectivity with programmability to facilitate fast feature innovation at a fraction of previous costs
- Reclaim CPU cores to allow more VMs, VNFs and workloads to be deployed per server



This includes data and control plane (networking, storage, and security) functions which can be fully implemented in, and offloaded by, a SmartNIC.

Let's take a deeper look at a SmartNIC roles in each — networking, storage and security. First, networking features such as overlay tunneling, load balancing and network monitoring and network virtualization services (NFV, vSwitch vRouting) are all able to be hosted on a SmartNIC. For storage, SmartNICs can be used to deliver performance for storage networking with RDMA and NVMe over Fabrics offloads, either integrated as flash storage controllers or within hyperconverged infrastructure to handle both storage and networking tasks. SmartNICs can also implement various types of security features that can operate independent of the host, like cryptography, deep packet inspection or blocking DDoS attacks before servers get overwhelmed. Implementation of a smart NIC to offload these tasks from the CPU allows for significantly better server scaling while greatly improving efficiencies.

CPU CORE WORKLOAD

As the shift towards network functions move from purpose-built hardware to software-driven and virtual appliances, we see that general-purpose CPUs are tasked with more specialized services. While CPUs are great for general application tasks, they are not specialized for handling network and storage traffic. Running these services on the CPU impacts overall application performance. Disaggregation of software and hardware, along with the virtualization of servers and networks, further adds to the load placed on CPUs and results in decreased efficiency. In figure 1 we see the results of this decreased efficiency as CPU cores are lost to perform virtualization services and security tasks. However, utilizing a SmartNIC to performing network and data path processing can boost performance by offloading CPUs in servers.

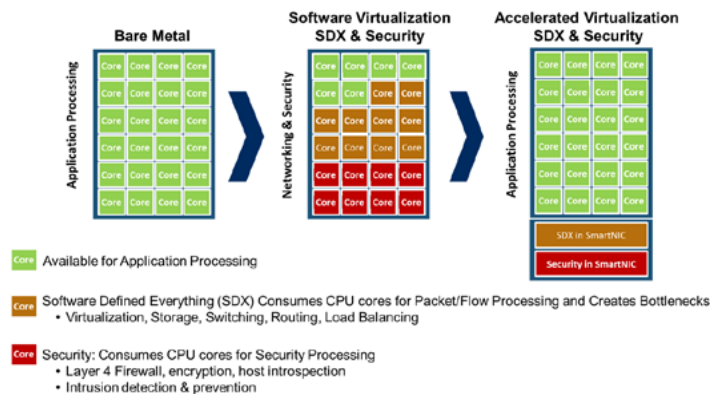


Figure 1. CPU core consumption with and without a SmartNIC.

ASIC, SOC, OR FPGA

As mentioned earlier, a smart NIC can be ASIC, FPGA, or System-on-a-Chip (SOC) based and there are specific advantages and tradeoffs to each. For the best performance, most acceleration functions must run in hardware, but for the greatest flexibility, the control and programming of these functions needs to run in software. If we look at price, programmability and flexibility for each it will help us to differentiate.

Starting with the Mellanox ASIC-based [ConnectX](#) Intelligent NICs, which are very cost effective and deliver the best price / performance. As an Intelligent NIC, it has built-in advanced hardware offloads allowing for extra flexibility, efficiency, and ultimate performance to outperform legacy or foundational NICs and many other “SmartNIC” designs. The [ConnectX-5](#) Intelligent NIC has a programmable data path that is relatively simple to configure, supporting a hardware-accelerated data plane that integrates with a software-based control and can offload overlay networks, RDMA, packet processing, storage tasks, big data, and software-defined networking (SDN), while lowering the burden on the host CPU.

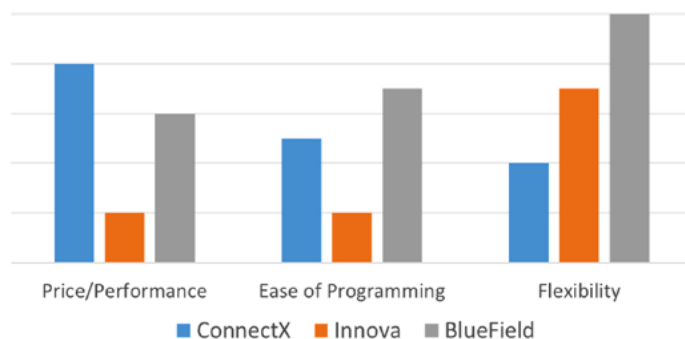


Figure 2. Characteristics of different SmartNIC types.

Mellanox [Innova-2](#), an FPGA-based SmartNIC by contrast is highly programmable and can be made to support almost any functionality, specifically suitable for companies with Verilog FPGA development skills and provides acceleration to a wide range of applications, including Security, Storage, Artificial Intelligence, Media & Entertainment, Cloud and more. The on-board FPGA provides users with the ability to create custom-made accelerations. The Mellanox [BlueField](#) SoC-based SmartNIC is a nice balance between the two, good price / performance, offering full programmability for accelerating a multitude of security, networking and storage applications and delivering optimization for control and data path operation. Advanced features include an embedded virtual switch with programmable ACL, transport offloads and stateless encapsulation and decapsulation of NVGRE, VXLAN, and MPLS overlay protocols and superior RDMA and GPUDirect® RDMA accelerators, including NVMe-oF.

HOW TO CHOOSE A SMARTNIC

To achieve maximum efficiency in software-define and cloud data centers, SmartNICs will be replacing foundational NICs but which type best fits your scenario? First, it's helpful to examine which functions are needed to support specific workloads? Ultimately the Mellanox ConnectX family of Intelligent NICs provides all the function, performance and flexibility to offload most transports, handle virtual switching and offers some predefined programmability. ConnectX can accelerate data plane tasks in hardware but runs the control plane in software. It does this at a fraction of the power and cost of competing SmartNICs.

Mellanox Innova FPGA SmartNIC is highly programmable can be defined to perform a wide variety of very specific tasks. The flexibility does come with additional costs, but when used in lieu of expensive CPUs the cost is easily recovered.

The Mellanox BlueField SmartNIC is highly customizable combining the ConnectX Intelligent NIC advanced offloads with 64-bit Arm® processing power to offer full programmability and capable of running Linux. This allows for simplified C-programming, permits the control-plane to be run on the adapter, and makes it easy to port application to the SmartNIC.

Mellanox Intelligent NICs and SmartNICs can accelerate and offload network traffic, virtualize storage resources, share GPUs over the network, support RDMA, and perform encryption, allowing Mellanox adapters to be a platform for innovation. This only scratches the surface of what a Smart NIC is capable of and with the ability to add intelligence into a network they can alleviate a bulk of network, storage and security data processing in high-bandwidth and compute-intensive applications, helping to boost performance and offering breakthrough efficiencies in data center and cloud-based computing.

About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables and silicon that accelerate application runtime and maximize business results for a wide range of markets including high-performance computing, enterprise data centers, Web 2.0, cloud, storage and financial services.

www.mellanox.com



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085

Tel: 408-970-3400 • Fax: 408-970-3403

www.mellanox.com