Industry Report

2018 NFV Report Series Part 1: NFV Infrastructure (NFVI) and VIM

The Trusted News and Resource Site for SDx, SDN, NFV, Cloud and Virtualization Infrastructure

Table of Contents

Executive Summary . . . 1

Introduction . . . 2

NFV, NFVI and VIM Primer . . . 3
- Brief History of NFV: Role of the ETSI ISG . . . 3
- NFVI – Role of NFV Infrastructure . . . 4
- NFVI - Virtualization Layer . . . 5
- NFVI - Virtual Infrastructure . . . 6
- NFV VIM - Virtualized Infrastructure Manager . . . 7

Organizations Accelerating NFVI and VIM . . . 10
- ETSI . . . 10
- Open Compute Project . . . 10
- Linux Foundation . . . 10
- Open Networking Foundation . . . 13
- Telecom Infra Project (TIP) . . . 13

Common NFVI and VIM Architectural Deployments . . . 14
- Centralized Cloud . . . 14
- Core DC + Edge . . . 14

2018 NFVI and VIM Market Update and Trends . . . 16
- Server and Compute Market – Driven by Cloud . . . 16
- Servers: Standard x86 Still the Foundation of NFVI . . . 16
- Edge NFVI deployments – the New Battleground for Intel, AMD, ARM . . . 18
- ARM, Alternative Architectures and HW Accelerators . . . 19
- Storage Hardware – Going Open . . . 20
- Switches: Proprietary versus Bare Metal + Open Software . . . 20
- Virtualization and the VIM Updates . . . 20
- VMs vs Containers, VMs with Containers? . . . 21

2018 SDxCentral NFVI and VIM Reader Survey . . . 23

2018 NFVI and VIM Report Conclusion . . . 28

2018 NFVI and VIM Products . . . 29

Featured NFVI AND VIM Report Products . . . 30
- AT&T: AT&T FlexWare . . . 30
- Enea Software: Enea NFV Access . . . 31
- Hewlett Packard Enterprise (HPE): HPE NFV Blueprints . . . 32
- Lenovo: Lenovo NFV Solutions . . . 33
- Nokia: Nokia Airframe Data Center Solution . . . 34
- VMware, Inc.: vCloud NFV . . . 35

contents

NFVI Data Center Hardware Platforms . . . 36
- Cavium, Inc.: ThunderXTM ARMv8 Processors . . . 36
- Dell EMC | Service Provider Solutions: Dell EMC NFV Ready Bundle for VMware . . . 36
- **Hewlett Packard Enterprise (HPE): HPE NFV Blueprints [FEATURED]** . . . 32
- Intel: Intel Select Fast Track Kit for NFVI . . . 37
- Kontron (Communications Business Unit) . . . 37
- **Lenovo: Lenovo NFV Solutions [FEATURED]** . . . 33
- Mellanox Technologies: BlueField SmartNIC . . . 38
- **Nokia: Nokia Airframe Data Center Solution [FEATURED]** . . . 34
- Netronome: Netronome Agilio CX 10, 25 and 40GbE SmartNICs . . . 38

NFVI Other Hardware Platforms . . . 39
- **AT&T: AT&T FlexWare [FEATURED]** . . . 30
- Ciena: Ciena 3926m Service Delivery Switch . . . 39

NFVI Software . . . 40
- 6WIND: 6WIND Virtual Accelerator . . . 40
- ADVA Optical Networking: Ensemble Connector . . . 40
- Cisco Systems: Cisco Enterprise Network Functions Virtualization . . . 41
- **Enea Software: Enea NFV Access [FEATURED]** . . . 31
- Huawei: FusionSphere OpenStack . . . 41
- Linux Foundation: Open Container Initiative . . . 42
- Mirantis: Mirantis Cloud Platform . . . 42
- Radisys: FlowEngine . . . 43
- Red Hat: Red Hat OpenStack Platform . . . 43
- Telco Systems: NFVTime . . . 44
- **VMware, Inc.: vCloud NFV [FEATURED]** . . . 35
- Wind River: Titanium Cloud Product Portfolio . . . 44

VIM Products . . . 45
- Canonical / Ubuntu: Ubuntu OpenStack . . . 45
- Ericsson: Ericsson Cloud Execution Environment . . . 45
- ZTE Corporation: ZTE TECS . . . 46

OPEN WINS WHEN OPEN WORKS.

Open works here.

Learn more about VMware Integrated OpenStack at **vmware.com/openworks**

Virtualize your network infrastructure with the proven, production-tested OpenStack® platform that's ready to work right now. VMware® Integrated OpenStack is easy to deploy with current infrastructures—and built for what's next. Optimize workloads, accelerate revenue streams, and prepare for the 5G future with the secure, carrier-grade solution that lowers TCO by 47 percent.*

Proven

Our reliable, feature-rich platform has been deployed by communication service providers around the globe serving millions of subscribers.

Seamless

Interoperability with existing infrastructures helps you set up, deploy, and run NFV infrastructure faster and more efficiently.

Cost-effective

TCO is 47 percent lower than that of a deployment based on completely open source–based components.

vmware®

REALIZE WHAT'S POSSIBLE.™

*ACG Research. "The Total Cost of Ownership (TCO) Advantages of VMware vCloud NFV OpenStack in the Virtualized Packet Core of Service Providers' Mobile Networks." April 2018.

Introducing AT&T FlexWare.

Firewalls, routers, and more in one.

LEARN MORE

JUNIPER NETWORKS | AT&T Business

Transform your telco infrastructure

NFV PLATFORMS OPTIMIZED FOR 5G EVOLUTION

HPE NFV Blueprints simplify telco cloud deployments and enable scalability, reliability and optimal performance at each layer of NFVI Stack. Future-proof your business for 5G and beyond by leveraging proven designs to reduce complexity and accelerate time to value through automation, whilst adhering to open industry standards.

Learn more: hpe.com/dsp/infrastructure

Accelerating next

Hewlett Packard Enterprise

Executive Summary

Over the course of the last 5 years, since ETSI and the world's leading communications service providers (CSPs) launched the ETSI ISG for NFV (Network Function Virtualization) in early 2013, vendors and CSPs alike have been working to make NFV a reality. Disaggregation of network services from proprietary hardware, deployment of a fully virtualized infrastructure, modern cloud designs and advanced, agile orchestration and management all promise lower CapEx and OpEx, with improved agility and time to market.

In 2018, the momentum behind NFV continues, even as reality sets in that full NFV adoption will take longer as CSPs slowly transition their infrastructure, business processes and culture—NFV adoption is not simply a technology problem. Regardless, NFV, along with cloud adoption, will continue its uptake and we will cover the state of NFV across our three report series. Our 2018 reports will cover a wide range, from industry-standard commercial off-the- shelf (COTS) hardware as NFV Infrastructure (NFVI), to hypervisors and the Virtualized Infrastructure Managers (VIMs), to Virtual Network Functions (VNFs) and the management and orchestration (MANO) necessary to deploy these functions.

This specific report covers the latest in NFVI and VIM development, the key foundations of any NFV deployment. We continue to see diversity in NFVI hardware, from the dominant standard 19-inch rackmount servers to other forms like blade servers and an increasing focus on hyperconverged architectures (HCI), especially for edge deployment supporting 5G and IoT services. As well, alternate compute platforms based on ARM architectures battle for CSP deployments with x86 architecture from both Intel and AMD are popping up at the Edge and on customer premises equiment (CPE). We'll cover more on the CPE, vCPE (virtual CPE) and uCPE (universal CPE) trends in our SD-WAN and Virtual Edge report later this year.

Along with server compute considerations, there's more appreciation for HW acceleration to improve I/O performance and acceptance that HW-accelerated/Smart NICs might be needed, along with software approaches from open-source projects like DPDK and FD.io.

And we observe continued uptake of options beyond VMs as Linux containers gain favor with CSPs and vendors alike, with more realization that hybrid deployments consisting of bare metal, VMs and containers will likely be the reality and that the effective management of all three will be critical to NFV success.

On the VIM front, VMware vRealize suite (as part of vCloud NFV) and its ESXi hypervisor are gaining more mindshare and acceptance from the CSP community. And while OpenStack remains the de facto option, some CSPs are looking at trialing both and, in some cases, running both for different workloads. Many already VMware in their IT infrastructure, but are now open to running network services on the same. Both VIMs continue to provide improved container services and as production-grade container-basedVNFs come on the market, both are poised to support such deployments.

2018 will continue to see more NFVI deployments and we expect more innovation on this front, as VMs and containers battle it out, as x86 and ARM jostle for relevance and dominance, as the Edge reveals new requirements for new application workloads and innovative HCI platforms come to the market.

Introduction

Welcome to the 2018 Network Functions Virtualization Infrastructure (NFVI) and VIM (Virtual Infrastructure Manager) Report. This is first in a series of three reports covering NFV, with the second focusing on NFV MANO and Assurance while third report covers VNFs and popular use cases. This report specifically provides an update into the trends and progress of NFVI, as well as the VIM (which we view as closely related to NFVI).

NFV continues to be a compelling move for Communication Service Providers (CSPs) worldwide because of its promise of fulfilling the networking needs of a service provider on standard server and storage infrastructures, coupled with an agile virtualized environment and flexible orchestration systems.

The NFV framework from the ETSI ISG was launched over 5 years ago, in January of 2013. And it consists of the following major components:

- **NFV Infrastructure (NFVI)—**The physical resources (compute, storage, network) and the virtual instantiations that make up the infrastructure.
- **NFV Management and Orchestration (NFV MANO)—**The management and control layer that focuses on all the virtualization-specific management tasks required throughout the lifecycle of the VNF. In ETSI's model, the Virtualized Infrastructure Manager (VIM) lives in this component, but because the VIM is so tightly coupled to the NFVI, we will cover it along with the NFVI in this report. We will cover MANO as well as service assurance our second NFV report.
- **Virtualized Network Functions (VNFs)—**The software implementation of a network function. The third report in our NFV series, the VNF report, will go into a lot more detail and will also cover the most popular use cases today.

Our analysis includes examination of hundreds of our news and analysis articles, in-depth interviews we have conducted with technology vendors and end users, and the results of the SDxCentral NFV Survey.

In addition to an overview of the technology and an analysis of customer expectations, we also collected data from the leading vendors. The product information from many leading technology vendors is available at the end of this report.

What you can expect from this report:

- An overview of NFVI and VIM, describing the evolution of architectural components and potential benefits.
- Organizations that have influenced (and continue to influence) NFV evolution
- The NFVI and VIM landscape today
- Details on different vendor offerings in the NFVI and VIM space

Thank you for reading or downloading this report, we hope you will find it a useful resource as you look to understand and adopt NFV technologies.

NFV, NFVI and VIM Primer

For SDxCentral members unfamiliar with the history of NFV and ETSI, or who are looking for a refresher, we'll provide an overview in this section and get you up to speed.

Brief History of NFV: Role of the ETSI ISG

NFV originated in the service provider community as operators looked for ways to cut costs and accelerate the roll out of profitable services to monetize their networks and grow their revenues. Expensive hardware-based appliances, which are complex to deploy and manage, were limiting the providers' ability to quickly trial new services and reduce costs.

Service providers felt if they could decouple the network services from the hardware, it would allow them to deploy networking components that could truly fit and support a fully virtualized infrastructure, including servers, storage and even other networks.

As NFV gained momentum, operators gathered within the European Telecommunications Standards Institute's (ETSI) and created the Industry Specification Group (ISG) for NFV to accelerate the progress of virtualizing network functions. Launched in January of 2013, the ETSI ISG for NFV has been working to develop the requirements and architecture of virtualized network functions in a telecommunication's network. It included these components of the NFV framework:

- NFV Infrastructure (NFVI) – The physical resources (compute, storage, network) and the virtual instantiations that make up the infrastructure.
- Virtualized Network Functions (VNFs) – The software implementation of a network function.
- NFV Management and Orchestration (NFV MANO) - The management and control layer that focuses on all the virtualization-specific management tasks required throughout the lifecycle of the VNF.

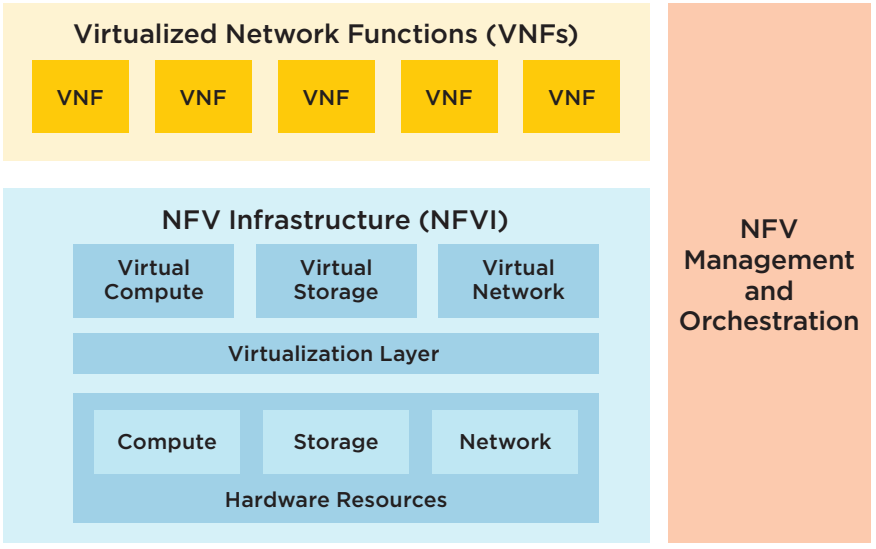

Adapted from ETSI publication GS NFV 002: Network Functions Virtualization (NFV); Architectural Framework

HIGH LEVEL NFV FRAMEWORK

Since the original ETSI model for NFV was released, vendors and operators have also put together their own technology programs, including developing their own open source projects which they then donate to the community (e.g. ONAP which is a combination of ECOMP from AT&T and Open-O from China Mobile).

NFVI – Role of NFV Infrastructure

With the basic premise behind NFV out of the way, let's focus on the two major aspects of the NFV infrastructure that we'll cover in the report: the NFVI (the foundation platform itself that includes the hardware, operating systems, hypervisor etc) and the VIM (the resource management component of the NFVI). Basically, NFV Infrastructure (NFVI) consists of the physical and virtual compute, storage and networking resources that VNFs run on.

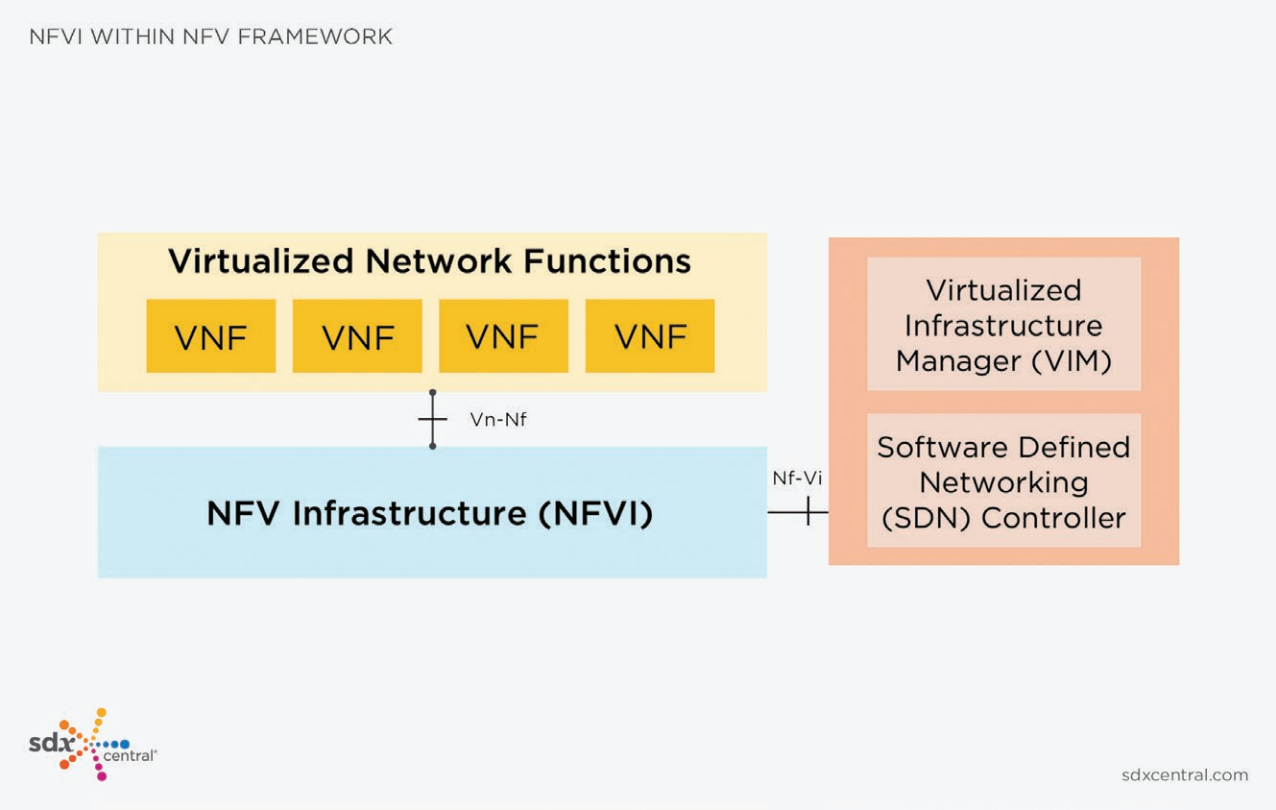

The NFVI layer primarily interacts with two other NFV framework components: VNFs and Virtualized Infrastructure Manager (VIM). The VIM, on the other hand, is responsible for provisioning and managing the virtual infrastructure. This means that the VNF to NFVI interface (Vn-Nf) constitutes a datapath through which network traffic traverses, while the NFV to VIM interface (Nf-Vi) constitutes a control-path that is used solely for management but not for any network traffic.

NFVI consists of three distinct layers: physical infrastructure which is made of servers and other hardware elements, as well as the virtualization layer and the virtual infrastructure.

