

High-Performance Hardware: Enhance Its Use in Software-Defined Networking

Software-defined networking (SDN), along with the power of virtualization, automation, orchestration, and other software-based techniques, promises to transform the internetworking industry. However, hardware still remains a big part of every network. Commercial off-the-shelf hardware, such as servers with x86 processors that use software-based forwarding capabilities or emergent white-box switches, will find a place in many use cases. But a more important consideration is the evolving role of high-performance hardware in SDN. As service providers optimize their networks to take full advantage of the very best tools in SDN, network functions virtualization (NFV), and other technologies, they need to strategically assess their hardware choices, based on functions and performance requirements and the intended business outcome for individual applications and services.

Low-cost, off-the-shelf hardware can support many standard applications. But high-performance hardware, such as core data-center switches that support more stringent I/O requirements, high throughput, and high reliability, is still crucial. Of course, every provider needs to drive network costs down, and using more mass-produced hardware can be one way to do that. But the SDN and NFV solutions that you deploy must satisfy the top requirements across applications, services, service provider architectures, and topologies, now and in the future.

Defining High-Performance Hardware

This white paper discusses high-performance hardware that falls within one of the following two general categories:

- Network hardware that uses purpose-built designs - These designs often involve specialized application-specific integrated circuits (ASICs) to achieve significantly higher performance than is possible (or economically feasible) using commercial, off-the-shelf servers based on state-of-the-art, x86-based, general-purpose processors.
- Network hardware that uses standard x86 servers, enhanced to provide high performance and predictable operation - A subset of this category is networking hardware that supports special software techniques that bypass hypervisors, virtualization environments, and operating systems. The result is much higher and more predictable performance than is possible with standard cloud-computing deployments.

Talk within the internetworking industry suggests that virtualized network functions can be deployed much like cloud-based applications. But there is a big difference between deploying network functions as software modules on x86 general-purpose servers and using a common cloud-computing model to implement network virtualization. Simply migrating existing network functions to general-purpose servers without due regard to all the network requirements leads to dramatically uneven and unpredictable performance. This outcome is unacceptable for subscribers and eliminates the provider's ability to provide service-level agreement (SLA) guarantees.

The unpredictable performance based on deployment of different types of virtual network functions (VNFs) was demonstrated by the European Telecommunications Standards Institute (ETSI) group responsible for NFV performance and portability testing and best practices.

Service provider networks differ from common cloud-computing environments in two important ways:

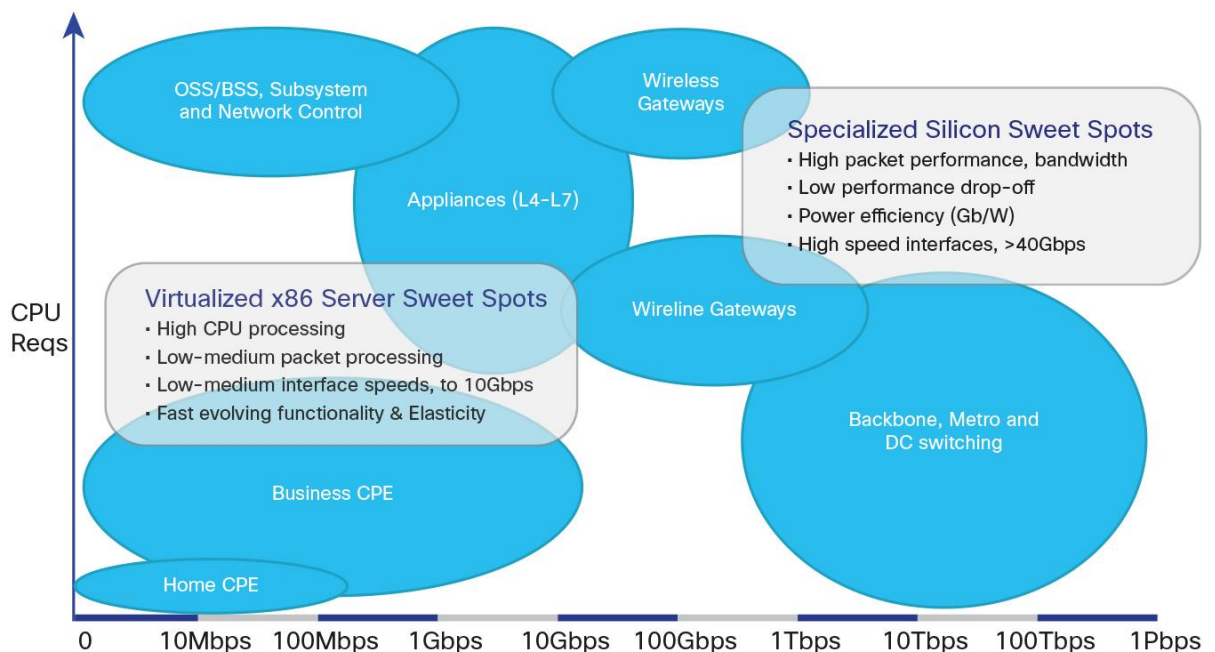
- Data-plane workloads need high and predictable performance.
- An end-to-end, global network management view is required across many and varied network shapes and hierarchies.

