

INDUSTRIALIZING NETWORK FUNCTIONS VIRTUALIZATION WITH SOFTWARE-DEFINED INFRASTRUCTURE

Telecommunications operators can deploy network functions virtualization and other general server workloads on the same platform with hyperscale, software-defined infrastructure from Ericsson, based on Intel® Rack Scale Design and powered by Intel® Xeon® Scalable Processor (Skylake).

Executive summary

In an effort to increase agility, decrease total cost of ownership (TCO), and cut complexity on the journey toward modernization, telecommunications operators are investing in network functions virtualization (NFV). This approach allows network functions to run on industry-standard, Intel® architecture-based servers instead of dedicated hardware. NFV holds tremendous promise in terms of cost savings and digital transformation, but organizations that want to deploy it broadly face several challenges. For example, they might have multiple NFV environments alongside legacy technologies, or they might need to support different versions of NFV environments. In addition, large organizations want to be able to support multiple cloud solutions with one hardware platform.

To resolve these challenges and get the full benefits of NFV, organizations need hyperscale, software-defined infrastructures (SDIs) that make it easy to deploy new resources by drawing from pooled compute, storage, and networking assets. Without such a hyperscale platform, NFV still requires distinct hardware islands for each virtualized network function (VNF) that requires specific optimized hardware, which makes it difficult to deploy NFV pervasively and at scale.

The Ericsson Hyperscale Datacenter System 8000 (HDS 8000), along with Ericsson software, provides a hyperscale NFV infrastructure (NFVi) that can support multiple telecom NFV environments, IT workloads running on VMware® or bare-metal solutions, and cloud environments on OpenStack—all on a single platform. Because the Ericsson HDS 8000 platform is based on Intel® Rack Scale Design, its hardware resources can be quickly allocated or re-allocated to automatically compose servers for different needs and different environments.

This document describes the Ericsson NFVi and shows how it can be used to industrialize NFV in telecom environments, large and small, regardless of network location.

Network complexity and the potential of pervasive NFV

Telecom network operators are under constant competitive pressure to deliver services with the speed, flexibility, and efficiency of the cloud, together with the deterministic behavior of telecom networks. Efficient rollout of new revenue-earning network services is critical in an increasingly network-connected world. Operators must do this despite complex network infrastructure where islands of fixed function, legacy hardware devices reside alongside newer infrastructure running NFV and cloud environments.

In an effort to increase agility and reduce costs and complexity, telecom service providers are virtualizing network functions. But virtualization alone can't completely solve the problems that come with isolated siloes of underlying physical infrastructure. For example, you can't quickly or easily adjust compute or memory on the servers hosting virtual machines (VMs) that run VNFs. In addition, most telecom network functions are designed for bare-metal solutions. Likewise, some IT workloads are not suited to virtualization because they require bare-metal deployments for maximum performance. While virtualization is a necessary step, something more is needed. Telecom service providers can

address these challenges with a hyperscale SDI that can support agile and widespread use of NFV and standard workloads alongside legacy hardware. One of the most comprehensive solutions is the Ericsson NFVi with Ericsson HDS 8000, which is based on Intel Rack Scale Design, a flexible, modular architecture that disaggregates data center resources for greater agility and efficiency. As a hyperscale SDI, the Ericsson HDS 8000 can run NFVi, bare-metal network functions, cloud environments such as OpenStack, or standard IT workloads on a single platform.

The Ericsson solution enables important technical and business benefits, such as:

- Cloud flexibility and agility
- Reduced equipment costs and power consumption
- Accelerated time to market for new services and products
- Increased operational efficiency
- Network appliance multi-tenancy, allowing use of a single platform for different applications, users, tenants, and cloud environments

These benefits are possible because the Ericsson NFVi, running on Ericsson HDS 8000, disaggregates compute, storage, and networking resources into virtual pools that can be flexibly reallocated.

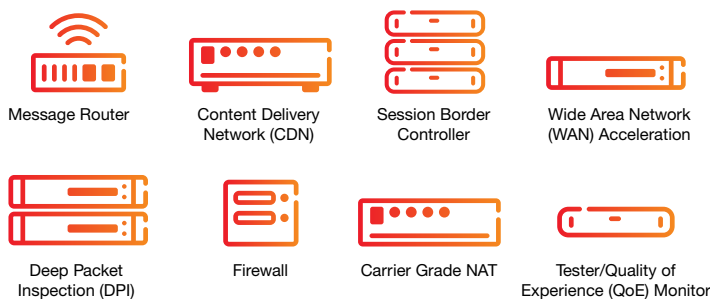
NFV and IT workloads running on the Ericsson solution can be optimized for efficiency and performance, with automated deployment and orchestration, all managed from a single pane of glass.