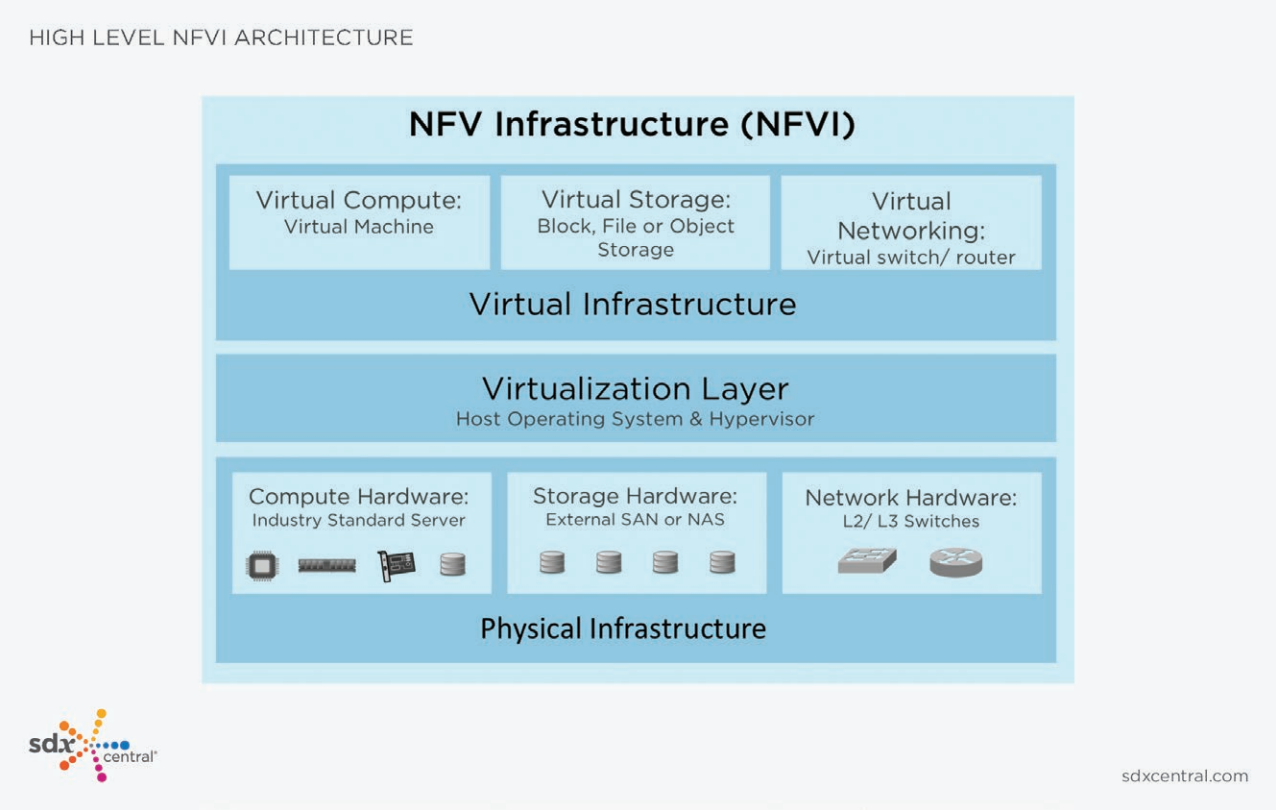

NFVI - Virtualization Layer

The NFVI virtualization layer sits on top of the hardware and is a software platform that typically involves a hypervisor on which VMs (virtual machines) can run. In the last year, Containers have become potentially viable for NFV production, and so this layer has expanded to incorporate container run-time engines. However, most container-based deployment are in trials or very limited deployments and are still viewed as immature for full-scale deployment. VMs and hypervisors make up the majority of NFV deployments today.

Hypervisors split up the resources of the physical machine and offer the equivalent of a physical machine to the application. In fact, the three main functions of the hypervisor are to A) split up the resources of the physical machine, B) provide isolation between different VMs and C) emulate all the necessary peripherals e.g. NIC cards. The main hypervisors used for NFV are VMware vSphere (ESXi), and KVM (there are other VM technologies from vendors like Microsoft and Oracle but these have not featured much in the NFV space):

- **vSphere:** vSphere is a commercial hypervisor from VMware. It is very mature (15+ years since the first release) and stable. For workloads that are not cloud native, vSphere also has several rich features that enable VM migration across hypervisors and high-availability. vSphere is what's called a Type 1 or bare-metal hypervisor in that it runs directly on the hardware.
- **KVM (Kernel-based Virtual Machine)**: KVM is a mature 10-year-old open source hypervisor project. It is a Type 2 or hosted open source hypervisor where it runs on top of a Linux operating system (called a Host OS to differentiate it from the operating system that runs in the VM, called a Guest OS). For this reason, when using KVM, you also need to choose the Host OS. Commonly used Host OS operating systems are RHEL, CentOS, SUSE or Ubuntu.

NFVI - Virtual Infrastructure

As we continue moving up the stack, the next step beyond the virtualization layers is the virtual infrastructure that consists of virtual machines, virtual storage and virtual networking.

Virtual Machines

Virtual machines are (VMs) created by the hypervisor as discussed above. Hypervisors present APIs to create, destroy, migrate and manage virtual machines. In the case of KVM, this is done through a library called libvirt, and for vSphere, it is done through a virtual machine manager called vCenter. These VMs are what VNFs are hosted within and that act as the workhorse powering the network services core to the NFV concept.

Virtual Storage

The responsibility of virtualizing block or file storage lies with the SAN, NAS or SDS (software-defined storage) software, and the VM is simply presented with a LUN (logical unit number) or file share. The virtualization layer may add additional features such as snapshots, backups, thin provisioning, VM sharing, replication etc. In the case of object storage, best suited for bulk storage, it is already virtualized and exposed through REST APIs, so nothing special needs to be done.

Virtual Networking

A hypervisor contains a virtual switch or router that has four functions: A) switch between VMs in the same physical node without having to go to an external switch most commonly used to create service function chains, B) provide overlay networking services where the networks exposed to the VMs are completely abstracted with their own address space and topologies different from the physical or underlay network, C) provide security services such as the enforcement of access control lists, and D) enable gateway connectivity to the internet. There are many options for virtual networking technologies used in NFV.

NFV VIM - Virtualized Infrastructure Manager

The VIM manages the NFVI and the serves as a conduit for control-path interaction between VNFs and NFVI. The VIM inventories, provisions, de-provisions and manages virtual compute, storage and networking while also communicating with the underlying physical resources. The VIM is responsible for operational aspects such as logs, metrics, alerts, root cause analysis, policy enforcement, service assurance etc. It is also responsible for interacting with the orchestration layer (MANO) and the SDN controller (if present) for configuring networking.

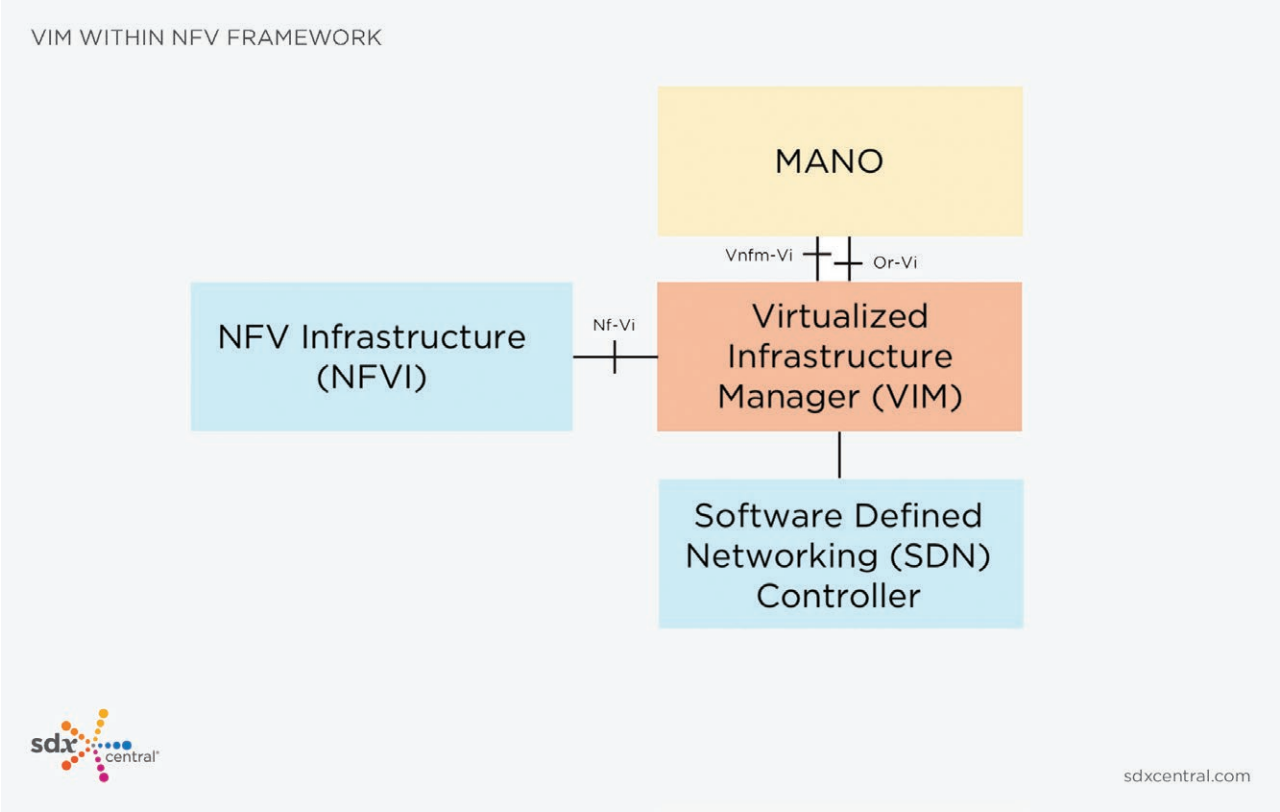

There are two main VIM software stacks prevalent in NFV: OpenStack and VMware vCloud Director. Other candidates like CloudStack have not been prominent in the last 2-3 years in the NFV space.

VIM - OpenStack

OpenStack is one of the largest open source projects in the world, and is eight years old with its 17th release, Queens, which provide virtual GPU support and also added more container-integration capabilities. To give a sense of the expansive scope, the Queens release had 159 companies contributing code with over 1400 contributors (source: Stackalytics). OpenStack is well-represented by a vibrant community with 7 platinum members, 21 gold members, and many more corporate sponsors and supporting organizations. The governance of the OpenStack Foundation allows for transparency and meritocracy; and many of the DevOps related infrastructure innovations by the OpenStack community have been copied by many other projects. However, the decentralized nature of the organization also results in a lack of consistency between projects, inadequately staffed projects and increasing complexity of the overall stack. There are numerous commercial flavors of OpenStack that are available, some of which are listed at the end of this report. Notable ones include those from Canonical/Ubuntu, Cisco, Ericsson, Huawei, IBM, Juniper, Mirantis, Red Hat, Suse, VMware, Wind River (their Titanium product was open-sourced recently as part of a new Edge-focused Linux Foundation project—more details will be provided later in this report).

market summary

Here's a quick summary table on the core and major optional OpenStack services:

Service	Project	Description
Compute	Nova	Nova is the heart of OpenStack and is responsible for scheduling a virtual machine onto a particular compute node based on a wide range of criteria (availability zones, resource utilization etc.) Nova also communicates with libvirt to create, destroy or manage the VM (e.g. increase resources, snapshot etc.)
Block Storage	Cinder	Cinder provides persistent block storage for VMs. It is an API translation layer that, through vendor specific drivers, communicates with SAN storage, NAS or SDS to carve up or destroy a LUN or file share and assign it to a VM. It also manages storage by triggering actions such as snapshots, backup etc.
Object Storage	Swift	Swift is the object storage project in OpenStack. Unlike Cinder which is an API layer, Swift actually implements the entire software defined storage stack. Though popular for enterprise use cases, we do not expect it to be very useful for NFV.
Networking	Neutron	Neutron can be used as an API layer to communicate to an external SDN controller via plugins or it can be used as a full blown SDN controller that communicates with switches through vendor specific drivers. Neutron also includes a few VNFs e.g. LBaaS, VPNaaS etc.
Image	Glance	VM images are stored in a Glance repository. The repository can reside on existing storage e.g. Cinder or Swift.
Bare Metal	Ironic	This is equivalent to Nova, but for bare metal resource provisioning. This project is also of limited use in NFV.
Identity	Keystone	Keystone is used to authenticate API access to any given service. It supports token-based authN and user-service authorization.
Telemetry	Ceilometer, Gnocchi, Aodh	These projects form a data collection service, time series database as a service and an alarming service. The goal is to reliably collect data about the utilization of physical and virtual resources, store that data and trigger actions based on policies. In the NFV stack, these activities are actually managed by the MANO layer, so these projects are not very useful as-is. However, Gnocchi and Aodh can also be connected to other collectors such as collected where they can be used for monitoring and service assurance, and so become a lot more valuable for NFV.
Orchestration	Heat	Heat uses human readable files to manage the entire lifecycle of infrastructure and applications. In addition to Heat templates, Heat also supports AWS CloudFormation files. Several MANO projects take advantage of Heat for infrastructure orchestration.

Service	Project	Description
Dashboard	Horizon	Horizon provides a graphical user interface to OpenStack. For installations that match the ETSI architecture, this is unlikely to be used since all interaction will be through APIs. However, there are NFV deployments that do not include MANO and therefore rely on the dashboard. OpenStack also has a CLI. Additionally, everything in OpenStack can be managed via REST APIs. In fact, REST APIs are a superset of what's available via the dashboard or CLI.
Data Processing	Sahara	A platform for the provisioning, management, and usage of clusters running popular processing frameworks.

VIM - VMware vCloud Director

VMware has been quite aggressive in the CSP market over the last year and have convinced many CSPs to build more NFV solutions around its technologies. The perception is that vSphere is more mature and safer—in particular, that it doesn't suffer from the upgrade challenges that many OpenStack installations seem to run into. The VIM that works with the vSphere hypervisor is called vCloud Director. vCloud Director is a proprietary cloud management platform from VMware. It is meant to be a day-one and day-two cloud management solution for compute, storage, networking and application level resources. The combination of vCloud Director and vSphere forms the VMware vCloud NFV Suite. And with vSAN software defined storage and NSX SDN controller, these technologies form a complete NFVI + VIM + SDN controller solution.

As described, vCloud NFV is essentially suite of multiple software products that work together including:

- **vSphere with Operations Management:** vCloud NFV utilizes vSphere to manage all the key virtual resources and adds improved adaptive operations management and automation (guided remediation and customizable actions).
- **vCloud Director for Service Providers (with option of VMware Integrated OpenStack):** vCloud Director is the VIM component of the vCloud NFV suite. VMware discontinued vCloud Director for its enterprise offerings choosing to focus vCloud Director instead on the CSP market, and has been adding new capabilities for vCloud Director (latest version was 9.0 in September 2017). VMware also adds the option to run its own flavor of OpenStack VIM (VIO) when clients desire to do so.
- **vRealize Operations Advanced Edition:** This component from the vRealize Suite provides improved visibility and assists with identification and remediation of performance and configuration issues. It also supports capacity optimization and planning as well as workload balancing.
- **vRealize Log Insight and vRealize Network Insight:** Log Insight is the log management component of the vRealize suite. It aggregates logs from the OS, apps, storage, network devices etc. and provides querying, dashboards, analytics and alerting. There is even machine learning to group related data together. Network Insight provides improved analytics across virtual and physical networks and helps provide recommendations for implementing micro-segmentation with NSX deployments.

Other elements of vCloud NFV include popular VMware components such as VMWare Site Recovery Manager, NSX, VSAN. vCloud NFV lets users tap into the stability and maturity of vSphere/ESXi along with the benefits of VMware's vRealize suite of products for management, visibility, planning and reliability.

Looking beyond VMs, there is now thinking that instead of using a VM-centric approach, a container-based approach might reduce overhead while achieving improve flexibility and manageability over underlying resources. For now, VMs are the predominant NFVI. We will discuss the rise of containers later in this report as we look at current trends.

Organizations Accelerating NFVI and VIM

There are quite a few standards bodies involved in driving NFV, including ETSI, Open Compute Project, Linux Foundation, MEF, TMForum, Open Networking Foundation (ONF), and OpenStack. Not all of them are involved in driving the NFVI layer; for example, MEF and TMForum focus more on MANO and assurance. In this section, we'll touch on organizations that impact NFVI and VIM report and cover others in the other parts of our report series.

ETSI

The European Telecommunications Standards Institute is a non-profit organization that produces globally-applicable standards for Information and Communications Technologies (ICT). Their activities are open to member and nonmember companies and individuals.

ETSI took the lead on NFV in 2012 with a breakthrough "Network Functions Virtualization" white paper. Since then ETSI has established an Industry Specification Group (ISG) for NFV that has grown from seven to more than 300 member companies, and has published more than 100 NFV documents.

ETSI has also gone beyond just publishing standards and into the realm of non-SDO (Standards Development Organization) activities. The NFV ISG has been busy creating architecture documents, white papers and hosting plugfests. It's taken on hosting open source projects such as OSM (Open Source MANO), develops architecture documents, publishes white papers and conducts plugfests. It is also actively collaborating with other projects, such as Linux Foundations' OPNFV in working through interoperability issues.

Open Compute Project

The Open Compute Project (OCP) is an open-source project that was started in 2011 by Facebook, Intel and Rackspace. And as evidenced by the strong attendance of over 3,400 people at its recent 2018 OCP Summit in March, there's a groundswell of activity from both the vendors and overall community. OCP aims to bring open-source principles to hardware design. OCP designs servers, bare metal switches, storage and rack solutions for datacenter applications. It is a broad project that deals with hardware, mechanical design, cooling, power, firmware, APIs and in some cases software and APIs as well (ONIE, Open Network Install Environment, and SAI, Switch Abstraction Interface) for switch management.

In addition to hardware server designs, the OCP has also been instrumental in driving white box networking, the details of which can be found in **our 2018 next-gen data center networking report which was recently published**.

Recently OCP has also branched out to the telco edge with clear interest from vendors and CSPs alike. For the edge, the servers have unique power, geometry, ruggedness etc. requirements. OCP, by all accounts, has been quite successful, with recent analysis indicating that it has had over **$1B worth of market impact**.

