

Align Operations for Network Functions Virtualization Environments

Achieve Operations Equilibrium: A Strategy for an NFV Operating Model



Contents

What You Will Learn	3
Introduction	3
NFV Operating Environment	4
Computing and Storage Resources	4
Network Resources	4
Virtualization and Abstraction	5
OpenStack	5
Tenants	6
VNF Tenant Architecture	6
Target Operation Support System Environment	6
Continuous Integration and Deployment: Agile Processes and DevOps	6
Strategy for an NFV Operating Model	7
VNF Service Management	7
Organizational Changes	7
Order Fulfillment	8
Automation and Orchestration	9
Assurance	9
Fault and Operations Management	10
Security	10
Billing	11
NFV DevOps	11
Asset and Configuration Management	11
Change and Release Management	11
Performance and Capacity Management	12
DevOps Tool Capability	12
Conclusion	13
Cisco Service Offerings	13
For More Information	13
Acknowledgements	13

What You Will Learn

Organizations are rapidly migrating to network functions virtualization (NFV) to address the competitive challenges of cost and time to market. This transformation requires significant changes in products, standards, and interoperability. But perhaps the most crucial challenge of NFV deployment is alignment with the operational model. This document identifies the operational challenges of NFV and offers best practices for developing an NFV operations strategy.

Introduction


Network functions virtualization (NFV) offers significant business benefits to an organization, but it also poses challenges to the traditional operating environment. In a traditional hardware-centric environment, topology changes are infrequent. But in a virtualized environment, the pace will be faster than many processes that require manual operations can accommodate. Organizations need to be able to dynamically discover, create, allocate, and reconfigure resources.

Traditional operation models enable operators to perform basic tasks such as root-cause analysis and impact analysis when faults occur. In a virtualized environment, the challenges are greater, with operators needing to monitor the entire network for application faults and performance. Operators also need to be able to automatically add resources when performance degrades, to meet service-level agreements (SLAs).

Traditionally, service operations have been challenging for operators because of the need to span multiple technology domains. In a virtualized environment, the challenge is magnified because in addition to spanning multiple technology domains, service operations also must cover both virtualized and physical infrastructure, and they must reach across networks to cover the entire fabric.

This document presents operational considerations and an operation model strategy for organizations deploying an NFV environment. Figure 1 shows the components of the NFV architecture discussed in this document and their relationship to the business.

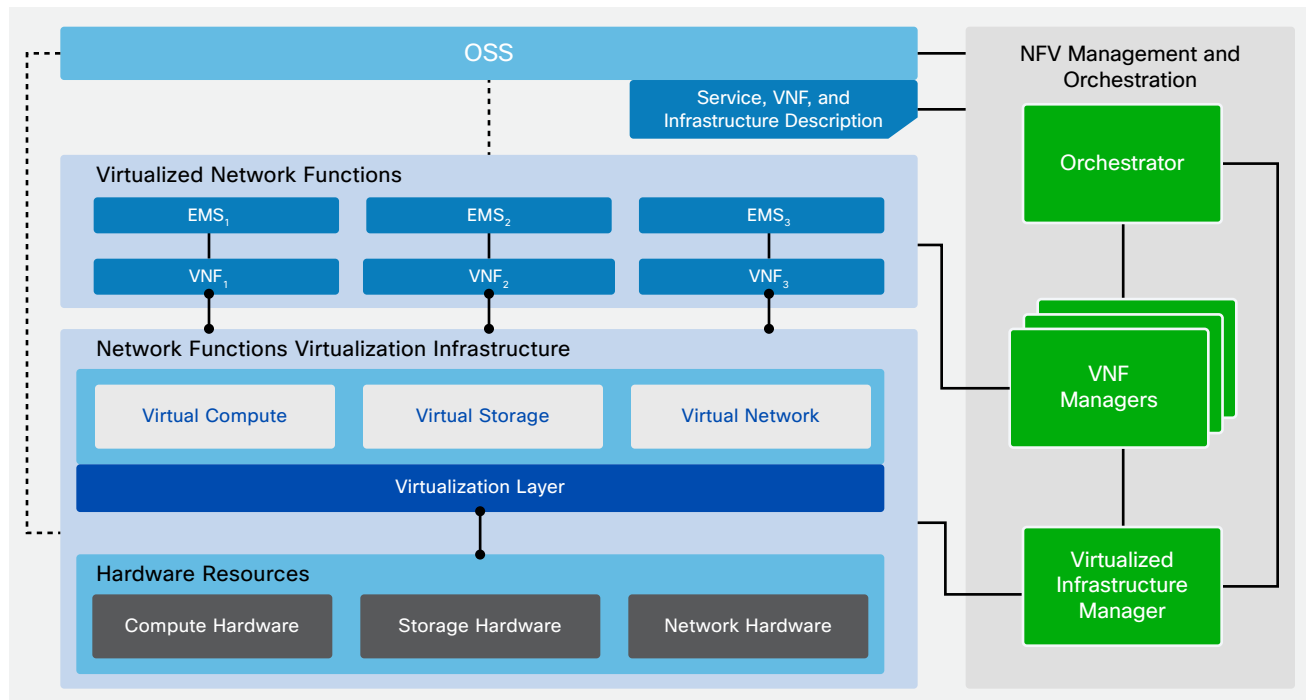
Figure 1. NFV Multidimensional Architecture