The unpredictability that was demonstrated is mainly due to data-plane workloads often being I/O bound, memory bound, or both and software layers containing important configuration details that may affect performance. Also, memory-intensive activities require equivalent chipset support for bare metal and full virtualization environments. So these issues are not specifically about hardware. Instead, they also concern how the overall environment is handled by software. Operating systems, hypervisors, and other infrastructure that is not integrated into best practices for data-plane applications will continue to contribute to unpredictable performance.

Assessing Hardware Requirements Based on Bandwidth and CPU Needs

A good way to begin your assessment of hardware requirements is to examine network functions in two dimensions: I/O bandwidth or throughput needs and computational power needs. When considering which network function to virtualize and where to virtualize it, examine the CPU load required and bandwidth load required throughout different layers of the network. This can help you determine that some, but not all, network functions are suitable for virtualization. In Figure 1, overlays in lighter green include specific hardware requirements in different parts of the network, arranged according to CPU and bandwidth requirements.

Figure 1. Assessing Hardware Requirements and Where to Use NFV



In Figure 1, the x-axis maps the I/O bandwidth performance requirements and the y-axis maps low to high CPU requirements. With current x86 processors, the functions that may work best on commercial, off-the-shelf x86 platforms require high CPU processing, low-to-medium packet processing, low-to-medium interface speeds (to 10 Gbps), and fast-evolving functionality and elasticity. These functions include: subsystem control (such as IP Multimedia Subsystem [IMS] applications), network control (for example, Domain Name System [DNS] applications), operations support systems (OSSs), business support systems (BSSs), many Layer 4-7 network appliances, and CPE.

The farther a function is to the right of Figure 1 the higher its I/O bandwidth and forwarding performance needs are. Functions with high I/O and forwarding performance needs are not good candidates for virtualization on standard x86 hardware. Instead, they require specialized, high-performance hardware or custom ASICs to achieve the desired performance or the optimal cost-performance trade-off. These functions depend on high packet performance, high bandwidth, low performance drop-off, power efficiency, and high-speed interfaces (40 Gbps and above). Examples include network backbone, metro, and data-center switching and high-performance network gateways.

Determining what type of hardware is appropriate for those functions that don't map towards the extremes of the chart requires a review of many factors. These include cost, quality of end-user experience, latency, networking performance, network predictability, and architectural preferences.

Figure 2 provides a view of applications that may work well in virtualized deployments on x86 servers. Applications to the left require more high-performance hardware, perhaps with custom ASICs. Applications to the right are good candidates for virtualized deployments on commercial, off-the-shelf hardware.