Organizations that deploy the solution today can enjoy the hyperscale and SDI benefits while laying a foundation that lets them take advantage of future capabilities enabled by the ongoing development and evolution of the Ericsson solution. This paper discusses both current capabilities and those that are coming in the near future to show a rounded picture of the impact that the Ericsson solution can have in telecom deployments.

NFV refresher

NFV is the virtualization of network functions that were historically performed by dedicated hardware. This new approach eliminates the need for proprietary network services devices because it decouples network functions from the underlying hardware so that the functions can be hosted on VMs running on industry standard servers based on Intel architecture. The goal of NFV is to transform the way that network operators architect networks by allowing consolidation of multiple network equipment types onto industry standards-based platforms (see Figure 1).

Traditional Network-Appliance Approach



- Fragmented non-commodity hardware
- Physical install per appliance per site
- Large hardware-development barrier to entry for new vendors, constraining innovation and competition

Network-Virtualization Approach




Figure 1. NFV allows for consolidation of multiple dedicated network appliances onto industry-standard, Intel® architecture-based servers, storage, and switches




Figure 2. Models for deploying VNFs rise in complexity and TCO as they involve more vendors

TCO rises with increased complexity

An NFV specification is developed and maintained by the ETSI Industry Specification Group (ETSI ISG), which is “working intensely to develop the required standards for NFV as well as sharing their experiences of NFV implementation and testing.”¹¹ (For more information about NFV and the work of the ETSI ISG, visit www.etsi.org/technologies-clusters/technologies/nfv.)

Examples specific to telecom operators include WAN acceleration, voice over Long-Term Evolution (VoLTE), user data consolidation, and carrier-grade network address translation (NAT).

Multiple NFV models

NFV is gaining momentum in the communications-service-provider community, and it will likely transform the way that networks are built and operated. However, the technology introduces new challenges to service providers, who must make the right decisions to build in the benefits of NFV architecture. Adoption often involves working with multiple vendors to service the hardware layer, the software layer, the virtualization layer, and the VNFs themselves. The choices that companies make in selecting technologies for each layer are important because they impact the NFV-transformation journey in several ways, such as system-integration costs, performance, operational complexity, and lifecycle management. These choices result in various levels of TCO benefits or risks.

The availability of a wide variety of products from multiple vendors at each level in the architecture results in several deployment models, as shown in Figure 2. The least complex model is one in which a single vendor provides the hardware, virtualization software, and VNFs. While less complex, this model eliminates the ability of organizations to choose best-of-breed technologies. The most complex model is one in which many different vendors provide these three components. This model offers the most choice, but system integration and overall management become very challenging. Most organizations that are on the NFV journey are somewhere between these two poles and have virtualized a small number of network functions using solutions from a small number of vendors. Telecom operators face additional challenges as they seek to expand their use of NFV. The technologies that they choose need to support high throughput and low latency; they must deliver performance and scalability, and they need to co-exist with legacy applications in a hybrid environment that includes physical and virtual appliances. The entire stack, and the VNFs running on it, must be orchestrated efficiently and consistently with as much automation as possible. Additionally, the NFV technologies should provide a migration path from today’s solutions—based on proprietary physical network appliances—to more virtual network-appliance solutions based on open standards.

The Ericsson NFVi, based on Ericsson HDS 8000 with Intel Rack Scale Design, can address these challenges and remove barriers to NFV adoption for telecom operators. For example, organizations that have already virtualized one or more network functions can continue their journeys by consolidating those multiple NFV environments onto a single platform based on open-source software and underpinned by the Ericsson HDS 8000. They can then deploy more VNFs over time onto the same highly scalable, future-proof platform, and they can even use multiple deployment models on the same data center infrastructure.

Why Intel® Rack Scale Design?

Telecom operators face significant challenges as they try to expand their use of NFV and meet demands for continuous connectivity. Due to rapid growth in the number of clouds, services, and devices, operators and cloud service providers are experiencing a workload explosion even as they push to reduce expenses, become more agile, optimize resource usage, and scale capacity without interruption. For operators who wish to deploy VNFs broadly across their networks, performance, resource optimization and scalability are vital.

Intel is answering these challenges with Intel Rack Scale Design. Data centers implementing solutions that conform to Intel Rack Scale Design significantly improve efficiency and enable rapid service provisioning,

thereby helping organizations gain efficiency, flexibility, and agility. Furthermore, a standards-based solution allows communications service providers to move away from inflexible, closed, and non-interoperable systems that are locked to specific vendor hardware.

Intel Rack Scale Design is:

1. A set of industry-standard API specifications and open-source reference software that allows you to pool resources
2. A reference architecture for a modern rack-architecture design that original equipment manufacturers (OEMs) and infrastructure vendors like Ericsson can implement to create SDI solutions for enterprise companies, telecoms, and cloud service providers

The Intel Rack Scale Design APIs are based on and extend DMTF's Redfish™ open industry-standard specification and schema. The schema specifies a RESTful interface and utilizes JSON and OData to help customers integrate solutions within their existing data center environments.