The major challenge in operating an NFV environment is the new operating model required to run a service chain consisting of one or more virtual network functions (VNFs). New approaches are needed to address increasing software development and management dependencies, rapid changes, and operating methods used in virtualized environments. New capabilities are needed to manage development and release of VNF and NFV components. Other new capabilities are needed to improve visibility into virtualized aspects of the solution and the linkage to physical-component counterparts. You can buy technology and tools, but the new pace of the data center will require changes in your organizational structure, levels of expertise, mindset, and processes. This transformation will take work and time.

Figure 2 shows an example of a framework provided by the European Telecommunications Standards Institute (ETSI¹) NFV Industry Specifications Group, a consortium of service providers and vendors that address NFV architectures and orchestration for NFV. This document also takes into consideration the ETSI Management and Orchestration (MANO) framework that categorizes the requirements for an NFV architecture into three functional layers to provide guidance in aligning the operational environment.

Figure 2. ETSI NFV Framework



NFV Operating Environment

NFV adoption introduces new multivendor elements into the operating environment that are needed for automation, workflow, and virtualization. These elements include processes, tools, and expertise that integrate into existing provisioning, operations, and lifecycle management. Essentially, the NFV environment consists of bare-metal storage, network, and computing resources that can be managed in a virtualized way on a virtualization and abstraction layer using automation and workflow tools.

After the physical and virtual platforms are enabled, services can be easily and even dynamically enabled on the platform to support rapid service enhancement, dynamic provisioning, and efficient operations.

Computing and Storage Resources

One of the main expected benefits of NFV is the capability to move from a purpose-built application-specific integrated circuit (ASIC) model to a commodity x86 server. In reality, some use cases may not be appropriate to run on low-end virtual machines,

and may require more computing power from virtual CPUs (vCPUs), memory, and network throughput.

The new virtualized platform must provide excellent performance for virtualized appliances and services, as well as power to run automated tasks (for example, Intelligent Platform Management Interface [IPMI]). It must also interface with next-generation assurance systems (Simple Network Management Protocol [SNMP] and NETCONF and YANG) in a scalable and efficient way. A centralized management console will facilitate capacity planning and forecasting.

Network Resources

A service provider's core business is to manage a variety of networks and services. The adoption of a software-defined network (SDN) introduces another layer—a network overlay or virtual network—that will require additional expertise from operations and engineering, as well as tools that can provide the end-to-end visibility to the network operations center (NOC) so that incidents can be resolved quickly and efficiently.