Linux Foundation

The Linux Foundation has many open source projects directly relevant to NFV and in particular NFVI. Recently numerous projects were combined under the umbrella of LFN (Linux Foundation Networking), a new entity to increase harmonization between the numerous networking projects hosted within the Linux Foundation:

- **ONAP:** policy-based orchestration of virtual network functions
- **OPNFV:** an integration project bootstrap the development and adoption of NFV components across open networking platforms; focused on integration testing, CI/CD.
- **OpenDaylight:** which builds the most widely used open source SDN controller
- **FD.io:** an umbrella project for various enhancements to the Data Plane Development Kit (DPDK), a programming framework for network acceleration that supports hardware platforms
- **PDNA:** which is developing an open source data analytics platform for network device and service telemetry
- **SNAS:** another network management framework to collect, aggregate and access real-time network telemetry with APIs to query, visualize and summarize collected data

Other projects that are relevant to NFV include:

- **IOVisor:** Dataplane acceleration software
- **OVS:** Open vSwitch
- **CORD:** Central Office Re-architected as a Datacenter, an edge-focused platform for developing carrier solutions, based on ONOS
- **ONOS:** an open-source SDN Controller
- **Open Switch (OPX):** Network operating system
- **CNCF:** Cloud Native Computing Foundation that hosts projects such as Kubernetes for container orchestration
- **Tungsten Fabric (formerly OpenContrail):** an open-source SDN controller and network virtualization fabric released by Juniper (part of its acquisition of the startup Contrail) and now part of the Linux Foundation projects
- **DPDK:** software acceleration technology for improved packet handling
- **FRR:** Free Range Routing is a comprehensive routing suite that had its roots in the open-source routing Quagga project
- **VCO:** Virtual Central Office – a sub-project tied to OPNFV that is focused around deployment of solutions such as virtual CPE in central office locations (similar to CORD but based around the OpenDaylight SDN controller
- **DANOS:** **DANOS is a recent addition to Linux Foundation, with code contributed by AT&T**, from it's dNOS (disaggregated Network Operating System) project. DANOS is a NOS for the white boxes within the network and will be combined with FRR, OPX and the ONF's Stratum as well as OCP/Microsoft's SONiC to provide a complete solution suite for white box networking.

The Linux Foundation has recently been **moving towards harmonization** and "re-aggregation" of the numerous projects under its wing, and reaching out to other groups like OCP, MEF, etc to try to corral the disparate efforts going on in the open-source community. This is in response to CSP and enterprise end-users starting to complain about the large number of projects, which are getting increasingly harder to keep track off and which might result in fragmentation and dilution of overall efforts and momentum in the open-source ecosystem.

OPNFV

OPNFV PROJECT

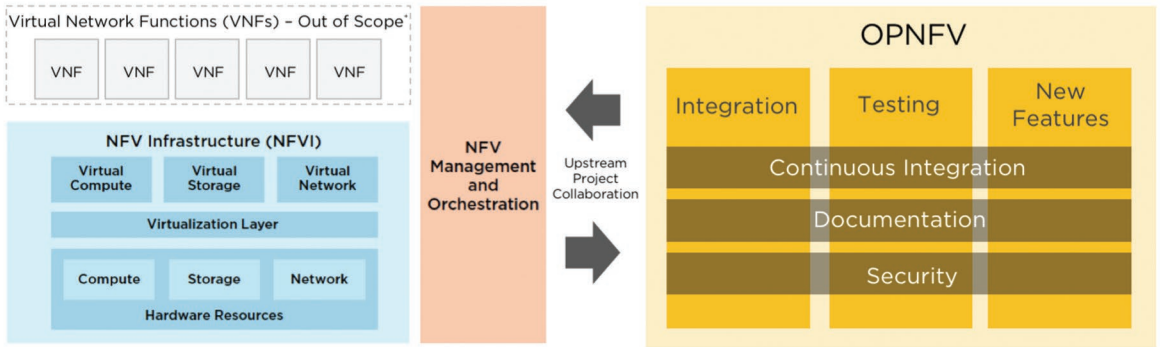

*VNFs may be part of test suites, but are not officially integrated and tested in OPNFV

sdxcentral.com

OPNFV is part of LFN and originally formed in 2014 with the goal of accelerating NFV adoption through system level integration, deployment and testing. Their continued mission is to create a reference NFV platform to accelerate the transformation of enterprise and service provider networks. Most recently, they rolled out the OPNFV Verified Program (OVP)—a verification and testing program for vendors and suppliers who are creating systems consistent with the OPNFV reference platform.

OPNFV is a non-traditional open source project that focuses on system integration and testing rather than software development. It is a very active project with broad community participation. The output is a set of well integrated interoperable reference software implementations that address NFV needs. OPNFV integrates components from upstream projects including OpenDaylight, ONOS, Tungsten Fabric (OpenContrail), OVN, OpenStack, Kubernetes, Ceph, KVM, Open vSwitch, and Linux. OPNFV also includes other data plane projects such as FD.io, DPDK, and ODP, and has moved to be more processor-agnostic recently, able to run on both x86 (Intel) and ARM commercial and white-box hardware.

CORD

Central Office Re-architected as a Datacenter (CORD) is also an open source NFV project that is hosted by Linux Foundation but driven by the ONF (which merged with ON.Lab, the original creator of CORD). CORD combines NFV, SDN, and cloud platforms to provide CSPs with a modern infrastructure to host new NFV applications. Major supporters of CORD include AT&T, SK Telecom, Verizon, China Unicom and NTT Communications.

CORD aims to build standardized PODs that provide "as-a-service" capabilities for CSPs to build network solutions around, and includes OpenStack or Docker, ONOS SDN controller, and VTN as part of their deployment. Key use cases that CORD focuses on today include R-CORD (residential services), M-CORD (mobile services), and E-CORD (enterprise services).

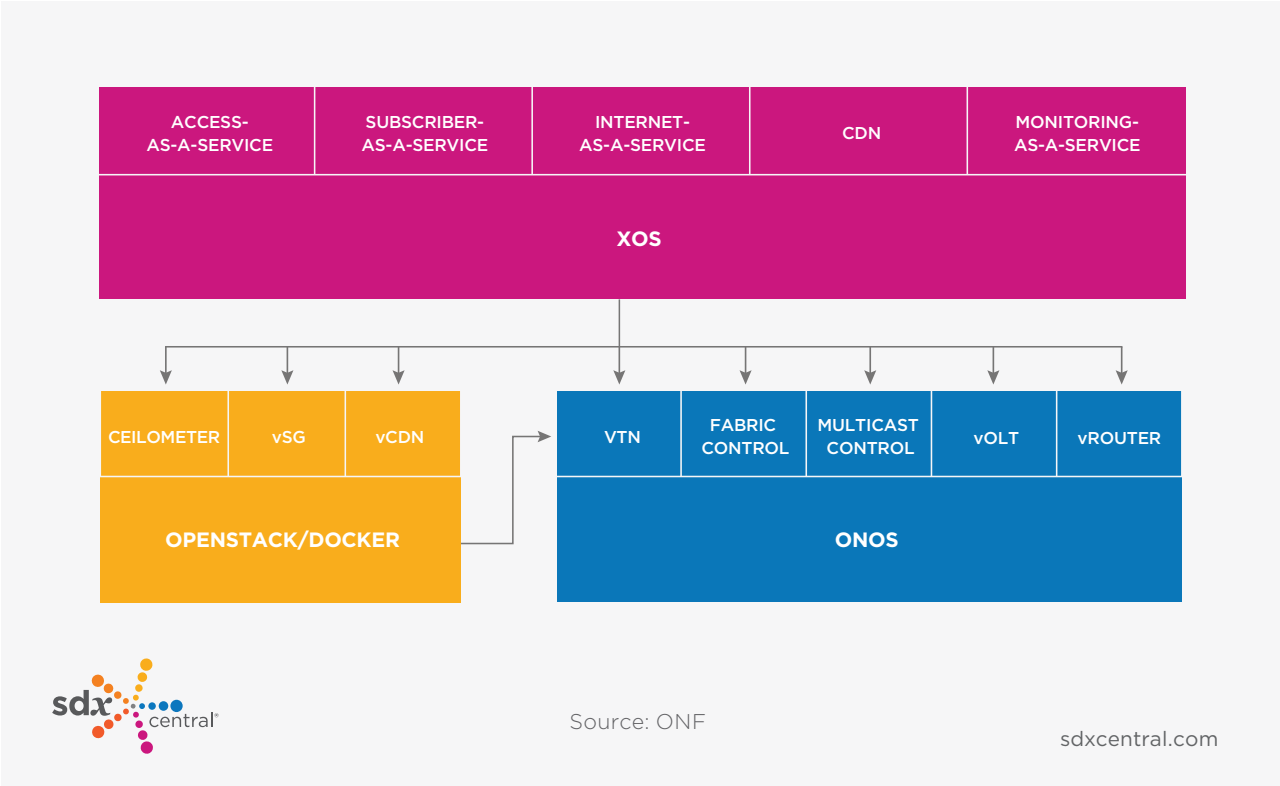

VCO

The Virtual Central Office is a project similar to CORD that focuses on delivering residential, business and mobile services within the central office. The **VCO project aims to build an "OpenDaylight-based reference architecture"** on top of OpenStack. It was first demonstrated in June 2017 at the OPNFV Summit in Beijing.

Key components of the VCO project include:

- NFVI hardware: Bare-metal switches, Open Compute Project (OCP) servers and industry standard x86 servers
- NFVI software: Cumulus bare-metal switch software, Ceph scale-out storage and RHEL with KVM hypervisor and DPDK
- VIM: OpenStack to manage virtual machines, virtual storage and virtual networks (also called overlay networks, managed in conjunction with ODL)
- SDN controller: OpenDaylight to manage underlay networks, overlay networks and VNF configuration
- VNF manager: OpenStack Tacker for managing VNFs and forming service function chains
- Service assurance: OPNFV Barometer, OPNFV Doctor, NetScout and Grafana to monitor metrics, events and alarms
- Residential vCPE VNFs: vBNG, virtual firewall (vFW), Anti-DOS, deep packet inspection (DPI) and vRouter
- Enterprise vCPE VNFs: VPN, vFW, vRouter

SDxCentral research will be monitoring both CORD and VCO to see how these projects will evolve and harmonize over time.

Open Networking Foundation

We've discussed the Open Networking Foundation in its close collaboration with the Linux Foundation on some key projects in the NFV space, including ONOS and CORD. The current ONF is merger of two organizations, the original ONF, which was focused more around standards such as OpenFlow and on use-case development, and ON.Lab, an open-source group form by researchers from Stanford, Berkeley and numerous vendor research teams. ONF is the main driver behind the ONOS and CORD projects and has recently announced additional projects in the form of Stratum and Trellis.

Stratum is an open-source implementation of a thin switch layer focused on defining the "contract" between the upper layers and what the switching data plane supports, through the use of the P4 language. It is supported by Google as well as Tencent, and includes CSPs like China Unicom, NTT, as well as networking vendors including Big Switch, white box vendors like Delta, Edgecore and QCT and merchant silicon vendors including Broadcom, Cavium, Mellanox as well as Xilinx.

In addition to the above-mentioned projects, the ONF was also a driver behind the P4 language, which is now merged as a new project under the Linux Foundation. P4 is gaining traction as a programming language to control the underlying network fabric and looks to gain momentum as a replacement for the venerable OpenFlow specification which still has traction is some networking areas but fallen mostly out of favor in the data center.

Telecom Infra Project (TIP)

The TIP, started by Facebook in 2016, is an open source project that intends to reinvent access, backhaul and core network equipment. Its focus is broader than just hardware and includes software, people, processes and other areas such as system integration and site optimization. TIP could be considered to fill an adjacency to NFV by working on hardware and software not covered by NFV (and OCP); therefore, it fills an important gap in the overall NFV transformation.

Common NFVI and VIM Architectural Deployments

Over the course of the last year, NFVI and VIM architectural deployments have not changed dramatically, and still primarily confirm to two different models. However, we expect to see increased focused on a Core + Edge model, as 5G and IoT take hold and the Edge becomes the new battleground for NFV deployments.

Centralized Cloud

For use cases like vIMS, a centralized cloud is adequate, as there is little need to have extensive services or compute capabilities close to the subscribers. In general, where latency and bandwidth are not major issues for the application, a centralized cloud will be easier to manage and cheaper to deploy.

Core DC + Edge

For certain use cases like vCPE , or increasingly, IoT applications (augmented reality, virtual reality, vehicle automation), the cloud is spread between the central data center and the edge.

The edge could be the Central Office in the case of thin CPE (i.e. the customer premise equipment is simply providing dumb connectivity) or the customer premise with thick clients (customer premise equipment that is able to host a few virtual machines).

Multi-access Edge Computing (MEC) is another use case where the cloud will likely be split between the core data center and the edge. MEC is an ETSI ISG that works on providing IT and cloud-computing capabilities within the Radio Access Network close to mobile users. The goal is to provide services close to the subscriber where requirements such as: low latency, big data analytics/ machine learning, location aware services and services that need network context are important.

Two types of implementation are possible for NFVI+VIM that straddle the core DC and the edge. In the first approach, the VIM is decentralized. In other words each edge has its own VIM. The central cloud has a federation layer that coordinates between the MANO and various VIMs. OpenStack Kingbird takes this approach.

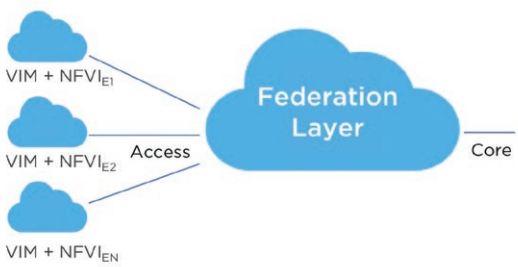

In the second approach, only one instance of the VIM exists that is deployed in the central cloud, and the NFVI is spread out across different edge clouds. The single VIM instance is able to manage this highly distributed NFVI. Nova Cells is a way to accomplish this.

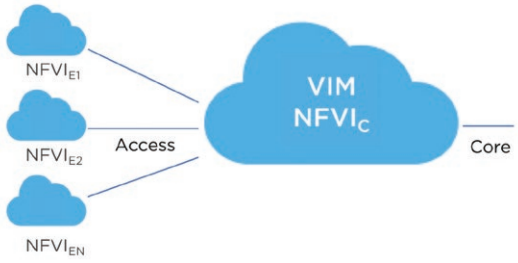

Increasingly, in our research with CSPs, some favor looking into a container-only or container-centric (with a thin VM layer for hardware abstraction) at the edge for reduced overhead. Today's telco applications tend to be VM-centric but as new applications are written and designed for the edge, we expect these will be based around modern micro-services architectures, which lend themselves well to being hosted in containers. Furthermore, we are seeing a trend at the edge around converged and hyperconverged (HCI) hardware architectures as possibilities that provide lower-cost, higher-reliability and easier deployability and manageability. At the same time, because of the nature of some of the NFV edge applications being focused around computer vision and assistive compute, GPU and other forms of HW acceleration are viewed as a key component of edge NFVI platforms.

And so, we'll expect to see different deployment architectures take root in the next few years, as the different NFV use cases go into scale production deployments worldwide. Key elements that will drive where the VIM (and NFVI) sits and how functions are partitioned across different topological and geographic locations will continue to be driven by latency requirements, bandwidth capabilities, security and compliance needs and most importantly, the nature of the use case.

2018 NFVI and VIM Market Update and Trends

We'll take a look at the recent changes in the NFV market, particularly as it relates to the evolution of NFVI and VIM. We'll start by examining the trends in cloud data centers and the general compute market, since these underlying trends are just as relevant for NFV and NFVI.

Server and Compute Market – Driven by Cloud

After a lost decade in which **sales languished** after the recessionary crash of 2008, the server market took off in the last half of 2017 with pent-up demand unleashed by a long-awaited Intel processor update. **Intel bundled the Skylake microarchitecture and a host of other enhancements** into a new portfolio of Xeon products christened the **Scalable processor family** that provided enough performance juice and workload-optimized SKUs to finally entice system buyers into opening their wallets. As we'll explain, the update was particularly propitious for CSP, cloud providers and others using x86 servers as NFV platforms since it includes **processor features** and **configurations tailored communications infrastructure**, **cloud services** and virtualization. The combination proved to be the right catalyst to ignite a languishing server market.

The latest market estimates with sales from Q4 2017 show the extent of the demand as the market exploded on a wave of sales to both hyperscale cloud vendors, which have provided the only bright spot in the market for several years, and enterprise IT. Both **Gartner** and **IDC** pegged worldwide server sales increasing more than 25 percent over a year earlier, although Gartner said units sold only increased 8.8 percent, a sign of the growing dominance of commodity 2U systems preferred by hyperscale data centers.

According to an IDC researcher, "Hyperscalers remained a central driver of volume demand in the fourth quarter with leaders such as Amazon, Facebook, and Google continuing their datacenter expansions and updates," Indeed, aside from IBM, whose numbers were skewed by a **significant mainframe product refresh**, white box ODMs were the fastest growing vendor segment according to IDC. Gartner's figures showed the same trend with **Inspur Electronics**, a favorite source of cloud vendors for commodity hardware, taking the honors for the fastest growing vendor in both Q4 and all of 2017. **Research firm Canalys found** that the total sales of all IT infrastructure hit an all-time high in 2017, up 7.3 percent for the year, which they also attribute to data center expansion at cloud-service providers. Short of a recession, which few economists expect in the near term, expect 2018 to be another record year for server, network and storage vendors.

Servers: Standard x86 Still the Foundation of NFVI

The principal reason NFV is so popular with carriers, large enterprises and cloud service providers is that it enables replacing relatively expensive, proprietary, locked-in network devices with industry standard servers and software. Indeed, sales of NFVI hardware is growing much faster than the server market writ large with **IHS-Markit estimating** that revenue from servers, storage and switches dedicated to NFV will increase from $696 million in 2016 to $3.1 billion in 2021, a CAGR of 35 percent.

Regardless the application, all market estimates make clear that server sales are still dominated by x86 hardware, with IDC pegging its share at 85 percent of server revenues. Even that lofty number is arguably down from earlier estimates that have shown Intel taking 90 percent or more of the server market. However, as we detail below, such dominance is probably a thing of the past as sophisticated cloud vendors look to new architectures such as ARM that accelerate particular workloads or reduce power, while AMD re-enters the server market with a compelling new microarchitecture and product family.

Despite the interchangeable, commodity nature of x86 systems, the server market is quite diverse and can be segmented into three significant form factors: rack servers, blade chassis and hyperconverged products. Each is a feasible platform in an NFV design, albeit only one of which is widely used by hyperscale facilities operated by cloud vendors, carriers and online service providers.