By implementing Intel Rack Scale Design APIs and pooling data center resources, organizations can flexibly compose their hardware into specialized systems to address the needs of diverse workloads and applications. The flexibility and software-based interfaces allow the composed systems to scale and handle rapid changes to use cases and workloads.

Another benefit of pooled system resources is that these can be replaced individually when implemented in separated replacement units. This means that,

instead of replacing or upgrading an entire server, switch, or storage pool, organizations can replace individual resource units, which helps control costs while allowing greater freedom to optimize resources. A pooled system-management engine, which is part of Intel Rack Scale Design, can use a mix of integrated pooled resources and disaggregated pools of resources contained in separated resource units for the composed systems that are used by applications.

Overview of the Ericsson NFVi

The Ericsson HDS 8000 that runs the NFVi solution is founded on Intel Rack Scale Design and Intel technologies, and it comprises both software and hardware components. Whenever possible, the solution uses open APIs between components, making it easy to customize the solution to meet unique needs. Solution components include the following:

- **Ericsson Cloud Execution Environment** is the NFVi's virtualization, control, and management layer, and it fills the role of virtual-infrastructure manager as defined by the ETSI NFV model. As a cloud platform, Ericsson Cloud Execution Environment ensures that several applications can share the infrastructure's compute, storage, and networking resources even when, as with VNFs, the applications were not designed to run in a cloud environment. Ericsson Cloud Execution Environment is optimized to run VNFs through custom configurations, such as a virtual switch, which is accelerated by the Data Plane Development

Kit (DPDK). Additionally, Ericsson Cloud Execution Environment includes telecom-centric capabilities that are not included in regular OpenStack distributions, such as virtual-LAN (VLAN) trunking and continuous monitoring high availability, which enables automated migration of a VM to a functioning host in the event of host failure. Ericsson Cloud Execution Environment is based on open-source components: Linux, KVM for virtualization, and OpenStack for multi-virtual infrastructure manager (VIM) support.

- **Ericsson Cloud SDN** is a network-virtualization solution that provides seamless intra-data center and inter-data center connectivity for virtual, physical, and container-based workloads. It enables flexible and highly automated connectivity management, allowing use cases such as service chaining and flexible intra-data center networking. The embedded software-defined networking (SDN) function of Ericsson Cloud Execution Environment can also be easily integrated with external SDN functions to enable automation and central governance of end-to-end external networking use cases. It uniquely combines an industrialized OpenDaylight controller with advanced routing capabilities in a community-driven solution. Ericsson Cloud SDN integrates with Ericsson Cloud Execution Environment to control the DPDK-accelerated cloud SDN switch running inside it, in addition to physical network equipment such as the top-of-rack (TOR) switch and data center gateway.

WHAT IS DATA PLANE DEVELOPMENT KIT?

The Data Plane Development Kit (DPDK) is a set of software libraries for accelerating packet processing workloads on commodity off-the-shelf hardware platforms. It greatly boosts packet processing performance and throughput, allowing more time for data-plane applications.

DPDK can improve packet processing performance by up to 10 times. It's possible to achieve more than 175 megabit-per-second (Mbps) throughput on a single Intel® Xeon® processor, and double that with a dual-processor configuration.² As a result, network equipment providers (NEPs) can lower development costs, use fewer tools and support teams, and get to market faster. Visit <http://dpdk.org/> for more information.

- **Ericsson Cloud Manager** provides cloud management and orchestration capabilities and fulfills the roles of NFV orchestrator and VNF manager in the ETSI NFV management and orchestration (MANO) specification. In addition, it also provides service assurance and analytics, plus multi-tenant, multi-site, and multi-VIM support. Ericsson Cloud Manager is capable of managing and orchestrating infrastructure across many physical data centers. It combines the simplicity and flexibility of IT activities with the scale of telecom operations to configure, coordinate, and manage VNF resources and associated services across highly distributed cloud environments. This is made possible by a closed-loop orchestration that flexibly adapts to the changing environment. With Ericsson Cloud Manager, service providers can manage the lifecycles of services and resources as the underlying VNF capabilities evolve. The solution's comprehensive workflow-automation engine executes both predefined and user-defined workflows. By having a flexible catalog driving its workflow engine, Ericsson Cloud Manager can enforce the consistent execution of workflows within and across domains to expedite the rollout of new products, services, and VNFs.