Figure 2. Applications Appropriate for Virtualized x86 Platforms

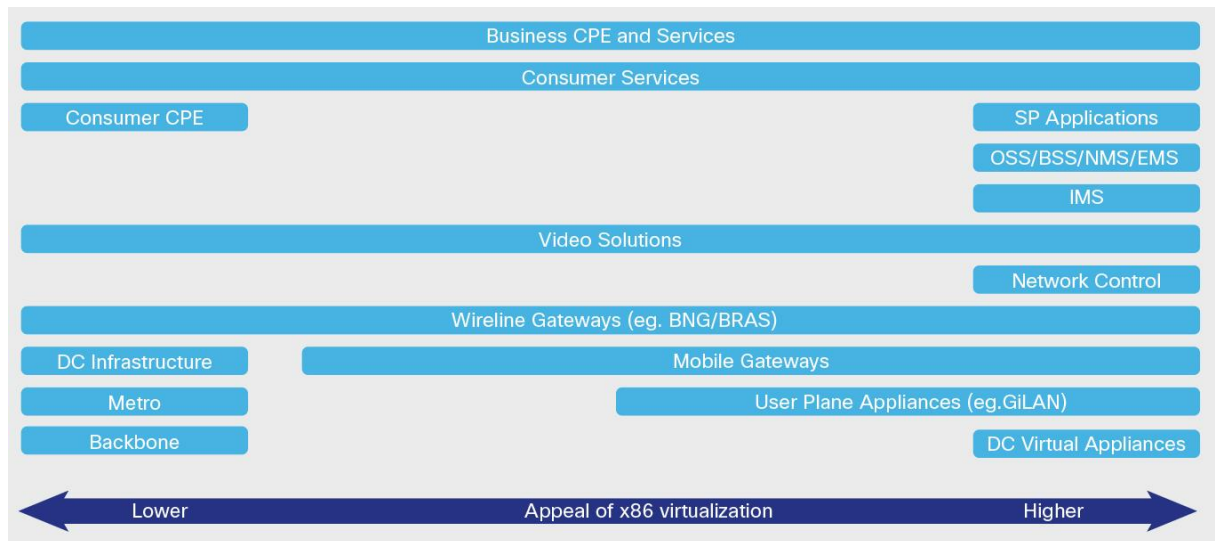


Figure 2 shows some obvious choices between high-performance and x86 hardware. However, the network shape, geography, architecture, and other factors may influence the final hardware choice. These factors include the population density of the area being served, the cost of power, and the cost of fiber. To address these factors, high-performance hardware can provide a larger footprint of your backbone to deliver high performance and low latency across a wide geographic area, fast line-rate forwarding, and a fast response to changing conditions.

If you're serving a smaller geography that is closer to your backbone switch, less high-performance forwarding hardware may be needed.

But many applications will still require highly predictable performance to keep your customers satisfied. Conditions vary tremendously for service providers in different regions and continents, and architectures and infrastructures vary to match. So it's also important to consider which architecture makes the most economic and market sense.

Hardware Requirements and Moore's Law

Remember, commercial off-the-shelf hardware and high-performance hardware and custom ASICs all follow Moore's Law: Their interface speeds continue to increase, and their systems grow smarter each year. Custom ASICs will continue to have higher performance than x86 or other platforms designed for general-purpose applications.

The Introduction and Scaling Question

You might not need high-performance hardware for certain functions initially. But what happens as a function or application scales? It might require a high-performance platform or be more economical on a purpose-built platform. So you can start out with commercial off-the-shelf hardware, then transfer the workload to the high performance hardware later. If you have focused on establishing a clean abstraction of the services from the underlying hardware infrastructure using SDN principles, it is easier to change or evolve your network deployment, keeping it independent of the upper services and applications. This is the true promise of SDN.

Along with an abstraction layer, an orchestration system is a vital component of your SDN environment. This system is used to manage across virtualized and physical infrastructures, and it can easily migrate workloads, when necessary, without exposing complexity to the upper layers of the services and applications. When there is a clean abstraction between orchestration, application, and infrastructure layers, different types of hardware can be switched quickly and easily as changing needs dictate.

For More Information

- [NFV Performance & Portability Best Practices, ETSI GS NFV-PER 001 \(Authentication required\)](#)
- [Network Optimization Through Virtualization: Where, When, What, and How?](#)
- [Evolved Programmable Network](#)
- [Cisco Evolved Services Platform](#)
- [SDN for Service Providers](#)
- [Network Functions Virtualization](#)



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)