Rackmount Systems for NFVI

Rack systems, particularly compact 1U and 2U designs, are the workhorse of most NFVI implementation due to their versatility, with a wide variety of available options, density, relatively low cost and the workload profile of NFV software which doesn't place extreme demands on processing performance or memory capacity. Particular system configurations are a function of the NFV features and network throughput required and can be customized via the following key components. Here's a sample of what a CSP would view as a standard server for running NFV workloads:

Component	Example Spec	Rationale
Processor	Dual-socket **Intel Xeon Gold 6130** (16C, 32T)	Dual socket servers typically offer the best balance of price, performance and energy use.
Memory	192-384 GB DDR4-2666 RAM	Memory prices have remained stubbornly high meaning that sizing should be closely matched to your expected VM and application memory profile along with the desired number of VMs per system. For example, if your VM flavor is 8GB and you plan to run 48 VM/ server, then any memory over 384GB is a waste. CSPs are working through understanding the needs of the various VNFs to figure out the ideal mix.
Flash storage	2 x 400GB NVMe or SSD	First decide whether your NFV application requires fast storage I/O, then estimate how much.
Hard disk storage	6 x 4TB 7500 rpm SATA	Depending on the need for persistence you may or may not need internal storage (in addition to boot storage). If you do, the amount is related to the particular VM type and application usage.
Network Interface Card (NIC)	4 x 10GbE or 2 x 25GbE	I/O performance is limited by the server motherboard design and specs. Total NIC bandwidth should be sized to match aggregate I/O requirements. Dual interfaces are advised for path redundancy.
Fans and power supply		The goal of cloud infrastructure is to improve resilience via software, to achieve, for example, 99.999% service availability while using three-nines hardware. Nevertheless, depending on your scale, sparing strategy, cost of support from server vendor and whether your VNFs are cloud-native, you might consider dual-redundant fans & power supplies instead.
Management interface	Baseboard Management Controller (BMC)	Servers allow out-of-band management using a processor distinct from the main CPU that is responsible for managing the health of the server and a BMC is well worth the minor price increment.

Blades

Blade servers are effectively a packaging exercise in which buyers trade off multi-vendor heterogeneity, lock-in risk, cost and some flexibility for pre-integrated convenience and density. Although popular with some enterprises, the blade market has stagnated as cloud vendors and service providers opt for general purpose rackmount systems. Blades support modules for compute, storage and switching in a single chassis that can span 2U to 16U in size making them popular with organizations that want to combine multiple discrete systems into a single box. Blade servers generally share the same components and specification options as rack mount systems, although thermal limitations mean that the most power-hungry devices like 205W Xeon Platinum CPUs might not be supported.

Blades can be used for NFVI, particularly in smaller data centers or remote sites, however as we'll point out below, other, purpose-built NFVI products are even better fits for edge deployments.

Hyperconverged Infrastructure

Hyperconverged infrastructure (HCI) products are an alternative to blades for those prioritizing simplicity and integration but want a scale-out platform in a conventional rack form factor. Like blades, HCI products combine compute and storage into a single box, but typically lack physical (but not virtual) network features beyond standard NICs since they are typically linked using external top-of-rack (ToR) switches. Looking purely at the hardware, HCI products are almost indistinguishable from conventional high-density servers; the real difference is in software where HCI includes a hypervisor, virtualization management stack, virtual storage system and related optional features for data replication, backup, DR and security.

HCI is typically used as the foundation for virtual infrastructure, which of course includes NFVI; however, the HCI price premium, due primarily to the software stack, can leave NFVI users paying for features they don't need. Also, HCI doesn't offer the same breadth of configuration options as discrete servers and is generally more difficult to upgrade and scale since units must be purchased as a bundle, which could create a mismatch between workload requirements and available capacity leading to either poor performance or underutilization. However, in certain situations such as edge deployments for MEC use-cases, HCI is being considered due to its compactness and pre-integration for improved reliability and deployability.

Related to HCI is the new concept of composable systems. Underlying this idea is that fact that PCIe (bus technology) can be turned into a connecting fabric that can be used as a rack-level backplane for converged infrastructure. By exposing each server's motherboard bus to the entire rack, PCIe fabrics combined with hardware virtualization software enables so-called composable infrastructure that does in hardware what VMs and NV do in software

HP Enterprise and startups like Liqid have developed rack-level composable systems that allow hardware including CPU cores, memory, SSDs, even GPUs to be sliced-and-diced into arbitrarily-sized resources The appeal of composable systems is the ability to deploy custom-sized server resources without the overhead, noisy-neighbor resource contention and potential compatibility problems of a hypervisor. So far, composable products are a niche that appeals to a small group of organizations but may prove useful in MEC deployments.

Regardless, over the course of the next year, we'll be watching closely to see if HCI and composable systems gain traction for edge deployments as CSPs roll out IoT and other 5G applications and services.

Edge NFVI Deployments – The New Battleground for Intel, AMD, ARM

For most of this century, Intel has virtually had the data center market to itself, however the situation is changing with competition arriving on multiple fronts. Most significantly, since it shares the same x86 instruction set, is AMD, which re-entered the server market in a big way with the introduction of its **Epyc chip** based on a new core microarchitecture (**Zen**) and system design.

Epyc is particularly intriguing for NFVI implementations since it is designed for **1- and 2-socket systems**, but features up to 32 cores with 64 threads per socket and up to 2TB of DDR4 memory across 8 channels. Even better for network applications is Epyc's generous inclusion of 128 lanes of PCIe Gen-3 I/O capacity. Indeed, **AMD later introduced a half-sized Epyc 16C/32T variant** targeting embedded systems like NFV that supports up to 8 10GbE and 16 SATA interfaces which should be compelling for emerging NFVI scenarios such as wireless base stations and other locations on the network edge.

The Intel ecosystem is also eyeing NFVI edge deployments with compact systems using **Intel's Xeon D-series SoCs**. For example, **Dell EMC recently introduced** the **Virtual Edge Platform** based on the D-2100 processor with up to 16 cores and 32 threads and 6 network interfaces, two 10GbE, four 1GbE and a single gigabit Ethernet management channel. Although the product initially targets enterprise SD-WAN deployments in remote offices, SuperMicro has introduced a comparable set of motherboards and systems using the D-2100 and designed for SD-WAN, vCPE and NFV edge applications. Expect other manufacturers to follow with a variety of both carrier-grade servers and

lower-power edge devices, with a variety of connectivity options (Ethernet, Wi-Fi, LTE, DSL, etc.) designed for NFVI. Some of these products will likely include specialized hardware acceleration such as GPUs for video transcoding or machine learning.

The other element in edge deployments is the customer edge, where white-box disaggregated uCPE (universal CPEs) are now gaining ground with CSPs worldwide, including **Verizon** and AT&T. AT&T recently **transitioned its disggregated dNOS (disaggregated network operating system) into the Linux Foundation** as Project DANOS. At the same time, Wind River also provided its Titanium Cloud solutions into open-source as part of another Linux Foundation project, the **Akraino Edge Stack**. We will provide more detailed coverage of uCPE, vCPE and SD-WAN in our SD-WAN and Virtual Edge report publishing in the fall of 2018.

Beyond x86 architectures, let's take a look at a rising wave in NFVI, ARM architecture, which has shown increasing momentum, particularly for edge applications.

ARM, Alternative Architectures and HW Accelerators

Although x86 systems, aka Industry Standard Servers, have been synonymous the data center since the fall of RISC, Intel is finally getting squeezed on both the low and high end. Several vendors have developed server-class ARM SoCs, with systems now beginning to emerge, while IBM and its OpenPOWER ecosystem have rejuvenated the POWER architecture into a viable cloud platform. The **latest POWER9 hardware** is overkill for NFVI and primarily targeting database and AI applications; however, its integration of NVIDIA's high-speed NVLINK GPU interface means that it could be compelling for organizations combining NFVI with GPU applications, i.e. video or machine learning, particularly should scaled-down systems become available.

More interesting is ARM's potential to disrupt the NFVI market since its combination of power efficiency, high core count and multiple embedded I/O subsystems make it suitable for virtualized workloads, network applications and edge deployments. For example, **Microsoft is testing Cavium ThunderX2** and Qualcomm ARM SoCs in its **OCP-compliant platform. Cavium is also working with China Unicom** to develop an **M-CORD** NFV platform for 5G deployments that combines its Thunder ARM SoCs with the XPliant programmable Ethernet switch silicon. **Qualcomm, Mellanox and 6Wind have combined to demonstrate** a platform using Qualcomm's Centriq 2400 ARM SoC that targets **VNF using 6Wind's software**.

The Qualcomm example is notable since 6Wind was a pioneer in using the Data Plane Development Kit (DPDK) to standardize network application development using a standard, portable API that works with both x86 and ARM. DPDK is advantageous since NFV is all about packet processing where I/O performance is critical and improvements come from either CPU enhancement or offloading network processing to special-purpose hardware. Hardware offload includes things like smart NICs, with features for virtual switching, encryption, compression and port mirroring or a co-processor which often features an FPGA designed to handle particular workloads. In response to the industry's distaste for proprietary hardware acceleration interfaces, there are several emerging standards such P4 (programs switches and I/O processors) and Intel QuickAssist (interfaces to FPGAs).

DPDK is but one form of hardware data plane acceleration that also includes ODP and IO Visor. DPDK is a Linux module and libraries that bypasses network processing bottlenecks in the kernel to boost packet processing performance and throughput by up to an order of magnitude. Indeed, **Intel claims it can achieve of 347 Mpps** (million packets per second) throughput on a single Xeon processor. A separate **study** by Intel showed that DPDK improved OVS throughput by 75 percent.

In contrast to DPDK, IO Visor takes the opposite approach by residing in the kernel instead of bypassing it. It uses two methods to accelerate packet processing: eBPF, which is a programmable in-kernel virtual machine to extend kernel functionality and XDP or extended data path, which uses eBPF to create a programmable, high performance packet processor that extends the standard Linux networking stack.

Encrypted network traffic has historically been an area where NFVI systems have to worry about traffic slowdowns due to the overhead of floating point calculations. Such bottlenecks have previously been addressed by cryptographic accelerator cards that offload processing from the main processor, however recent significant advances in Xeon (and now AMD Epyc) performance, coupled with new hardware features such as AVX (advanced vector extensions) mean that there's little added CPU overhead when doing crypto operations on the primary processor. For example, **Intel benchmarks show** that the latest Xeon Scalable (Skylake) family improves crypto performance by 75 to 200 percent over the prior generation chips for AES operation and secure hashing (SHA) respectively. OpenSSL performance likewise improved by almost 50 percent on average, thus there's little reason to add specialized crypto hardware for NFVI implementations.

Storage Hardware – Going Open

NFVI can use similar storage designs to conventional enterprise systems, namely embedded, server-based storage, NAS arrays or a SAN with dedicated storage systems. Each of these can allocate virtual volumes for NFVI VMs and applications. Due to the cost and density advantages, cloud vendors and carriers have turned to storage optimized server chassis that combine a conventional 2S system with dozens of disk slots and paired with software defined storage system to manage data placement, I/O and application access.

And just like the Server components, the open-source movement (OCP) has similarly come to storage. For example, the Facebook OCP-compliant **Bryce Canyon chassis** supports 72 3.5" drives accessible by 25 Gb or 50 Gb NICs. For its part, Microsoft is focused on disaggregating storage hardware from data and volume management via its **Project Denali**. **In the Denali model**, a host CPU handles low-level functions such as address mapping and wear leveling that are traditionally done on a drive, while also managing logical storage elements such as volumes, storage networking and data replication. Such software defined storage technologies are useful in some NFVI implementations, particularly those of service providers handling large data files and streams such as video.

Switches: Proprietary versus Bare Metal + Open Software

NFVI connects into a data center switch fabric using ToR L2/L3 switches that have historically been proprietary designs with closed, or tightly controlled programming and management interfaces. Recently, bare metal, white box switches with open APIs, boot loaders and operating systems have emerged as viable alternatives, primarily due to the efforts of cloud service providers looking for faster and cheaper options to expensive enterprise gear.

As we covered in our **recent SDxCentral report on next-generation data center networking**, white box switches are conceptually similar to servers in that they are built from widely available components; however, until recently, they've robust, lacked full-featured software. White box products use commonly available switch ASICs, i.e. merchant silicon, that like a PC with Windows, disaggregates hardware from software. Bare metal devices support open network boot loaders and **install environments like ONIE** that enable users to choose from OSs like **Cumulus Linux**, **Open Network Linux** (ONL) or **PicOS**. Industry efforts to promote open hardware, such as the Open Compute Platform (OCP), have become an outlet for cloud providers to push hardware vendors to develop products adhering to their standards. OCP is also an outlet for cloud-sponsored research, such as Microsoft's **Switch Abstraction Interface** (SAI), which is now an OCP project, and improves cross-platform switch programmability.

Such bare metal systems are ideal for NFVI due to their ability to be customized, controlled via SDN platforms and low cost. As we discussed above, merchant silicon such as Cavium's XPliant and others will be a critical part of NFVI deployments at the **network edge for things like 5G** and video streaming, OTT services since they can be integrated into compact, single-chassis systems suitable for remote deployments in often-harsh environments.

Virtualization and the VIM Updates

NFVI is built on a foundation of VMs running on a hypervisor controlled by a virtualization platform like VMware vSphere or OpenStack. Choice of virtualization layer is critical since it not only provides isolation for guest OSs and

applications, but virtual network resources including interfaces and internal switches. VMware vSphere, KVM and real-time flavors of KVM are the primary options for NFV, and while the Linux-based platforms have historically been preferred due their cost, open source customizability and availability of a real-time variant, VMware is now an equally viable NFV platform.

A key reason for the ascendance of VMware in the NFV market is the company's acceptance of open source software and its work to integrate the vSphere platform with OSS software including OpenStack and the Kubernetes cluster manager. **VMware claims** to have about 100 production NFV deployments at more than 50 CSPs and carriers that collectively include more than 300 million mobile subscribers. The company is wisely targeting 5G infrastructure, which it rightly sees as driving new NFV deployments in the future.

One of VMware's strategic advantages in NFV is NSX, its virtual network overlay that unifies and centralizes network traffic management and policy controls with VM management. By providing a packaged, supported SDN platform, NSX eliminates much of the DIY installation and deployment overhead that often accompanies open source virtual networking alternative. And with NSX-T (a companion product to NSX), VMware intends to have NSX available not just on its own vSphere/ESXi platform, but also on KVM and across other systems for improved coverage across all environments. VMware has also developed a comprehensive management and security suite, vRealize, **that also supports** various Linux distros, OpenStack, Windows/Hyper-V and AWS EC2. The combination of vSphere with vSAN software defined storage, the NSX SDN controller and vRealize forms a complete NFVI + VIM + SDN controller solution.

In speaking with CSPs, though VMware is making significant inroads, OpenStack remains the preferred VIM for their NFV deployments. And OpenStack continues to advance, with the **recent Queens release (Feb 28, 2018)**, OpenStack drives forward improved container support and also provides capabilities for edge services:

- **Edge computing support** — OpenStack-Helm and LOCI, support edge computing applications as described in a recently-published whitepaper from OpenStack: "**Cloud Edge Computing: Beyond the Data Center**". Helm containerizes OpenStack services and puts them into a Kubernetes pod. LOCI provides a set of lightweight container images that reduces the footprint for images.
- **Zun container service** — a new OpenStack project that allows users to efficiently start and run containers with little overhead. It also adds networking, storage and authentication capabilities to containers by integrating with Neutron, Cinder, Keystone and other core OpenStack services.
- **LOCI** —LOCI makes **Open Container Initiative**-compatible images of OpenStack services that can work with OpenStack-Helm or used individually to deliver standalone services like Cinder block storage.

Further, the Queens release adds support for virtual GPUs (vGPUs), which might prove useful in edge deployments as AI, augmented and virtual reality use cases come to bear. Nevertheless, OpenStack has momentum, even though we still hear regularly from CSPs about the difficulty of deploying, upgrading and maintaining OpenStack. There aren't many choices beyond OpenStack (or VMware), though some CSPs are playing with the idea of going straight to containers and Kubernetes.

VMs vs Containers, VMs with Containers?

NFVI network services have historically been deployed as full VMs; however, the rapid maturation of container technology means that future NFVI deployments should seriously consider running as containerized apps instead. The industry is standardizing on the Docker container image and runtime along with the emergence of **Kubernetes as a widely supported de facto standard** container manager and orchestration platform mean that containers have broad vendor support on every significant server and cloud platform.

Intel has a white paper describing the use of containers and Kubernetes for NFV and that details enhancements it made to the core OSS stack to address technical shortcomings in networking features. These include:

- **Multus Container Networking Interface** (CNI) Plugin for multi-network interface support in a Kubernetes pod.
- **Single Root I/O virtualization (SR-IOV) CNI Plugin** to support SR-IOV devices with DPDK.
- **Node Feature Discovery** (NFD) to detect hardware features available on each node in a Kubernetes cluster, and advertises those features using node labels.
- **CPU Core Manager for Kubernetes** to provide core affinity, i.e. pinning, for NFV workloads on Kubernetes that constrains workloads to particular CPUs and cores.

Intel's work illustrates both a critical barrier to using containers for NFV, a lack of networking features, and the industry's commitment to addressing them. Similarly, the Linux Foundation **OPFNV Euphrates project** is another example of work underway to enable production-ready features and deployment tools for operating OpenStack-based NFVI.

Telcos such as **AT&T, Ericsson**, Verizon and others have been **experimenting with containers for NFV using OpenStack** for several years, but recent improvements in container technology mean that anyone with a significant NFV deployment should be doing likewise.

Over the course of 2018, the NFVI and VIM space will continue to see innovation across these areas we just covered: open-hardware both in data centers and at the edge, hyperconvergence and composable systems, x86 vs ARM, VMs vs containers, and we'll continue to track these and provide updates in the next edition of this report. In the meantime, we'll take a look at what our site members think about NFVI and VIM.

2018 SDxCentral NFVI and VIM Reader Survey

The third annual network function virtualization (NFV) survey gives us a snapshot of how the market is developing and how virtualized architecture implementations and service delivery models are maturing. The SDxCentral Research team ran the survey on the SDxCentral site during the month of March 2018[1]. In general, the responses of technology vendors were broken out from end-user customers. In addition, 2018 results were compared to 2017 and 2016 stats to try to identify trends or significant changes in perception or experiences that have occurred through the years.