- **Ericsson Command Center** provides hardware-infrastructure management. It configures and manages the compute resources, storage capacity, and network connectivity of both Ericsson HDS 8000 and systems from other vendors. It also manages the power systems, firmware, and related functions of Ericsson HDS 8000. Whether organizations use Ericsson Command Center to manage a standalone system or an entire data center infrastructure, a single configurable dashboard provides a “single-pane-of-glass” view of the current state of the hyperscale system and the rest of the data center. With what administrators learn about the infrastructure, they can analyze needs and plan operations down to the component level to better match services to the most appropriate hardware.
- Ericsson HDS 8000 is the hardware component of the Ericsson NFVi described in this paper.

- Enables a multitude of different applications and virtualization solutions hosted within the same physical infrastructure
- Is suitable for high-bandwidth and processing-intensive applications

With the Ericsson HDS 8000, compute, storage, and networking resources are physically assembled into performance-optimized data centers (PODs). PODs are configured and controlled through management software, which enables the platform to be a software-defined, cloud-based infrastructure. A POD physically consists of compute, storage, and network resources, provided as sled or rack units:

- Compute sled units (CSUs) or compute rack units (CRUs) with Intel Xeon processors, random-access memory (RAM), and Intel® Ethernet Adapters
- Storage sled units (SSUs) and storage rack units (SRUs) with hard-disk drives (HDDs) or SSDs
- Networking sled units (NSU) or networking rack units (NRU) with Intel Ethernet Adapters and controllers
- Equipment-access switches (EAS)
- Chassis that provide physical enclosures with fans for the CSUs and SSUs

Ericsson HDS 8000

Ericsson HDS 8000 is the hardware backbone of the NFVi. It is based on Intel Rack Scale Design and servers built with Intel® Xeon® processors, Intel® Solid-State Drives (SSDs) Data Center P3700 Series or Intel SSD DC P4600 Series, and Intel® Network Adapters. Ericsson HDS 8000:³

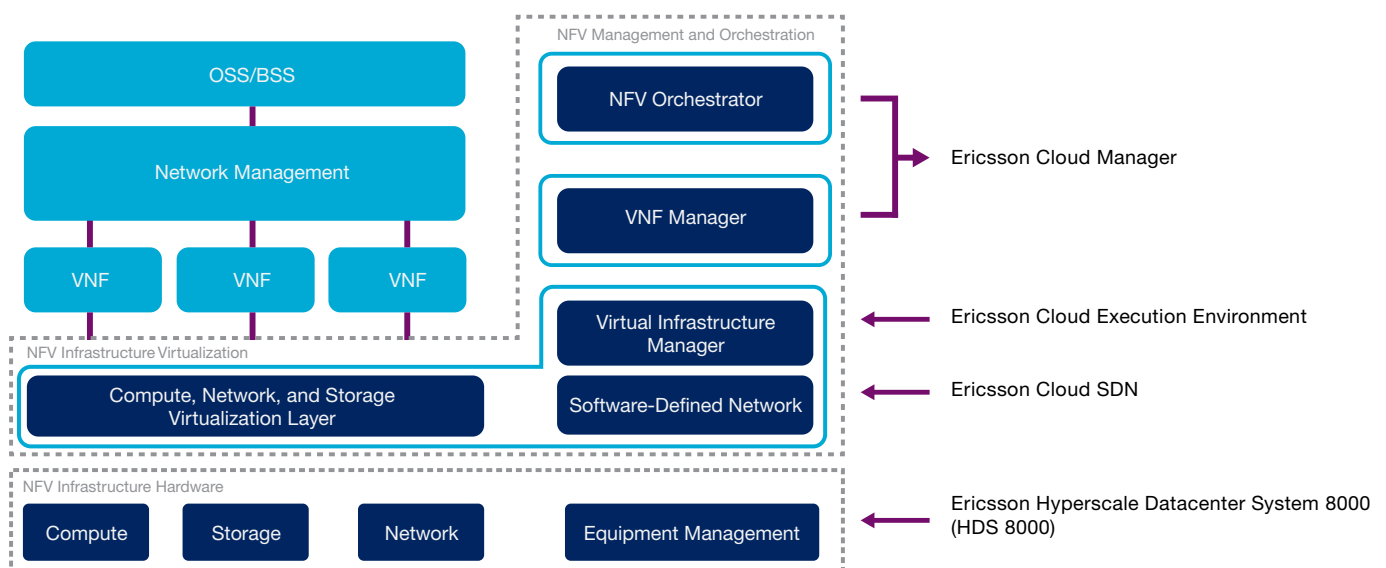


Figure 3. The Ericsson NFVi with Ericsson HDS 8000, based on Intel® technologies, provides orchestration between all roles defined by the ETSI model for NFV; the NFVi encompasses the hardware and software components found within the dashed line

INTEL® XEON® SCALABLE PROCESSOR WITH ADVANCED SECURITY TECHNOLOGIES

Intel® Data Protection Technology with Advanced Encryption Standard New Instructions (Intel® AES-NI) is an instruction set found on Intel Xeon processors that increases encryption performance and reduces processor load. That capability allows encryption to be implemented without noticeable performance overhead. Intel AES-NI can also strengthen encryption by protecting against “side-channel” snooping attacks, which makes it harder for malware to find vulnerabilities in the encryption. This efficient, hardware-based solution allows you to deploy industry-standard AES encryption more widely, so you can protect data on networks, on storage devices, and within applications without sacrificing performance.