NFV Interest

63% of end-user respondents indicated NFV was changing the way they architect their networks, making it "very relevant to enterprises." Another 26% felt it was "relevant," but not that different from their "data center virtualization" initiatives. Only 5% thought NFV was "only relevant to service providers, not enterprises." This demonstrates that NFV which originally started as a pure CSP initiative, does have supporters and visibility within the enterprise community.

Of those who believe NFV is relevant to enterprises, 23% noted they have already deployed NFV solutions, which is up from 15% of respondents last year. 41% claimed their organization had "not yet deployed [NFV], but was currently testing or qualifying solutions." Again, this is up from last year, when 35% said they were "testing and piloting NFV solutions."

Responses from service providers and technology vendors, based on what they have observed with customers, echo the uptick in adoption. 30% said customers were "experienced users that had deployed NFV solutions for over a year," 20% said customers had more recently deployed NFV "within the past year," and 36% said customers had "not yet deployed, but were currently testing and qualifying solutions."

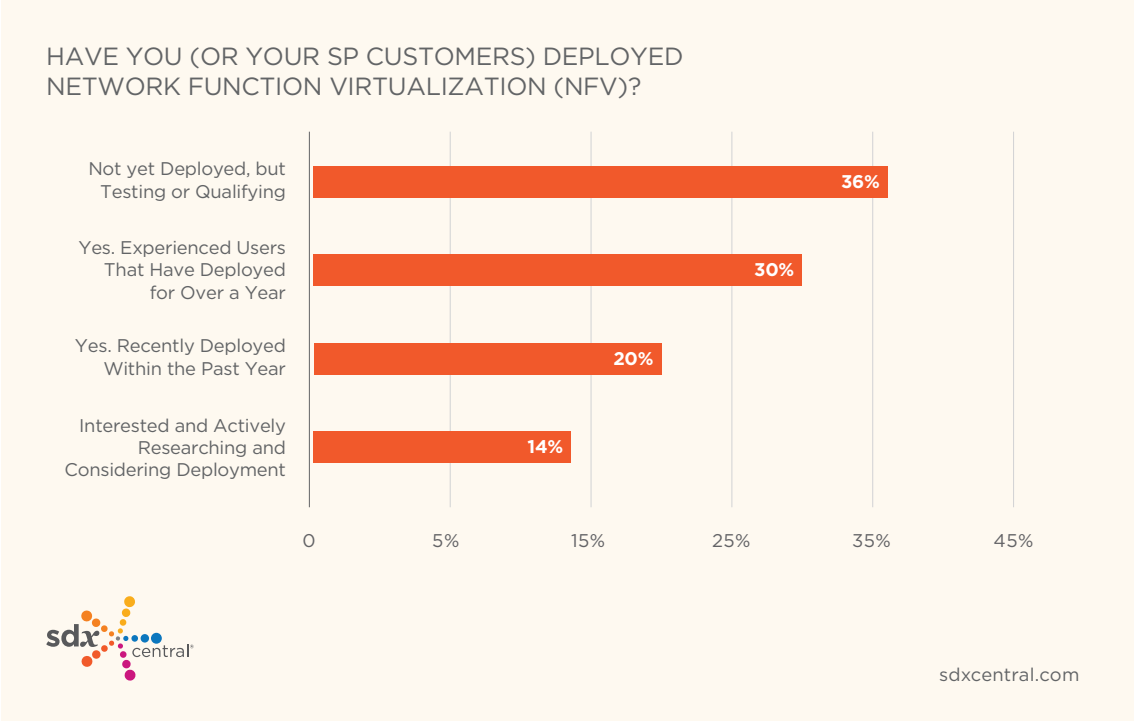

[1] There were 85 respondents with 65% end-users (enterprises and service providers) submitting for themselves and 35% technology vendors submitting information on behalf of their customers

There was a notable drop in end-user respondents who said they were "interested and actively researching and considering deployment of NFV" this year. Only 6% said they were researching, which is way down from last year, when 41% of participants indicated they were in the consideration phase of the adoption process.

Some of the drop in the research phase can be explained by the uptick in customers moving to the testing and deployment phases. The rest of the drop, however, might be that those who were considering NFV solutions last year, are no longer planning to deploy them anytime soon. 18% of this year's respondents said they "don't plan to deploy within the next year," which is significantly more than last year, when only 6% said they had determined NFV wasn't going to work for them and 3% saw no appeal in NFV technologies.

NFV Drivers

When asked an open-ended question around why NFV is applicable to the enterprise, end-user respondents noted that NFV's flexibility and potential ability to reduce CAPEX/OPEX made it relevant.

Service provider and technology vendor respondents responding on behalf of their SP customers named "reducing costs" as the biggest business driver for NFV. 34% thought that saving "money by using pre-built, validated NFV frameworks and virtualization," was why so many were interested in NFV. 29% named "agility," noting that "NFV allows for more resource mobility and efficiency and allows service deployment in hard-to-reach places." The third option that rounded out the list of the top drivers was accelerating "time to market," with 26% of respondents identifying "velocity" and "NFVs flexibility, which allows for faster service deployment," as NFV's biggest draw.

Interestingly, this is a slight difference from 2017 and 2016, when agility and flexibility topped the list. However, this is the first year that operational and capital expenses were consolidated into a single "cost reduction" choice, so that probably accounts for the difference. In 2017, 16% of vendors choose a "reduction in operational expenditures" as the top driver and 6% picked a "reduction in capital expenditures." In 2016, 23% of participants indicated that NFV's operational savings were more of a driver than potential reductions in capital spending, at 13%. When combined, cost savings would have topped the list in 2016 and taken the second spot in 2017.

Barriers to NFV Adoption

What's hindering the adoption of NFV technologies? 79% of end-user respondents and 64% of service provider and technology vendor participants indicated it was a "lack of training and in-house talent" that was stopping the deployment of NFV technologies the most.

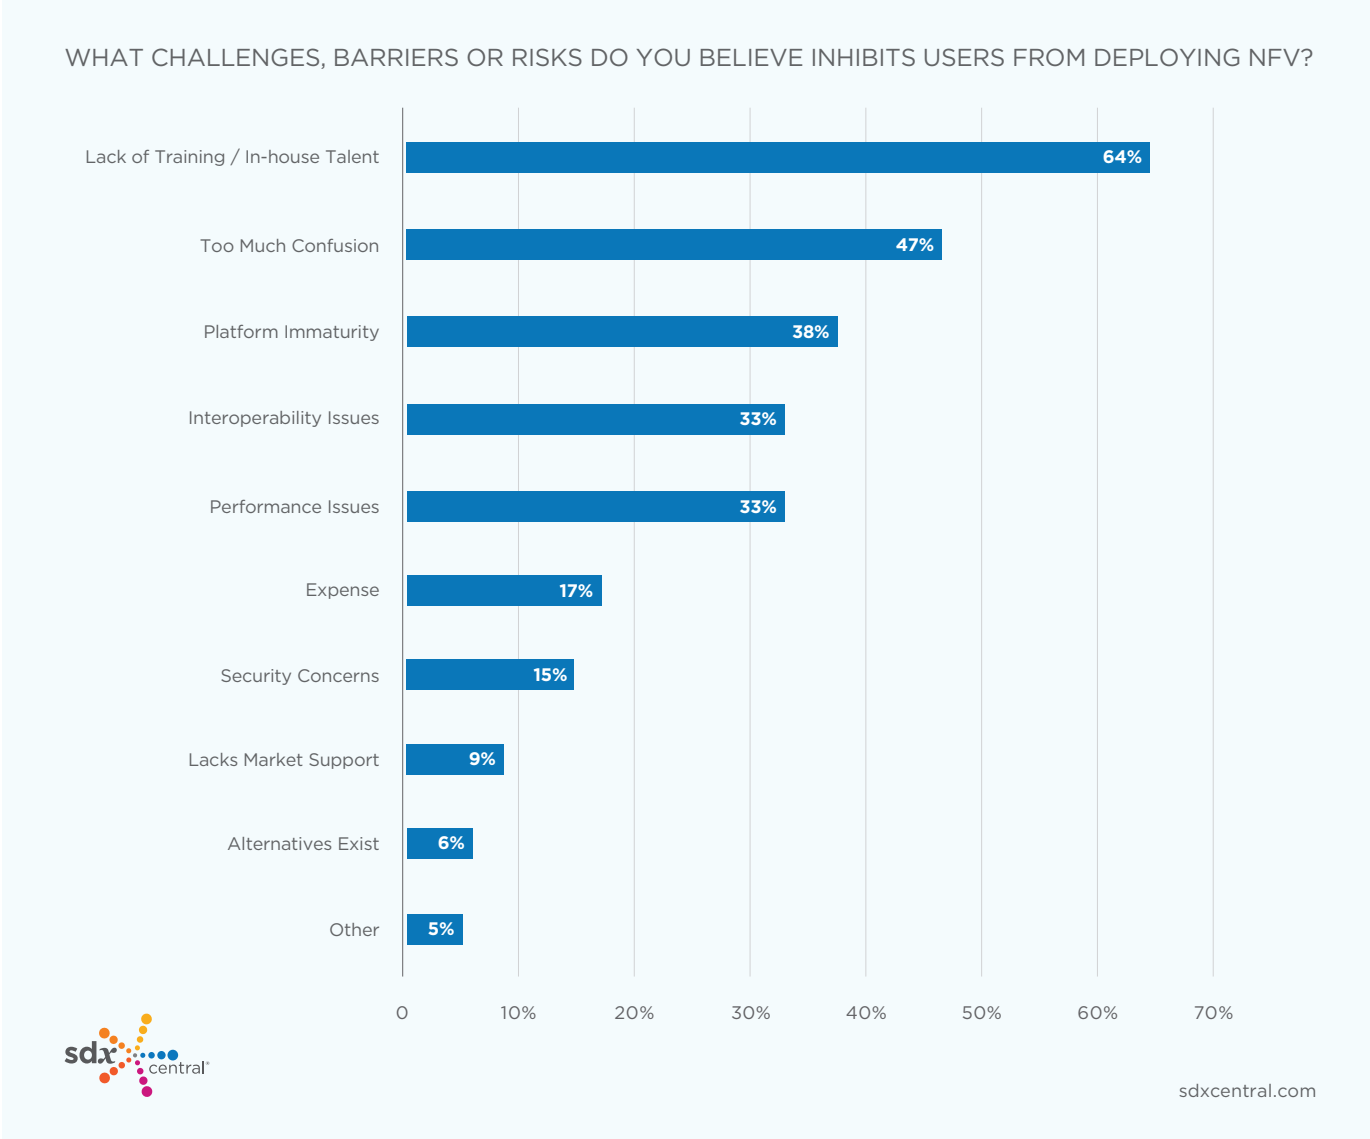

The second most cited barrier to adoption according to end-user respondents was "interoperability issues." Interestingly, last year (2017), the inability to easily "certify and integrate a multi-vendor solution," took the top spot, at 59%. Service providers felt interoperability was fourth on their list, along with "performance issues," at 33%. Instead, "too much confusion," thanks to a lack of standards, was the second barrier to adoption service providers chose, at 47%, while only 37% of end-users noted it was a big issue.

The other standout factor impeding the roll out of NFV solutions was the perception of "platform immaturity," as 32% of end-user respondents and 38% of service providers said vendor platforms are still slow to deliver on all their organization's requirements.

Containers vs VMs

Interestingly, service provider and technology vendor views are split on whether containers or virtual machines (VMs) will be the infrastructure for delivering NFV applications in the future. 47% felt containers will probably supplant VMs, but it's going to take time. At least it "will take longer than 18 months." 34% believed "containers will coexist with VMs, with no clear dominant choice." They felt that containers will grow, but VMs will retain sizeable share and the market might end up seeing more "containers run inside VMs." Only 15% thought that "containers were starting to replace VMs now and will become the default in the next 18 months."

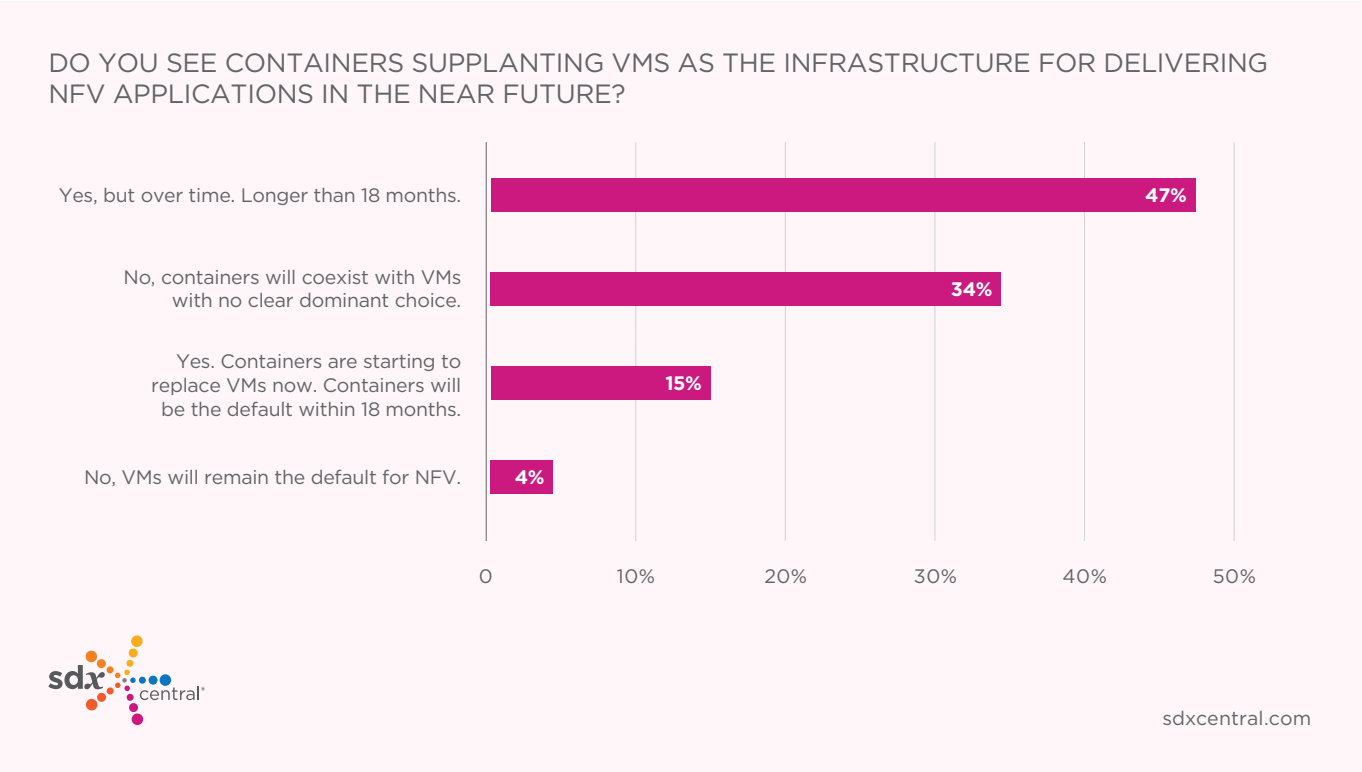

NFVI and VIM Vendors

When asked which solutions service providers currently use for their NFV infrastructure (NFVI), VMware was the overwhelming choice at 56%, which supports our observation and anecdotal evidence earlier from the CSPs that we've had conversations with. It appears that that ESXi as a hypervisor has become a viable choice for NFV deployments, and appears preferred in some situations even. The other vendors that received double-digit responses were Cisco at 35%, Ericsson at 26%, HPE at 24%, Nokia and Dell at 21% each, and Huawei at 12%.

For their NFV virtual infrastructure manager (VIM), service providers favored non-vendor-specific OpenStack for their deployments at 59%, which shows OpenStack's continued dominance and may speak to mixed/hybrid deployments. It also may be skewed by the fact that many early NFV deployments are at some of the world's largest CSPs, who either have the in-house expertise to deploy generic OpenStack and customize them for their own needs. Or they might pay

System Integrators (SI) with the expertise to do so on their behalf. As NFV uptake continues with Tier 2 and Tier 3 CSPs worldwide, we might see more choosing vendor-integrated and tested versions instead. In any case, the next most popular choices are vendor-integrated and supported versions including VMware (VIO and vCloud NFV) at 41%, RedHat's RHELOPS at 32%, Kubernetes or container alternatives at 29%, Juniper's Contrail Cloud at 24%, Nokia's OpenStack (based on Red Hat) at 21% and Ericsson at 12%.

Preview of Other Survey Results in NFV Report Parts II and III

What are the most common use cases of NFV today and where are they being deployed? What's the role of open-source in NFV? And are CSPs looking to expand their deployment of NFV or keep the status quo?

We'll cover these and more in the upcoming parts II and III of the NFV reports series, which will delve into MANO (and assurance) and VNFs respectively. Those will be released for downloads in the next 4-6 weeks, stay tuned!

NFVI and VIM Report Conclusion

Over the course of next year, we expect NFV to continue its penetration within the CSP space. There is sufficient early success that momentum will continue to build. And while we are seeing challenges in deployment, ongoing issues with overall performance, orchestration and VNF onboarding, we are confident that CSPs and technology vendors alike will resolve these issues given sufficient time. Neverthelss, as NFV spreads beyond the Tier 1 CSPs worldwide, requirements and challenges will likely change, with an increased focus on ease of deployability, manageability and upgradability.

On the NFVI front, we expect 2018 to be the year where we see continued innovation as VMs and containers battle it out, as x86 and ARM jostle for relevance and dominance, and as the Edge reveals new requirements for new application workloads and innovative HCI platforms come to the market.

We're looking forward to seeing how the market develops in 2018 and we'll be sure to watch closely and capture our observations in next year's edition of the report In the meantime, we hope you've found the content in this report useful If you have feedback for us, feel free to drop us a note at **research@sdxcentral.com**

2018 NFVI and VIM Report Products

The following section of this report profiles some of the more popular vendor offerings. As the inaugural report, we aim to cover a good sampling across the space but would welcome vendors to reach out to us. **Extended profiles of each product can be viewed online**. The information was gathered via a collaborative effort between SDxCentral's Research Team and the vendor's appropriate product experts. We welcome feedback and if you believe we've missed a product, please contact the research team to start a conversation about having your products included in future editions of this report.

While every attempt has been made to validate the capabilities listed in the profiles, SDxCentral advises end users to verify the veracity of each claim for themselves in their actual deployment environments. SDxCentral cannot be held liable for unexpected operations, damages or incorrect operation due to any inaccuracies listed here.

SDxCentral welcomes feedback and additional information from end users based on their real-world experiences with the products and technologies listed. The SDxCentral Research Team can be reached at **research@sdxcentral.com**.

Organizations with listings in this report include: 6WIND, ADVA Optical Networking, AT&T, Canonical/Ubuntu, Cavium, Inc., Ciena, Cisco, Dell EMC, Enea Software, Ericsson, Hewlett Packard Enterprise (HPE), Huawei, Intel, Kontron, Lenovo, Linux Foundation, Mellanox Technologies, Mirantis, Netronome, Nokia, Radisys, Red Hat, Telco Systems, VMware, Wind River, ZTE Corporation

Qualified Vendors Without Listed Products

The following are additional organizations with products which we believe should be included in this report but whom did not submit for various reasons. If you belong to one of these organizations, we'd love to hear from you and get you products listed in the future. And regardless of whether you belong to one of the following organizations, if you would like to submit to the report or to future versions of this report, please contact **research@sdxcentral.com**.

Advantech	MontaVista Software LLC
Amax	Silicom Ltd.
Inspur	

AT&T FlexWare

(Click to View More Details Online)

https://www.business.att.com/content/productbrochures/network-function-virtualization-product-brief.pdf

AT&T
PUBLIC | PRIVATE
208 South Akard Street
Dallas, Texas, 75202, United States
866.662.4548
http://www.att.com

Description of Company: AT&T is bringing it all together for our customers, from revolutionary smartphones to next-generation TV services and sophisticated solutions for multi-national businesses.

For more than a century, we have consistently provided innovative, reliable, high-quality products and services and excellent customer care. Our mission is to connect people with their world, everywhere they live and work, and do it better than anyone else.

AT&T in SDxCentral Company Directory

Description of Product: AT&T FlexWare, designed and deployed on the AT&T Integrated Cloud platform leveraging Software-Defined Networking (SDN) and Network Function Virtualization (NFV) technologies.

AT&T FlexWare in SDxCentral Product Directory

Value Proposition
AT&T FlexWare, designed and deployed on the AT&T Integrated Cloud platform leveraging Software-Defined Networking (SDN) and Network Function Virtualization (NFV) technologies.
Product Differentiation
Quickly deploy and configure virtual network functions from best-of-breed vendors. Improve total cost of ownership vs. existing proprietary equipment. Reduce the complexity of network infrastructure and vendor contracts. Streamline operations and simplify network management. Future proof your investment as your business evolves.

Top Use Cases
1. Security refresh: customers who are looking to refresh their firewalls and/ or add enhanced security options. 2. Equipment refresh: customers whose routers and firewalls are approaching “end-of-life”. 3. Enterprises looking to benefit from the value of multiple network services with the ability to mix and match applications (routing, security and WANx), features, and management options.

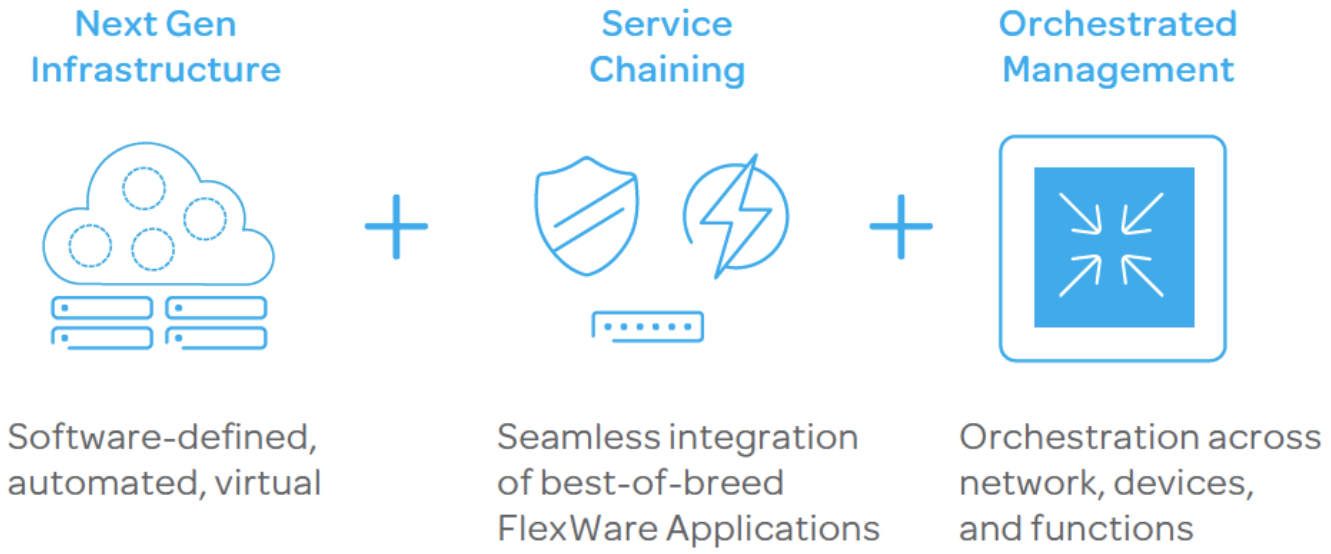

FEATURED

CATEGORY **NFVI Software**

Enea NFV Access

(Click to View More Details Online)

https://www.enea.com/globalassets/downloads/nfvi-platforms/enea-nfv-access/enea-nfv-access-application-note.pdf

Enea Software
PUBLIC | PRIVATE
P.O. Box 1033 Jan Stenbecks Torg 17
Kista, Sweden, SE-164 21, Sweden
info@enea.com
+46 8 507 140 00
http://www.enea.com

Description of Company: Enea develops the software foundation for the connected society, reducing cost and complexity at the network edge. We supply open-source based NFVI software platforms, embedded DPI software, Linux and Real-Time Operating Systems, and professional services. Solution vendors, Systems Integrators, and Service Providers use Enea to create new networking products and services faster, better and at a lower cost. 3 billion people around the globe already rely on Enea technologies in their daily lives.

Enea Software in SDxCentral Company Directory

Description of Product: Enea NFV Access is a lightweight virtualization software platform uniquely designed for deployment on edge devices at customer premise. It is streamlined for high networking performance and minimal footprints for both platform and VNFs, resulting in very high compute density. Enea NFV Access provides a foundation for vCPE agility and innovation, reducing cost and complexity for computing at the network edge.

Enea NFV Access in SDxCentral Product Directory

Value Proposition

- Accelerates time-to-market, thanks to a combination of ready-made software and deployment services
- Reduces bill-of-materials, thanks to its high compute density with advanced optimizations for both ARM and x86 white box edge devices
- Minimizes lifecycle cost, thanks to continuous integration and maintenance cycle
- Prevents vendor lock-in, based on 100% open with standard interfaces (ETSI NFV and OPNFV)
- Easy device management supporting FCAPS, and multiple interfaces for VNF management and SFC

Relevancy to NFVI/VIM

A lightweight virtualization software platform designed for deployment on edge devices at customer premise. It is streamlined for high networking performance and minimal footprints for both platform and VNFs, resulting in very high compute density.

Product Differentiation

- Minimal footprint: Designed for high compute density on edge devices and standard servers. OpenStack is not required in the standard setup
- Fast boot: Boot speed optimization for best-in-class availability
- NETCONF/YANG based device management framework supporting FCAPS functionality in the platform
- Customizable and extendable end-to-end service function chaining based on open standards

Linux Container Integration

- Supports virtualization with virtual machines and/or containers.
- Optimized KVM/QEMU and Docker Containers.
- Containerized OpenStack for solutions requiring OpenStack compatibility at customer premise.

Top Use Cases

1. uCPE deployments with centralized VNF lifecycle management and service function chaining over NETCONF/REST APIs, allowing standards based modeling and integrations.
2. Container based uCPE deployments with centralized VNF lifecycle management and service function chaining over the Docker API.

Supported Hypervisors

KVM

Industry Standards Participation

ETSI NFV ISG, Linux Foundation, OpenStack Foundation

Linux Foundation Projects Supported

LFN

Key Partners

Arm, Intel, Lanner, Advantech, Clavister

Key Customers

Active undisclosed projects: Whitebox vendors and operators in APAC region. Several engagements in EMEA and NAM through telecom systems integrators.

Characteristics	Enea NFV Access	Common uCPE solutions
Platform RAM Footprint	< 1 GB	4-12 GB
Platform Disk Footprint	< 1 GB	4-12 GB
Platform CPU Utilization	Down to single core	Down to 2-4 cores
Platform Boot Speed (Excl. BIOS)	< 3 seconds	10-30 seconds
Virtualized Network Throughput over vSwitch	10 Gb IMIX Line Rate	1 Gb IMIX Line Rate
Virtualized Network Latency over vSwitch	Average 10-15 µs	Average 25-75 µs

HPE NFV Blueprints

(Click to View More Details Online)

https://hpe.com/dsp/infrastructure

Hewlett Packard Enterprise (HPE)
PUBLIC | PRIVATE
3000 Hanover Street
Palo Alto, California, 94304, United States
csbmarketing@hpe.com
650.857.1501
http://hpe.com/dsp/transform

Description of Company: Hewlett Packard Enterprise helps CSPs transform the way they do business and grow in a fast-changing market – to become Digital Service Providers. CSPs are adopting IT and cloud technologies and methodologies to transform their network and operations. Through its portfolio of Telco-focused solutions and services, HPE helps CSPs increase network agility, enhance operations efficiency, and leverage customer insights to successfully pursue new opportunities and embrace new business models.

Hewlett Packard Enterprise (HPE) in SDxCentral Company Directory

Description of Product: HPE NFV Blueprints are validated reference configurations based on HPE infrastructure platforms and partner VIM offerings that are optimized for specific NFV use-cases. These blueprints are offered with toolkits and documentation to simplify the install and build of the NFV infrastructure stack. HPE NFV Blueprints showcase HPE's thought leadership on how to engineer open NFV-I solutions, incorporating multi-vendor components based on industry standards.

HPE NFV Blueprints in SDxCentral Product Directory

Value Proposition

HPE NFV Blueprints showcase HPE's thought leadership on open NFV-I solutions based on open industry standards.

Design Leadership: Built on HPE's proven expertise with NFVI deployments globally. These designs leverage HPE's breadth of compute, storage, networking portfolios.

Accelerate Time to Market: toolkits automate configurations of NFVI stack and can be customized by SIs as required.

Open Ecosystem: HPE's rich ecosystem of SIs, ISVs and NEPs provide SPs with choices that avoid vendor lock-in.

Relevancy to NFVI/VIM

HPE NFV Blueprints are validated reference configurations that automate and simplify the deployment of telco clouds by leveraging proven designs optimized for NFV use cases and include both the physical and virtual components of the NFVI stack.

Product Differentiation

HPE NFV Blueprints have been optimized for every layer of the NFVI stack and are designed to enable scalability via modularity, reliability with no single point of failure and accelerated performance. Leveraging infrastructure as code principles and HPE NFV best practices, HPE NFV Blueprints utilize automated deployment and configuration of the NFVI stack to help reduce complexity and eliminate human error. All system components are NEBS level 3 certified, tuned with optimal BIOS and NIC configurations for carrier-grade performance and incorporate open industry standards such as Redfish. HPE NFV Blueprints are designed and validated with close technical collaboration from leading VIM and SDN partners, and offered with end to end support.

Linux Container Integration

The virtualization layer within the HPE NFV Blueprints can be used to host Containers.

Top Use Cases

1. Core: vEPC, vIMS
2. Edge: SD-WAN
3. IOT, vFW

Supported Hypervisors

ESXi, KVM

Supported Network Acceleration

DPDK, PCI pass through, SR-IOV

Performance and Scalability

- **Maximum Compute CPU Cores (20 compute nodes):** 880 physical cores
- **Available Shared Storage Capacity with CEPH:** 26 TB
- **Total Data Throughput (Active–Active):** 2 Tbps

Automation API Support

The entire physical and virtual infrastructure can be managed by RESTful APIs. Physical APIs include Redfish API and Restful iLO. The virtual infrastructure APIs include OpenStack service APIs and SDN (Virtual Services Directory) APIs.

Industry Standards Participation

ETSI NFV ISG, Linux Foundation, OCP, OpenStack Foundation, TIP, Open19, DMTF, OpenSAF, OpenHPI, TMForum

Linux Foundation Projects Supported

LFN, Open Container Initiative

Key Partners

Intel, Red Hat, Wind River, VMware

Lenovo NFV Solutions

(Click to View More Details Online)

https://www3.lenovo.com/us/en/data-center/solutions/c/solutions?menu-id=Solutions

Lenovo
PUBLIC | PRIVATE
7001 Development Drive
Morrisville, NC, 27560, United States
hhumrickhous@lenovo.com
+19192378438
https://www3.lenovo.com/us/en/

Description of Company: Lenovo is a US$43 billion global Fortune 500 company and a leader in providing innovative data center, commercial and consumer technology. Our portfolio of high-quality, secure products and services covers servers, storage, networking, software (including ThinkSystem and ThinkAgile solutions), PCs (including the legendary Think and multimode Yoga brands), workstations, smart TVs and a family of mobile products like smartphones (including the Motorola brand), tablets and applications.

Lenovo in SDxCentral Company Directory

Description of Product: Lenovo delivers performance tuned NFVI solutions for CoSPs. Lenovo unveiled two integrated solutions at MWC'18 capable of providing robust backbones for virtualized networks powered by Red Hat OpenStack: 1) Lenovo Intel Select Solution for NFVI & 2) Lenovo Ref Architecture for NFVI. Lenovo open source solutions simplify & accelerate validation & deployment of HW & SW for NFV workloads to transform the network end to end - unencumbered by the protectionism from providers of proprietary solutions.

Lenovo NFV Solutions in SDxCentral Product Directory

Value Proposition

Lenovo enjoys a notable industry position with the capability to offer solutions to CoSPs spanning the data center to the central office out to the network edge. These two performance tuned NFVI solutions are typically deployed within the DC & integrate standard, open hardware & software suited for NFV deployments. We deliver disruptive, low cost, open source based solutions on Red Hat OpenStack that can greatly improve CoSPs ability to deliver new services in a more agile & efficient manner.

Relevancy to NFVI/VIM

Lenovo's delivery of robust, secure and performance optimized solutions for NFV workloads satisfies the need to transform infrastructures to support 5G and the IoT, enabling operators to transport massive amounts of data anywhere between any device.

Product Differentiation

Lenovo offers solutions to CoSPs spanning the data center to the central office out to the network edge. Lenovo's data center infrastructure offers major benefits for CSPs. By reducing HW requirements, they can deploy new networks faster, reduce costs significantly, and enable greater flexibility with the capability to scale with fluctuating user traffic.

Lenovo's two NFVI solutions feature built in acceleration for superior price performance and are factory integrated, turnkey packages including compute, storage, networking + virtualization software.

Lenovo is rated #1 for performance, customer satisfaction and server reliability for x86 servers (TBR Report, Dec, 2016; ITIC Global Reliability Survey, June, 2017).

Key Customers

Lenovo is engaged with several pilot projects throughout the world.

Top Use Cases

1. Lenovo's partnership, testing and validation with Intel, Mellanox, Netronome accelerated switching and packet processing solutions greatly improve network I/O performance, infrastructure efficiency and deployment flexibility.
2. Lenovo's partnership, testing and validation with Intel QAT (Quick Assist Technology) and Mellanox NIC + FPGA card support leadership encryption capabilities providing impenetrable security for data integrity and privacy protection.
3. Lenovo's partnership, testing and validation with Intel QAT dramatically compresses and decreases the amount of storage a file takes up. As a result, operators can more cost effectively and efficiently store increasingly large volumes of data.

Industry Standards Participation

ETSI NFV ISG, Linux Foundation, OCP, OpenStack Foundation, Lenovo participates in 30+ enterprise & data center standards organizations worldwide, including DMTF, Redfish, OpenHPC, Ceph.

Key Partners

Intel, Red Hat, Mellanox, Netronome, A10 Networks

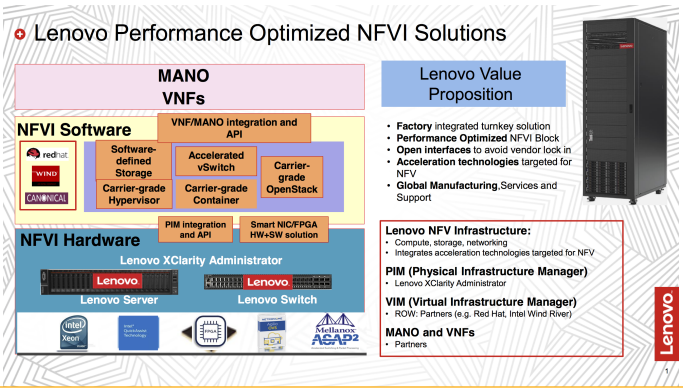

Nokia Airframe Data Center Solution

(Click to View More Details Online)

http://networks.nokia.com/solutions/airframe-data-center-solution

Nokia
PUBLIC | PRIVATE
Karaportti 3 P.O.Box 226
Espoo, FI, 02610, Finland
+358 10 44 88 000
http://networks.nokia.com/

Description of Company: Nokia is a global leader in technologies at the heart of our connected world. From the enabling infrastructure for 5G and the IoT, to emerging applications in VR and digital health, we are shaping the future of technology to transform the human experience. Powered by the research and innovation of Nokia Bell Labs, we serve communications service providers, governments, large enterprises and consumers, with the industry's most complete, end-to-end portfolio of products, services and licensing.

Nokia in SDxCentral Company Directory

Description of Product: Nokia AirFrame data center solution offers a flexible range of data center infrastructure portfolio, ranging from small implementations up to a hyperscale OCP data center. This innovative solution design is driven to enable the convergence of telco & IT domains and is geared to support a diverse range of demanding applications.