Intel® Trusted Execution Technology (Intel® TXT) is an instruction set built within Intel Xeon processors. It works with enabled software to help detect systems booting with unknown BIOS, firmware, or operating systems, including those that have been attacked by some of the more advanced, persistent rootkits that threaten modern enterprises today. Intel TXT and other associated tools provide a processor-based, tamper-resistant environment that compares firmware, BIOS, and operating system or hypervisor code to known-good configurations to establish a measured, trusted environment prior to launch. If trust is not verified, Intel TXT identifies that the code has been compromised, which lets you protect the system and remediate the problem. By starting with a root of trust and measured launch environment (MLE), Intel TXT offers you better protection from malware and important visibility into the integrity of your system.

The solution also includes the cables, racks, and optical backplane, in addition to a network fabric consisting of:

- A control network that allows management platforms to monitor, control, and configure the compute, storage, and network resources
- A data network, which provides the network fabric for the application workloads

This modular structure makes it possible to change hardware components as technology evolves while keeping the same infrastructure over a long lifecycle. Intel Rack Scale Design and the Ericsson HDS 8000 allow network and system admins

to abstract resources in a way that completely transforms the data center to be much more flexible, efficient, and manageable.

vPOD overview and benefits

Physical PODs consist of the hardware chassis, an optical backplane, and the related components needed to connect, power, and manage the system. Physical resources can then be organized as software-defined partitions referred to as virtual PODs (vPODs).

A vPOD defines a set of compute, network, and storage hardware resources with internal networking between the servers and internal connectivity from the servers to the

storage hardware. A vPOD’s pooled resources can be distributed across multiple racks, chassis, and storage sleds, and they are characterized by immediate and high availability. vPODs are configured and managed in Ericsson Command Center.

The vPOD architecture is key to how the Ericsson HDS 8000 can help optimize NFV and the enterprise data center. Because vPODs are actually virtual, software-defined partitions, they allow IT and telecom operators to pool hardware resources to support multiple, concurrent cloud platforms, such as a telecom cloud, IT cloud, and commercial cloud, all running across the same switch fabric.

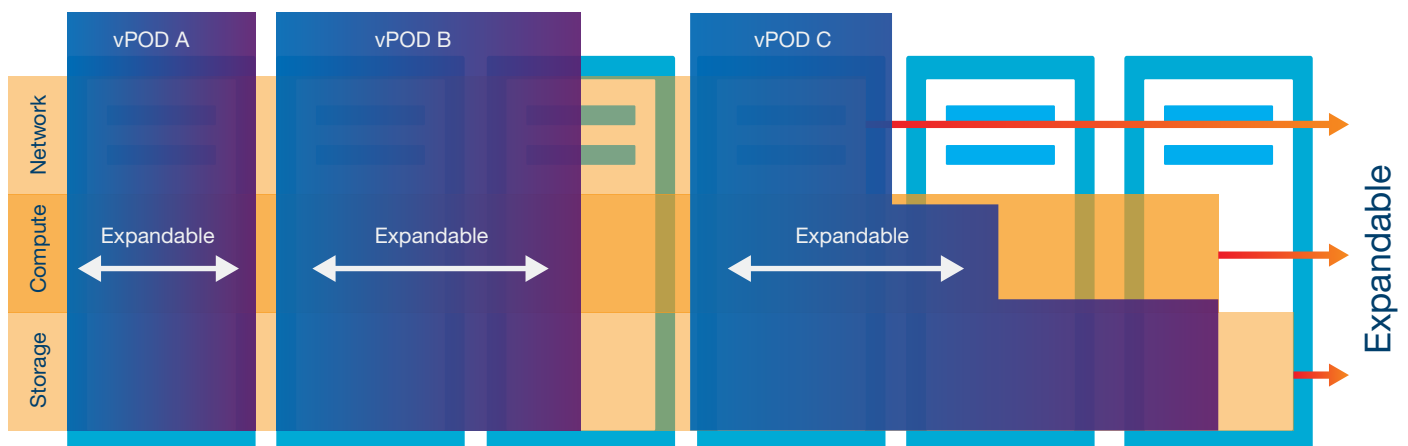


Figure 4. Based on Intel® Rack Scale Design, Ericsson HDS 8000 disaggregates compute, storage, and networking and allows resources to be assembled into highly agile vPODs

Admins can use the Ericsson Command Center portal to assign to a vPOD the hardware that is best suited to a particular workload. All storage, compute, and networking infrastructure is allocated as needed from different physical machines or racks. And because all infrastructure is software-defined, it is simple for an admin to scale or modify a vPOD to correctly size it for a given workload. One of the benefits of software-defined scaling of components into vPODs is the ability to tailor resource levels to fit a particular application workload. With this approach, servers can be optimally selected for any workload regardless of whether it changes over time. The unused components in the data center can be used by other applications or other users at the same time, which increases utilization.

Ericsson HDS 8000 can coordinate physical resources in near real time across machines and racks by using an optical backplane. This capability removes the limitations of traditional electrical-based connectivity between physical components. For example, with an optical backplane, memory and compute components can be connected over distances of hundreds of meters with minimal latency in performance.

The vPOD architecture and the NFVi provide a number of benefits for all environments running on them, including cloud flexibility and agility, better utilization, and single-pane-of-glass management.

Cloud flexibility and agility

The Ericsson NFVi with Ericsson HDS 8000 encompasses a comprehensive set of compute, network, and storage hardware in a rack based on Intel Rack Scale Design; it abstracts this hardware into orchestrated SDI. This design offers companies a level of agility previously available only to the cloud giants, like Google and Amazon. Because of their design, vPODs help organizations rapidly adapt to changing business needs. If an operator needs to launch a new virtual router, administrators can rapidly configure a vPOD to provision that service. If the service requires high performance, the vPOD can be configured with high-performance components. If it requires high

FUTURE-PROOF YOUR DATA CENTER

The Ericsson NFVi, when using Ericsson HDS 8000 based on Intel® technologies, is a continually evolving solution with an expansive visionary roadmap enabled in part by Intel® Rack Scale Design and the Intel® Xeon Scalable Processor. Organizations can deploy the infrastructure today and enjoy the benefits of SDI and data center modernization; and they can then grow along with the infrastructure as it adds new capabilities and features, some of which are described in this paper.

network bandwidth, administrators can assign it as much bandwidth as it needs from the abstracted pool. This agility means that enterprise businesses, software service providers, and communications service providers can use the flexibility of Ericsson HDS 8000 to modernize and transform their businesses in response to emerging opportunities.

Better utilization

The Ericsson NFVi increases overall resource utilization in the data center by disaggregating physical resources from multiple machines and racks into unified resource pools. By eliminating multiple, siloed infrastructure stacks, IT administrators and telecom operators can reduce or remove the need to purchase additional hardware for over-provisioning compute or storage. IT can also reduce hardware requirements across the organization by more efficiently redistributing software-defined compute, storage, and networking resources in response to changing business needs or peaks in usage. Previously idle resources can be used more effectively, which reduces the need to purchase additional servers or storage to scale or change a given workload.

Single-pane-of-glass management for all environments in the NFVi

Ericsson Command Center enables better visibility and reporting of data center resources from a single combined view, so admins can gain detailed knowledge into what is going on inside the data center. This insight can help companies reduce waste from unidentified resources

sitting unused or underused. Better visibility also helps organizations determine the best way in which to distribute existing resources for their current and planned workloads. With a more complete picture of company-wide compute, network, and storage infrastructure, organizations can plan for future needs in a more systematic way.

Using vPODs for various NFV environments

The vPOD-centric architecture of the Ericsson solution enables extraordinary flexibility because vPODs can be adapted for specific needs, even when business conditions change. This approach enables better utilization because under-utilized resources can be assigned where they are needed—to vPODs running more demanding workloads. Organizations can improve efficiency through centralized infrastructure management from a single pane of glass.

The Ericsson NFVi with Ericsson HDS 8000 is versatile, which is critical to telecom operators, who often need to support multiple versions of an NFVi environment based on their application certifications. With vPODs and the Ericsson HDS 8000, administrators can create parallel NFVi environments to support different versions of an application. The orchestration layer can then move workloads between the environments as needed, while administrators gradually move hardware into a vPOD running the newer version as demand on the older version declines.

As another example of the infrastructure's flexibility, administrators can run a vPOD with typical compute and storage resources and then connect accelerators, such as SmartNICs, which provide high throughput and low latency to those applications that require it, for example in the packet core domain or for user-plane traffic generally. Administrators can also move the accelerators to different vPODs or different types of servers within a vPOD as needed.