Nokia Airframe Data Center Solution in SDxCentral Product Directory

Value Proposition
Nokia is merging IT & telco cloud domains in its Nokia AirFrame data center solution. It supports IT cloud standards while also tailored for telco requirements. Our solution provides a solid foundation for NFV deployments, enabling flexibility, agility & automation needed for VNFs to perform at its best. The AirFrame solution includes SDN capable networking equipment, fully compliant w/ standard interfaces. The Nokia AirFrame data center solution is available in rackmount & OCP form factors.
Relevancy to NFVI/VIM
Complete data center solution for telco and IT
Product Differentiation
World first telco grade OCP based data center solution
Supported Hypervisors
ESXi, KVM
Supported Network Acceleration
SR-IOV, Specialized NICs, DPDK, PCI pass through
Industry Standards Participation
ETSI NFV ISG, Linux Foundation, OCP, OpenStack Foundation, TIP
Linux Foundation Projects Supported
LFN
Key Partners
Intel, Open Compute Project (OCP)

Top Use Cases
1. Scalable data center solution from hyperscsale to the local edge 2. OCP based AirFrame data center is designed for the lowest possible operational cost 3. Pre-integrated telco-grade data center solutions for CSPs, supporting applications ranging from Cloud Core to Cloud RAN
Key Customers
China Unicom, 3UK, Altan, Telenor

vCloud NFV

(Click to View More Details Online)

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/vmware-vcloud-nfv-datasheet.pdf

VMware, Inc.
PUBLIC | PRIVATE
3401 Hillview Ave
Palo Alto, California, 94304, United States
1-877-486-9273
http://www.vmware.com

Description of Company: VMware, a global leader in cloud infrastructure and business mobility, helps customers accelerate their digital transformation. VMware enables CSPs to master a software-defined approach to business, networks and IT with VMware Cross-Cloud Architecture and solutions for the data center, mobility, and security.

VMware, Inc. in SDxCentral Company Directory

Description of Product: vCloud NFV is a solution for Communications Service Providers (CSPs) who want to radically reduce cost and increase service agility through a cloud-based network deployment.

vCloud NFV in SDxCentral Product Directory

Value Proposition

VMware vCloud NFV drives CSP agility through network & operations transformation. CSPs are able to: accelerate new services & business models with a secure multi-tenant carrier-grade NFVI platform; simplify operations mgmt & deliver operational excellence with network & service visibility & monitoring; enable adaptive agile open networks with best of breed choice of ecosystem partners. VMware provides CSPs w/ a converged IT & network infrastructure for applications & services across all clouds.

Relevancy to NFVI/VIM

vCloud NFV is a modular ETSI conforming NFV platform optimized for converged multi-vendor and multi-domain NFV services. It includes extensive tooling for NFV service delivery, operations management, security, and resource management capabilities.

Product Differentiation

vCloud NFV employs a single architecture to manage multiple VIMs including vCloud Director and VMware Integrated OpenStack, adding flexibility to deployment and management, and providing secure multi-tenancy fully isolating and securing workloads.

vCloud NFV provides holistic, 360-degree visibility and insight into all network layers along with proactive and predictive analytics, intelligent log analysis, and active remediation of workloads and northbound integration with 3rd party OSS/BSS solutions.

CSPs can quickly and easily deploy vCloud NFV in their network with the current talent pool in their organization, due to the large number of VMware-certified individuals working across IT and network organizations.

Pricing

vCloud NFV is a modular multi-tenant NFV platform which includes vSphere, vSAN and NSX. vCloud NFV addresses operations management with vRealize Operations Manager and vRealize Log Insight. VIM includes vCloud Director and VMware Integrated OpenStack

Top Use Cases

1. Fixed Line: vCloud NFV enables fixed line providers to deploy SD-WAN & vCPE services reliably and at scale. With vCloud NFV, CSPs can deliver elastic bandwidth, higher QoS & highly personalized services while simplifying operations & lowering costs.
2. Mobile: vCloud NFV enables MNOs to transform their network architectures to be more agile & innovative with vEPC & vIMS solutions. MNOs use vCloud NFV to deploy cost-effective, scalable & agile solutions for differentiated new services & higher QoE.
3. IoT: vCloud NFV provides CSPs with a NFVI platform to transform network architectures, expand into new IoT verticals, such as Connected Car. With vCloud NFV, CSPs drive faster service innovation and delivery, & accelerate time to revenue.

Supported Hypervisors

ESXi

Supported Network Acceleration

DPDK, SR-IOV, Large Pages, CPU Pinning, CPU and Memory reservation and guarantee, NUMA affinity, Direct Pass Through

Industry Standards Participation

ETSI NFV ISG, Linux Foundation, OpenStack Foundation, OSM

Linux Foundation Projects Supported

Cloud Foundry, Cloud Native Computing Foundation, LFN

Key Partners

https://www.vmware.com/solutions/industry/telco-technology-partners.html

Key Customers

Telefonica, Telia, Vodafone, Ooredoo

CATEGORY **NFVI Data Center Hardware Platform**

ThunderX ARMv8 Processors

(Click to View More Details Online)

http://www.cavium.com/pdfFiles/ThunderX_PB_p12_Rev1.pdf

Cavium, Inc.
PUBLIC | PRIVATE
http://www.cavium.com/company.html

Description of Product: The ThunderX product family provides the best in class 64-bit ARMv8 Data Center and Cloud Processors, offering an unprecedented level of integration and industry leading SoC performance. The ThunderX processor families are supported by rich eco-systems including OEM/ODM COTS and OCP servers, and commercial and open source OS, virtualization, and NFV platforms

Value Proposition
The Cavium ARMv8 processors provide high performance and scalable offerings which address the various uCPE use cases from retail stores to enterprises, and Telco cloud NFV COTS servers use cases. The Cavium ARMv8 processors integrate many hardware accelerators for boosting performance and efficiency. These accelerators are abstracted by standard APIs like DPDK, FD.io, and ODP for interoperability. Cavium has rich eco-systems of ODM/OEM hardware and open source and commercial software and OS.
Relevancy to NFVI/VIM
Cavium's ARMv8 based processor families span COTS servers targeted ThunderX and ThunderX2 families to edge computing and uCPE targeted OCTEON TX SoC families. These ARMv8 processors provide high-performance hardware for NFV from Telco cloud to uCPE.

Top Use Cases
1. uCPE hardware including OEM and ODM white boxes, offering cost-effective and scalable uCPE hardware platforms. 2. COTS servers for edge and cloud datacenters running containers and virtualized VNFs, e.g. vEPC, vIMS, vBNG. 3. COTS servers for edge data centers running C-RAN vBBU.
Supported Hypervisors
KVM
Key Customers
China Unicom, China Mobile, Sprint

Dell EMC NFV Ready Bundle for VMware

(Click to View More Details Online)

http://en.community.dell.com/techcenter/service_provider_solutions/w/wiki/12384.dell-emc-nfv-ready-bundle-for-vmware

Dell EMC | Service Provider Solutions
PUBLIC | **PRIVATE**
http://www.dellemc.com/

Description of Product: The Dell EMC NFV Ready Bundle for VMware is a turnkey solution optimized to simplify and accelerate production deployments. The solution combines hardware, software, and Dell EMC engineering and is designed to create a more flexible, scalable, and agile platform for CSPs. It includes open standards-based Dell EMC cloud infrastructure hardware (compute, networking) and a choice of a Virtual Infrastructure Manager (vCloud Director or VMware Integrated OpenStack) with vSAN or Dell EMC ScaleIO.

Value Proposition
The Dell EMC NFV Ready Bundles for VMware is a turnkey solution optimized to simplify and accelerate production deployments for service providers.
Relevancy to NFVI/VIM
The Dell EMC NFV Ready Bundle for VMware integrates all of the European components into a flexible, customizable building block. Dell EMC and VMware co-engineered this infrastructure for optimized performance,

Top Use Cases
1. Telco Cloud, Webtech Service Provider
Supported Hypervisors
n/a
Key Customers
Not publicly available.

Intel Select Fast Track Kit for NFVI

(Click to View More Details Online)

https://builders.intel.com/docs/networkbuilders/intel-select-fast-track-kit-for-nfvi.pdf

Intel
PUBLIC | PRIVATE
http://www.intel.com/

Description of Product: Intel's Select Fast Track Kit for NFVI delivers workload-optimized solutions that accelerate and simplify the process of selecting the precise hardware and software for optimized solutions based on network functions virtualization infrastructure. This allows the ecosystem of partners and communications service providers a faster, optimized path to market.

Value Proposition
Intel Select Fast Track Kit for NFVI accelerates infrastructure deployment at reduced complexity and cost for a more efficient and future-ready data center operation by optimizing workloads with configurations verified to exploit performance advantages of technologies enabled by Intel Xeon processor Scalable platforms.
Relevancy to NFVI/VIM
The Intel Select Fast Track Kit for NFVI was created by analyzing the functional requirements of a wide range of NFVI use cases to create a validated and workload-optimized configuration for VNFs and other NFVI applications.

Top Use Cases
1. Simplified evaluation: Intel Select Solutions are tightly specified in terms of HW and SW components to eliminate guesswork and speed decision-making. 2. Fast and easy deployment: Increase efficiency in IT's testing process, speed time to service delivery and increase confidence in solution performance. 3. Workload-optimized performance: designed by Intel and our partners to deliver to a performance threshold for the workload and are built on the latest Intel architecture
Supported Hypervisors
KVM, QEMU
Key Customers
https://networkbuilders.intel.com/endusers

SYMKLOUD OpenStack Platform

(Click to View More Details Online)

https://symkloud.com/pdf/openstack-datasheet.pdf

Kontron (Communications Business Unit)
PUBLIC | PRIVATE
https://symkloud.com

Description of Product: The SYMKLOUD OpenStack Platform is a turnkey open source and white box solution that offers operators a disruptive deployment model to reduce opex and leverage the operational freedom to mix and match multi-vendor services without vendor lock-in. The complete OpenStack solution is packaged within a single 2U MS2900 Series converged hardware for a truly 'out-of-the-box' cloud environment.

Value Proposition
Based on the Canonical OpenStack Distribution, the SYMKLOUD OpenStack Platform enables Service Providers to deploy NFV infrastructure with greater speed and efficiency to roll-out new revenue-generating services, sooner. Kontron is committed to the open source community and is a member of the Linux Networking Fund (LNF) and OpenStack Foundation.
Relevancy to NFVI/VIM
The SYMKLOUD OpenStack Platform has been validated with a range of VNFs supplied by an ecosystem of ISV partners. The solution scales across telco central offices and datacenters, and can leverage existing compute infrastructure.

Top Use Cases
1. NFV/SDN deployments for telecom, video and hosting providers, including vEPC, vIMS, vRAN, vSBC, and other security and video delivery use cases.
Supported Hypervisors
KVM
Key Customers
Not publicly available

BlueField SmartNIC

(Click to View More Details Online)

https://www.mellanox.com/related-docs/prod_adapter_cards/PB_BlueField_Smart_NIC.pdf

Mellanox Technologies
PUBLIC | PRIVATE
https://www.mellanox.com/

Description of Product: The BlueField family of products is a highly integrated system-on-a-chip (SoC), optimized for NVMe storage systems, NFV, security systems, & embedded appliances. BlueField SmartNIC programmable network adapter card combines ARM processing power with advanced network offloads to accelerate a multitude of security, networking and storage applications, delivering world-leading performance, flexibility and efficiency.

Value Proposition
The BlueField SmartNIC adapter enables a more efficient use of compute resources, now dedicated to run applications rather than focus on networking or security processing, therefore supporting the growing demand for bandwidth per host. This software-defined adapter ensures ultimate flexibility by adapting to future protocols and features through simple software update.
Relevancy to NFVI/VIM
Using the BlueField SmartNIC for NFVi, the control and data planes of infrastructure (networking, storage, and security) are fully implemented in the smart network adapter and enables the host to focus solely on NFV processing.

Top Use Cases
1. Security Apps: Addresses concerns of modern data centers by combining HW encryption accelerators w/ embedded SW and fully integrated advanced network capabilities 2. Storage Apps: Adapter may operate as a co-processor offloading specific storage tasks from the host, isolating part of the storage media from the host, or enabling abstraction of software-defined storage logic using the BlueField Arm cores. 3. Networking Apps: Accelerated Switching & Packet Processing for Open vSwitch (ASAP2- for OVS) delivers flexible, highly efficient virtual switching & routing capabilities
Supported Hypervisors
n/a
Key Customers
http://www.mellanox.com/page/case_studies

Netronome Agilio CX 10, 25 and 40GbE SmartNICs

(Click to View More Details Online)

https://www.netronome.com/media/documents/Netronome_Corporate_Brochure.pdf

Netronome
PUBLIC | **PRIVATE**
http://netronome.com

Description of Product: Netronome Agilio CX 10, 25 and 40GbE SmartNICs transparently offload and accelerate NFVi layer processing functions such as Open vSwitch, Contrail vRouter, and Linux connection tracking, for example. Agilio CX SmartNICs are ideal for hardware acceleration and offload of OpenStack cloud networking and NFV applications in telecom service provider networks. Agilio CX SmartNICs simplify operational complexity while yielding an improvement of 3-5X in overall TCO for typical NFV workloads.

Value Proposition
Agilio CX 10/25/40GbE SmartNIC platforms fully and transparently offload virtual switch and router datapath processing for networking functions enabling compute servers used for NFV and cloud computing to save critical CPU cores for application processing while delivering significantly higher performance. The Agilio CX platform features standard low-profile PCIe SmartNICs and software, designed for x86 COTS rack servers, fitting needed OS, power and form factor requirements.
Relevancy to NFVI/VIM
The Agilio SmartNIC solution boosts NFVi efficiency by enabling each server to deliver up to 6X more output per server per dollar spent while maintaining complete vendor independence & mobility of VMs & VNFs that service next generation applications.

Top Use Cases
1. The Agilio SmartNIC platform accelerates and offloads the NFVi layer so that VNFs run up to 6X faster than traditional NICs. Enhanced telemetry capabilities in Agilio are critical for the successful deployment of 5G networks.
Supported Hypervisors
KVM
Key Customers
Cisco, Dell, Juniper, Ericsson, Symantec

CATEGORY **NFVI Other Hardware Platforms**

Ciena 3926m Service Delivery Switch

(Click to View More Details Online)

http://media.ciena.com/documents/3926m_Service_Delivery_Switch_DS.pdf

Ciena
PUBLIC | PRIVATE
http://www.ciena.com

Description of Product: Ciena's 3926m Service Delivery Switch is a compact, smart CPE platform that delivers Ethernet service capability with ultimate flexibility for modular network functions that eliminate unnecessary costs and add service agility. It is a carrier-grade VNF host based on the Service-Aware Operating System (SAOS) used in all of Ciena's packet switches, providing operational efficiency and consistent system attributes. SAOS delivers benefits across all Ethernet access and aggregation applications.

Value Proposition
The device allows the creation of value-added business or mobile backhaul services that combine connectivity with in-demand Virtual Network Functions (VNFs), as well as support for legacy TDM services. The 3926m can address today's most challenging network scenarios, providing flexibility and future-proof attributes that de-risk business decisions while allowing for fast time to market.
Relevancy to NFVI/VIM
The 3926m is a modular VNF host that facilitates large-scale NFV infrastructure.

Top Use Cases
1. Mobile backhaul infrastructure 2. Ethernet business service delivery 3. Data center NNI hand-offs
Supported Hypervisors
n/a
Key Customers
Not publicly available.

6WIND Virtual Accelerator

(Click to View More Details Online)

http://www.6wind.com/products/vrouter/6wind-virtual-accelerator/

6WIND
PUBLIC | **PRIVATE**
http://www.6wind.com/

Description of Product: 6WIND Virtual Accelerator is hypervisor acceleration software for Linux KVM that provides accelerated virtual switching and networking features for virtual infrastructures to enable Network Function Virtualization (NFV), data center virtualization, and network appliance virtualization. 6WIND Virtual Accelerator works with existing Linux and OpenStack flavors to deliver line rate processing without SR-IOV, PCI passthrough, or any changes to the infrastructure.

Value Proposition
6WIND Virtual Accelerator combines DPDK software expertise, the industry's highest Open vSwitch performance, Linux compatibility for easy integration with Linux applications and tools, a rich networking feature set with full-featured Layer 2/3 networking and switching including Flow-based (OpenFlow) and standard packet processing, complete management tools (deployment, configuration, monitoring and troubleshooting) and support.
Relevancy to NFVI/VIM
6WIND Virtual Accelerator is a vRouter deployed in KVM hypervisors for NFV Infrastructure acceleration and networking features. It delivers the virtual routing and switching features and performance required for hypervisor networking.

Top Use Cases
1. Network equipment provider customers for equipment virtualization use cases (uCPE, vEPC, Edge Computing and more) 2. Platform vendor customers with virtualized service platform use cases (OpenStack, SDN and Telecom)
Supported Hypervisors
KVM
Key Customers
Nuage, Kontron

Ensemble Connector

(Click to View More Details Online)

https://www.advaoptical.com/en/products/network-virtualization/ensemble-connector

ADVA Optical Networking
PUBLIC | PRIVATE
http://www.advaoptical.com/

Description of Product: Ensemble Connector is a high-performance Network Function Virtualization Infrastructure (NFVI) platform for hosting multi-vendor VNFs at the customer premises, in the gateway between network clouds, and in the data center. It also provides a robust network Operating System (OS) that eliminates the performance bottlenecks of traditional virtual switches, such as OVS. It operates on any white box as well as on ADVA's high-performance FSP 150 ProVM edge compute nodes.

Value Proposition
Ensemble Connector's competitive advantages include: No vendor lock-in, including hardware, VNF and MANO; Improved virtual switching with Carrier Ethernet, achieving gigabit throughput on Atom; Zero-touch provisioning (ZTP) for direct ship from vendor to end user; Improved networking with security and encryption: at Layer 2 and 3, including LTE access; Carrier deployments for more than 2 years; Multiple commercial models: Perpetual, subscription, feature-based, speed based
Relevancy to NFVI/VIM
Ensemble Connector is a high-performance NFVI for hosting multi-vendor VNFs at the customer premises, network, or data center. It provides a network OS that runs on any white box as well as on ADVA's high-performance FSP 150 ProVM nodes.

Top Use Cases
1. SD-WAN: Simplifies deployment of SD-WAN VNFs on uCPE platform. 2. uCPE: Enables high performance, optimized cost & support for any VNF. 3. Secure cloud: Can be hosted in any cloud environment, providing end-to-end security & transparency.
Supported Hypervisors
KVM
Key Customers
Verizon VNS –uCPE software, Masergy – uCPE software, Dartpoints – micro data center

Cisco Enterprise Network Functions Virtualization

Cisco Systems
PUBLIC | PRIVATE
http://www.cisco.com

(Click to View More Details Online)

https://www.cisco.com/c/en/us/solutions/collateral/enterprise-networks/enterprise-network-functions-virtualization-nfv/datasheet-c78-738570.html

Description of Product: Cisco Enterprise NFV Infrastructure Software is a lightweight virtualization platform that integrates full VM lifecycle management, monitoring, device programmability, and service chaining in a single, installable package to simplify operations and roll out new virtual network services quickly on any platform.

Value Proposition

The Cisco Enterprise Network Functions Virtualization service extends Linux by packaging additional functions to provide the flexibility and freedom of choice in next generation deployment and platform options. The NFV Infrastructure Software also supports a web-based management device portal. From this portal, the user can upload VNF packages, implement full lifecycle management, turn services up and down, connect to VNF consoles, and monitor critical parameters.

Relevancy to NFVI/VIM

The Cisco Enterprise Network Functions Virtualization service offering brings a full VM lifecycle management, monitoring, device programmability, and service chaining to Cisco, and third-party VNFs on NFV optimized platforms.

Top Use Cases

1. Increase revenue: Generate new revenue with faster availability of new, customer-facing digital services.
2. Network operations: Increase agility and flexibility to scale network services up or down per business demands.
3. Cost reduction: Reduce operational expenses by as much as 50 percent.

Supported Hypervisors

KVM

Key Customers

Vodacom, Reliance Jio, Altice (SFR), Telstra, Verizon

FusionSphere OpenStack

Huawei
PUBLIC | **PRIVATE**
http://www.huawei.com

(Click to View More Details Online)

http://e.huawei.com/en/cloud-computing/fusionsphere-openstack

Description of Product: FusionSphere OpenStack cloud platform builds a multi-virtualization platform (ESXi and Huawei XEN and KVM) and converged resource pools with automatic bare metal server provisioning. It coordinates with Huawei SDN controller to automatically orchestrate end-to-end overlay network services, which are based on hardware and software SDNs.

Value Proposition

By integrating with OpenSpack, FusionSphere supports third-party physical hardware and virtualization software, and can reuse the resources of existing data centers, improving the efficiency of scheduling and managing IT infrastructure resources. In addition to supporting cloud services based on OpenStack, FusionSphere also provides backup, disaster recovery(DR), live migration, resource scheduling across data centers, and customized telecom cloud extensions.

Relevancy to NFVI/VIM

FusionSphere Openstack consists of carrier-class server virtualization with distributed storage virtualization and network and security virtualization components to provide a unified, efficient, and open cloud operating system.

Top Use Cases

1. Automatic Network Orchestration: FusionSphere OpenStack cloud platform builds a multi-virtualization platform (ESXi & Huawei XEN & KVM) & converged resource pools w/ automatic Bare Metal Server provisioning.
2. Cloud Provisioning Service: A boot & provisioning module responsible for installing OSs on hosts, deploying service modules, monitoring service status, & arbitrating between active & standby services.
3. VM Bandwidth Configuration: Extended flavor APIs define VM bandwidth resources to meet the requirements of VNFs.

Supported Hypervisors

ESXi, KVM, XEN

Key Customers

Vodafone, Intel, Telefonica Vivo

Open Container Initiative

(Click to View More Details Online)

https://www.opencontainers.org/

Linux Foundation
PUBLIC | **PRIVATE**
http://www.linuxfoundation.org/

Description of Project: The Open Container Initiative (OCI) is a Linux Foundation project to design open standards for operating-system-level virtualization, most importantly Linux containers. OCI develops runC, a container runtime that implements their specification and serves as a basis for other higher-level tools.

Value Proposition

The Open Container Initiative (OCI) promotes a set of industry open standards and specifications around container technology. OCI's main focus is on the container image format and runtime areas where stability and standardization are more important by creating a formal specification to allow a compliant container to be portable across all major operating systems and platforms without artificial technical barriers.

Relevancy to NFVI/VIM

The Open Container Initiative (OCI) is a lightweight, open governance structure formed under the the Linux Foundation, for the express purpose of creating open industry standards around container formats and runtime.

Top Use Cases

1. Open standards based system development: Open APIs & standards will enable more vendors to build meaningful solutions that benefit the end user w/ more choice & less lock-in through easier interchangeability of components.
2. App portability: Administrators can commit to container technologies w/ their current choice of any particular infrastructure, cloud provider, devops tool, or other solution.
3. Docker support: The OCI specification is used by Docker, & it represents only about 5% of our code of the Docker platform.

Supported Hypervisors

n/a

Key Customers

Not publicly available.

Mirantis Cloud Platform

(Click to View More Details Online)

https://content.mirantis.com/Whitepaper-Mirantis-Intel-NFV-Open-Cloud_Landing-Page.html

Mirantis
PUBLIC | **PRIVATE**
https://www.mirantis.com/

Description of Product: Mirantis Cloud Platform (MCP) is a comprehensive cloud software stack that includes OpenStack (for bare metal and VM compute), Kubernetes (for containers), and related open source software. MCP enables a disaggregated NFV stack that is tuned for high performance and based on open source standards and non-proprietary infrastructure hardware. Our NFV solution comprises a hardened reference platform compliant with the ETSI-NFV reference architecture, built using 100% open source software.

Value Proposition

Because Mirantis provides 100% open source software, its customers can "disaggregate" more of their NFV technology stack, lower their costs, and avoid vendor lock-in. Mirantis Cloud Platform includes DriveTrain, a suite of open source lifecycle management tools such as Git, Gerrit, Jenkins, and Salt, to treat infrastructure as code with a centralized configuration, version control, etc. Using DriveTrain, software updates can be consumed rapidly, often with zero workload downtime.

Relevancy to NFVI/VIM

Mirantis Cloud Platform (MCP) provides open source NFVi+VIM based on OpenStack, Kubernetes, OpenContrail, Ceph, and more. By working with the largest global communications companies, Mirantis has optimized its platform for NFV use cases.

Top Use Cases

1. Within the ETSI architecture, Mirantis Cloud Platform provides the NFVi & VIM Components. For NFVi, OpenStack is used for virtualized computing, networking, & storage. For VIM, Mirantis Cloud Platform includes an SDN controller.
2. In a vCPE use case, MCP will provide the NFVi while VNFs from multiple vendors are enabled for validation, & an ONAP-compliant solution is used for management & orchestration.
3. Using the DriveTrain LCM within MCP, VNFs can be continuously delivered using the same toolchain as the infrastructure. VNFs can be coordinated using multiple orchestrators, & cloud providers can be used as targets.

Supported Hypervisors

KVM

Key Customers

AT&T, Vodafone, Reliance Jio, Ericsson, Saudi Telecom

FlowEngine

(Click to View More Details Online)

http://www.radisys.com/assets/flowengine-tde-2000

Radisys
PUBLIC | PRIVATE
http://www.radisys.com/

Description of Product: FlowEngine TDE-2000 is a High Density L2-L5 networking device that helps service providers to scale with agility, reduce network complexity, accelerate the move to SDN and NFV functionalities, and reduce CapEx & OpEx. It operates in data center & telecom environments that demand high network throughput, scalability and guaranteed performance in a small form factor. It provides load balancing, service function chaining and network packet brokering capabilities for critical service delivery needs.

Value Proposition
FlowEngine is easy to integrate with any SDN controller, so the customer is not locked with one SDN. FlowEngine delivers the industry's highest performance traffic distribution and supports extended DevOps for service readiness. Additionally, FlowEngine supports multiple functionalities within a single system: switching, load balancing, and service function chaining.
Relevancy to NFVI/VIM
FlowEngine implements load balancing algorithms to provide operators highly programmable access to load balancing policies & access control rules. It combines switching, load balancing, security, and routing functionality under open policy control.

Top Use Cases
1. FlowEngine enables intelligent flow management & dynamic service function chaining. It also offers a performance SDN data plane network service platform that can be programmed with OpenFlow controller and deliver up to 1.2 Tbps data plane processing.
Supported Hypervisors
KVM
Key Customers
Verizon

Red Hat OpenStack Platform

(Click to View More Details Online)

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/10/html-single/network_functions_virtualization_product_guide/

Red Hat
PUBLIC | PRIVATE
http://www.redhat.com/

Description of Product: Red Hat delivers a complete, open software foundation for CSPs. Red Hat OpenStack Platform provides the foundation to build a private or public IaaS cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

Value Proposition
The future of Telco is open: Red Hat's open telco framework connects three essential platforms- data management, application development and infrastructure- on which modern CSPs rely for their digital transformation. Red Hat's NFV Solution, based on contributions to OpenStack, KVM, and DPDK, is open source and standards-based; it underpins the framework and is strengthened by Red Hat's broad ecosystem of certified partners to provide a stable and interoperable foundation on which to build.
Relevancy to NFVI/VIM
NFVI is built on Red Hat OpenStack Platform and is the foundation for CSPs' network transformation projects. Red Hat CloudForms can be used to present information from multiple VIMs. Red Hat Enterprise Linux is a co-engineered host and guest OS.

Top Use Cases
1. EPC - Mobile operators have several use cases for Virtual Packet Core depending on what services they provide. Red Hat has published a ref. architecture for deploying the mobile packet core on Red Hat OpenStack Platform as the NFVI. 2. VoLTE - Red Hat OpenStack Platform is well situated to help deploy VoLTE projects by providing the infrastructure and scaling needs of VoLTE and other operator services that rely on IMS and functional service chains. 3. VCO - Red Hat's built a Virtual Central Office solution based on Red Hat OpenStack Platform & ODL for NFV to create solutions supporting residential, business & mobile.
Supported Hypervisors
KVM, Red Hat Virtualization
Key Customers
Orange, Verizon, Turkcell, Sprint, MyRepublic

NFVTime

(Click to View More Details Online)

http://www.telco.com/index.php?page=nfvtime-turn-key-vcpe-solution-suite

Telco Systems
PUBLIC | **PRIVATE**
http://www.telco.com

Description of Product: NFVTime is an open and neutral full NFV service environment that allows telecom service providers and network integrators a smooth launch path to NFV services. With NFVTime, SPs will easily provide their customers multiple managed, on-demand, business services that scale and grow with their needs.

Value Proposition

Distributed environments: supports distributed deployment models and operational environments

Any VNF: gives complete agility to select and run any application and use any type of whitebox without limitations or being locked to single vendor

Run on x86 and Arm architectures: runs on any hardware architecture utilizing the best performance and hardware infrastructure available

Rapid deployment and short time-to-market: includes all required component essential for deployment and operation

Relevancy to NFVI/VIM

NFVTime is complete NFV edge solution that allows service providers and enterprises to deploy NFV services in a short time-to-market with a minimal initial investment.

Top Use Cases

1. Open CPE: NFVTime NFVi turns any whitebox device into an open uCPE capable of running any VNF or application designed for KVM/OpenStack environment
2. SD-WAN & secure SD-WAN: runs any SD-WAN SW on its NFVi w/ out-of-the-box support & can be service chaines to vSec to add robust security using best of bread solutions
3. MEC: NFVi-OS supports Multi Access (Mobile) Edge Compute use cases, providing all the required hypervisor and device lifecycle management need for supporting vEDGE and MEC applications

Supported Hypervisors

KVM

Key Customers

Not publicly available.

Titanium Cloud Product Portfolio

(Click to View More Details Online)

http://www.windriver.com/products/product-overviews/Titanium-Cloud-Product-Overview-TiS/

Wind River
PUBLIC | PRIVATE
http://www.windriver.com/

Description of Product: Wind River Titanium Cloud is a fully integrated, ultra reliable, and deployment-ready portfolio of virtualization software products. The portfolio includes Titanium Core for large scale telecommunications data center deployments, Titanium Edge and Titanium Edge SX for dual and single node edge deployments, as well as Titanium Control for industrial control systems.

Value Proposition

The Titanium Cloud portfolio includes the industry's only fully integrated, ultrareliable, and deployment-ready virtualization platforms that enable service providers to deploy virtualized services faster, at lower cost and with guaranteed uptime. When service uptime is critical for profitability, Titanium Cloud products ensure virtualized services run when, where and how they need to, always.

Relevancy to NFVI/VIM

Titanium Cloud is a pre-integrated solution comprised of open source components such as OpenStack, Linux, Ceph, KVM, and DPDK. Customers benefit enormously from the time-to-market advantage of an integrated solution along with the ease of use.

Top Use Cases

1. MEC: Working with hardware MEC application providers, the Titanium Cloud portfolio enables MEC solutions to meet the high performance and ultra-low latency requirements.
2. C-RAN and V-RAN: In collaboration with customers such as Altiostar and Nokia, the Titanium Cloud portfolio provides software infrastructure that meets the low latency, high reliability needs of C-RAN and V-RAN applications
3. Virtual Business CPE: By leveraging the Titanium Edge platform, service providers can deploy cost-sensitive NFV applications such as a virtual business CPE (vBCPE) with high reliability on only two servers.

Supported Hypervisors

KVM

Key Customers

https://www.windriver.com/products/titanium-cloud/telco/#featured_customers

Ubuntu OpenStack

(Click to View More Details Online)

http://insights.ubuntu.com/wp-content/uploads/Canonical_Distribution_Factsheet_Web.pdf

Canonical / Ubuntu
PUBLIC | **PRIVATE**
http://www.ubuntu.com/

Description of Product: The Canonical Distribution of Ubuntu OpenStack is the fastest and easiest way to build an OpenStack cloud. It is based on Ubuntu, the most popular platform for OpenStack deployments, and the leading Ubuntu OpenStack reference architecture. It brings together all the components included in the Ubuntu OpenStack Distribution, plus the convenience of automation through the intuitive, web-based UI, and the the choice of which components to use for hypervisor, storage, and networking.

Value Proposition

Ubuntu OpenStack is a fully integrated, optimized combination of Ubuntu Server LTS, OpenStack Kilo, Juno or Icehouse and powerful tools to deploy, manage and scale cloud and datacenter networks. With new support for the LXD hypervisor, Ubuntu can provide 15 times the density for KVM for workloads, tremendously improving cost efficiencies in the cloud.

Relevancy to NFVI/VIM

Ubuntu is one of the most popular operating system for OpenStack. Designed to be the fastest, most reliable way to build and launch enterprise-scale clouds from provisioning to deployment and management.

Top Use Cases

1. OpEx Savings: Canonical's OpenStack Autopilot handles your production cloud across multiple physical machines through installation, expansion, and everyday operations.
2. General compute cloud: Run your web applications in a private cloud while also meeting auto-scaling requirements to cope with variable increases in demand.
3. Reduce Operational Expenditures: Flexwebhosting is one of the largest web hosting providers in netherlands with more than 60,000 customers.

Supported Hypervisors

KVM

Key Customers

Deutsche Telekom, Sky, LexisNexis, Time Warner Cable, Bloomberg

Ericsson Cloud Execution Environment

(Click to View More Details Online)

http://www.ericsson.com/ourportfolio/products/cloud-execution-environment

Ericsson
PUBLIC | PRIVATE
http://www.ericsson.com

Description of Product: Ericsson Cloud Execution Environment is the VIM in Ericsson's system verified NFVi solution. It provides deterministic performance while securing an always available cloud and NFVi environment. Ericsson Cloud Execution Environment is based on OpenStack with added features not yet available in open source, to meet the carrier grade needs of telecom service providers. The solution supports operators in quickly deploying virtual telecom, OSS, BSS, IT and media workloads while keeping TCO low.

Value Proposition

Telecom grade VIM capable of handling the most demanding operator requirements and mission critical services.

Other Features: High throughput with low latency, through an enhanced vSwitch and support for SR-IOV; Continuous end-user access to services through redundancy and high availability capabilities; Trusted tenant isolation; A unified cloud infrastructure operations & maintenance through fault and performance management

Relevancy to NFVI/VIM

The product is included in Ericsson's system verified NFVi solution which aligns with the ETSI NFV Architectural Framework and the ETSI Management and Orchestration (MANO) architecture. The product has been contracted with more than 120 customers.

Top Use Cases

1. Universal VIM: Run any workload, anywhere (edge, central office, datacenter) with highly optimized control plane ready for 5G
2. Automation: Built-in all aspects of operations for simplicity and low total cost of ownership
3. Multi-vendor capable: Completely open integration points for workloads, MANO and hardware

Supported Hypervisors

KVM

Key Customers

Swisscom, Telefonica, Telstra, NTT DOCOMO, Vodafone Hutchison

ZTE TECS

(Click to View More Details Online)

https://sdnfv.zte.com.cn/en/products/NFVI

ZTE Corporation
PUBLIC | PRIVATE
http://wwwen.zte.com.cn/en/

Description of Product: ZTE TECS is based on OpenStack and complies with Network Functions Virtualization (NFV) architecture, implements centralized scheduling, and management for virtualized infrastructure resources through the unified interfaces. In addition, the telco enhancements has been made on performance, reliability and security, to meet the requirements of telco cloud.

TECS includes six components: TECS Director, TECS OpenStack, TECS OpenPalette, TECS Compute, TECS CloveStorage and TECS Network.

Value Proposition
Carrier-grade features such as Huge Pages, NUMA, CPU Pinning, SR-IOV, DPDK,High Availability. Integrates PIM(Physical Infrastructure management) into VIM; batch settings of BIOS, boot sequence and network connection.Simplified Deployment and Operation.
Relevancy to NFVI/VIM
The NFV industry generally adopts OpenStack as the first platform of NFVI. TECS is based on OpenStack with the NFV accelerated features such as Huge Pages, NUMA, CPU Pinning, SR-IOV, DPDK. TECS passed all tests by the OpenStack Powered Platform.

Top Use Cases
1. China Mobile NFVI+vIMS and vRCS 2. China Telecom cloud management system 3. VEON 7 countries (Russia, Ukraine, Kazakhstan, Algeria, Bangladesh, Georgia, Pakistan) unified NFVI
Supported Hypervisors
ESXi, KVM
Key Customers
China Mobile, China Telecom, Telecom Austria, VEON, Telefonica

SDxCentral, LLC
3511 Ringsby Ct, #101
Denver, CO 80216 USA
www.sdxcentral.com

The Trusted News and Resource Site for SDx, SDN, NFV, Cloud and Virtualization Infrastructure

Rev B