Hardware composition for different needs

vPODs on the Ericsson HDS 8000 support various configurations. For example, one server in a vPOD could be optimized for processing power to handle CPU-bound applications whereas another might require more storage and another higher network-throughput capabilities. Administrators could then select the right server for each type of workload. This approach can help dramatically reduce under-utilization of resources. Figure 5 shows how administrators use Ericsson Command Center to create vPODs. Ericsson Command Center invokes Intel Rack Scale Design APIs to form one or more nodes from pools of disaggregated resources. These nodes are then assigned to a vPOD over the network fabric, and workloads can be assigned to the vPOD, all through a single pane of glass.

This section discusses example hardware compositions of vPODs for network functions that telecom operators might want to virtualize.

Example 1: Customizing a vPOD for virtualized Ericsson User Data Consolidation (vUDC)

The Ericsson User Data Consolidation (UDC) solution eliminates the complexity of managing user subscriptions by consolidating subscriber information from the different network silos into one common repository, creating a unified profile. This approach allows for a simpler and more scalable network topology and greater efficiency in managing the network's database resources. It also gives more flexibility to introduce new services.

Making use of the unified profile, Ericsson UDC enables a more complete view of the customer base, which facilitates monetization of data assets and helps to grow revenue through innovative use cases. Thanks to its embedded analytics and exposure capabilities, Ericsson UDC makes available the required knowledge about the customer base that is crucial to deliver better customized services.

Ericsson UDC can be deployed as a VNF on the Ericsson NFVi. Ericsson UDC includes the centralized user database, so it requires more memory than other telecom VNFs. However, native Ericsson UDC is limited to the amount of memory available to the Generic Ericsson Processor. To overcome this limitation, administrators can run virtualized UDC (vUDC) on a memory-optimized server on a vPOD, thus improving the application performance.

Ericsson UDC also requires more storage than most telecom applications, again because of the centralized user database, so storage size and throughput are important considerations. Administrators can allocate additional storage and NICs to the vPOD where vUDC runs in order to better ensure optimal storage and input/output operations per second (IOPS).

Example 2: Customizing a vPOD for virtual Ericsson Evolved Packet Gateway (vEPG)

Ericsson Evolved Packet Gateway (EPG), together with Serving GPRS Support Node—Mobility Management Entity (SGSN-MME), is a critical component of the Ericsson Evolved Packet Core. Ericsson EPG is part of the Ericsson end-to-end solution for rapid deployment of highly scalable and reliable Long Term Evolution (LTE) networks. Service providers use the Ericsson EPG as a gateway between their mobile packet-core network and other packet-data networks, such as the Internet, corporate intranets, and private data networks.

Ericsson EPG can be deployed as a VNF on the Ericsson NFVi. One of the primary requirements for Ericsson EPG is high packet throughput, so network-optimized servers in a vPOD are ideal. Significantly, the Ericsson EPG workload drops predictably at night, so administrators can proactively move servers in and out of the vPOD to accommodate changes in demand.

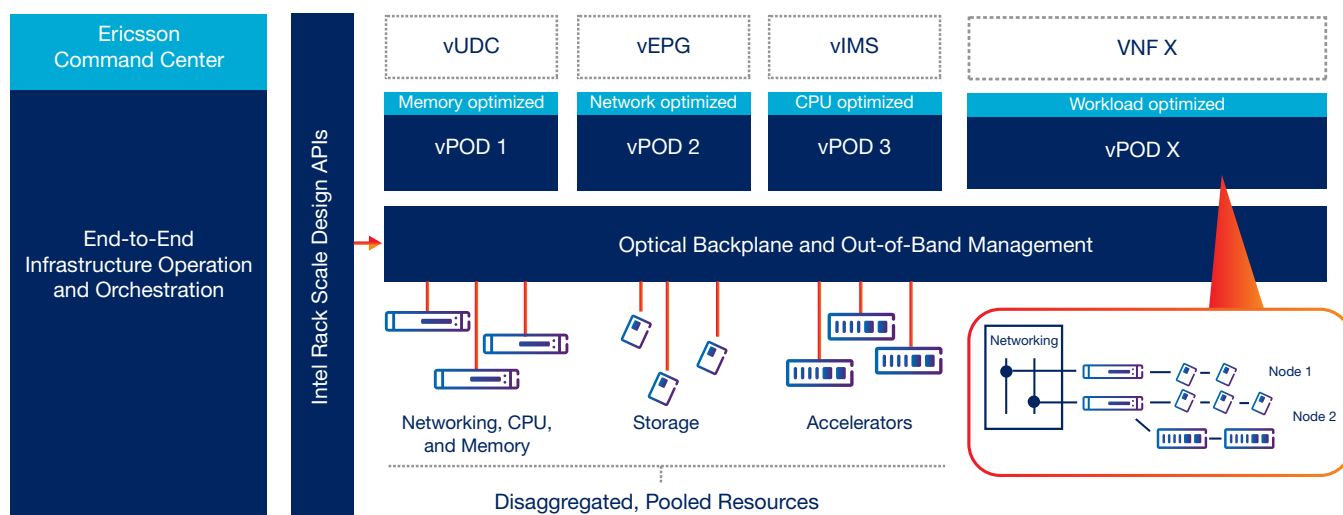


Figure 5. Ericsson Command Center invokes Intel® Rack Scale Design APIs to compose vPODs from pooled resources

This approach improves utilization because unused servers can be re-allocated where they are needed, such as to a vPOD running billing or analytics workloads.

To ensure high network throughput, virtualized EPG (vEPG) takes advantage of the virtual switch (OVS) accelerated by DPDK and of single root-I/O virtualization (SR-IOV) and Peripheral Component Interconnect (PCI[®]) pass-through. These capabilities are supported by Intel technologies like Intel[®] Virtualization Technology for Connectivity (Intel[®] VT-c) and Intel Virtualization Technology for directed I/O (Intel[®] VT-d) and by 10/40 gigabit Ethernet (GbE) Intel[®] Ethernet Converged Network Adapters. A SmartNIC could also be assigned to the server hardware to further enhance network performance and throughput.

Example 3: Customizing a vPOD for virtualized Ericsson IP Multimedia Subsystem (IMS)

Ericsson IP Multimedia Subsystem (IMS) is a core network solution that delivers rich real-time communication services for both consumer and business users over any access network and for any device type, including smart phones, tablets, wearables, laptops, and fixed phones. Examples of communication services are high-definition (HD) voice (such as VoLTE), Wi-Fi[®] calling, enriched messaging, enriched calling with pre-call info, video calling, HD video conferencing, and web communications.

Virtualized Ericsson IMS is CPU-

intensive, so the vPOD should contain a CPU-optimized server. It is easy to add more servers to the vPOD as needed to meet demand. While virtualized Ericsson IMS itself does not store user data, Ericsson IMS IPWorks could require a fairly high amount of storage—though less than vUDC. IPWorks provides centrally managed Domain Name System (DNS), telephone-number mapping (enum), and triple-a services for virtualized Ericsson IMS. Additional storage can be added to the vPOD as needed to help ensure that IPWorks runs at optimal performance.

Example 4: Customizing a vPOD for virtual Ericsson Multiservice Proxy (vMSP)

Ericsson Multiservice Proxy (MSP) is a multi-purpose, multi-technology network node that provides service control and enables policy control while representing an integration point for value-added services. It allows operators to differentiate user experiences, control costs, and monetize data traffic, and it can be run as a VNF on the Ericsson NFVi.

Virtual MSP (vMSP) requires optimal network performance, so administrators would assemble a vPOD that optimizes packet throughput. Server load balancers in vMSP would fully utilize the vPOD's vSwitch capacity, but they do not require all of the cores on a host. Because fewer cores are required, unused CPU capacity could be repurposed to other applications, such as billing or analytics.

Immediate benefits and future readiness

NFV has the potential to provide transformational benefits to telecom operators. Virtualizing network functions can deliver flexibility, simplicity, and cost savings that a traditional, hardware-based approach cannot match. However, telecom networks have unique needs. Challenges related to performance, latency, and scalability are particularly important to telecom operators and must be overcome before VNFs can be pervasively deployed in telecom networks.

The Ericsson NFVi, with Ericsson HDS 8000 based on Intel Rack Scale Design, provides a flexible, hyperscale foundation, and it removes barriers for the virtualization of demanding telecom networking functions. Organizations can continue their NFV journeys by consolidating multiple VNF environments onto the Ericsson solution where they can be managed from a single interface.

Over time, the benefits of Intel Rack Scale Design and Ericsson HDS 8000 will increase as these transformative technologies continue to develop. Organizations that deploy Ericsson and Intel technologies can realize the benefits of SDI today while positioning themselves to take advantage of technical advances in the future.

LEARN MORE

To learn how the Ericsson HDS 8000 can serve as an SDI for common IT workloads, read “Transform How You Manage Your IT Data Center Workloads.”

Learn more about Ericsson HDS 8000:

www.ericsson.com/hyperscale/cloud-infrastructure/hyperscale-datacenter-system

Learn more about Intel Rack Scale Design:

www.intel.com/content/www/us/en/architecture-and-technology/rack-scale-design-overview.html

¹ ETSI. "Network Functions Virtualization." www.etsi.org/technologies-clusters/technologies/nfv.

² Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

³ This section provides only an overview of the Ericsson HDS 8000. For a detailed description, contact your Ericsson representative.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

Ericsson is the trademark or registered trademark of Telefonaktiebolaget LM Ericsson. All other product or service names mentioned in this document are trademarks of their respective companies.

*Other names and brands may be claimed as the property of others.