¹ Source: http://www.etsi.org/deliver/etsi_gs/nfv/001_099/002/01.01.01_60/gs_nfv002v010101p.pdf

Virtualization and Abstraction

One of the main components of NFV is the capability to virtualize and abstract computing, storage, and network resources. Today's data centers run tens of thousands of virtual machines and are moving toward yet another level of virtualization based on Docker containers, potentially supporting in the range of a hundred thousand virtual machines. Additionally, Gartner predicts, "By 2018, converged infrastructure systems categorized as 'hyper convergence infrastructure' will represent 35 percent of total converged infrastructure shipments."

The level of virtualization being proposed to support the new set of applications in the NFV space will require significant change in the operating model. New tools and solutions are already evolving, reusing components from data center virtualization that are already mature and delivering value repurposed for NFV. However, the operationalization of that technology often lags as a result of technology complexity, changes in processes, and comfort with existing processes and tools.

Concepts such as the following may need to be inserted into existing processes and teams:

- Agility and DevOps
- Model-based orchestration (NETCONF and YANG)
- Situation management
- vCPUs
- Virtual network interface cards (vNICs)
- Network underlays and overlays
- Virtual forwarders (Open vSwitch [OVS], Cisco® Vector Packet Processing [VPP], Cisco Virtual Topology Forwarder [VTF], etc.)
- Intel Data Path Development Kit (DPDK) and VPP
- SDN controllers

Organizations need to understand how to plan for this new operating model, in which network elements share resources that are not controlled or designed for that dedicated function. New organizational models that break down traditional silos are needed. And organizations will need to deploy, manage, and maintain the new virtualized infrastructure next to the current existing platform and support a hybrid infrastructure.

Virtualization also accelerates network element provisioning and resource sharing for resources that are underutilized. New provisioning methods and

roles and responsibilities are needed, and automation may take over some traditional processes. The new model will increase economic efficiency, and service providers and their vendors may also have new activities to help ensure that capacity is available for consumption of computing, memory, and network resources in the virtualized environment.

OpenStack

OpenStack is becoming widely adopted in organizations, because it offers an opportunity for cost savings and for large ecosystems to enhance their capabilities, time to market, and potential use cases. NFV was developed for openness with a broad support community, similar to OpenStack.

However, although OpenStack helps reduce costs for organizations and improves the speed of innovation, it also introduces new operational challenges for the pace of typical service provider operations. The frequency of OpenStack releases and the potential increase in the frequency with which VNFs need to be upgraded can affect the change and release management and deployment processes in place today.

Because OpenStack is built on open-source projects, ideally a service provider should rely on a distribution partner that can provide a carrier-class solution that meets platform consistency and supportability requirements. These subscription contracts provide support for the platform for upgrade releases, and manage performance, security, and availability according to agreed-on service-level definitions.

Because NFV infrastructure (NFVI) environments must be elastic and grow based on demand, organizations should choose OpenStack distributors that provide capabilities to automate the installation of the operating system in computing nodes added to a cluster, allowing rapid growth.

In addition, an upgrade process is needed that not only allows the code to be kept current (to reduce bugs and security vulnerabilities and take advantage of new features), but also to allow existing VNFs to operate normally after an upgrade is completed. Sometimes organizations may need to rely on more than one version of an environment at a time, so the operation of multiple environments needs to be addressed. Availability requirements will play an important role in determining when upgrades can occur without outages or downtime, if needed.

Tenants

NFVI operators have unique challenges managing a shared environment for multiple tenants because each service can have unique requirements within a shared space. The NFVI operator must first have a well-defined VNF lifecycle process for adding, operating, and retiring services, with well-defined service objectives and agreements for releases, changes, security, capacity management, etc.

VNF service component owners may also have current process requirements that differ from the agile methods used by NFVI operations. Process integration and a strong understanding of operational activities by VNF owners and NFVI operators are needed. A recommended approach is to develop well-defined roles and responsibilities for each entity, including NFVI engineering and operations staff, VNF owners, and tenant engineering groups.

Responsible, accountable, consulted, and informed (RACI) charts showing areas of responsibility can help ensure that role, responsibility, and communication requirements are met. A common approach is for the NFVI operator to act as an organization (within the larger organization) to help ensure that NFVI agile processes can be employed for rapid change and adoption of virtual services.

VNF Tenant Architecture

Tenant architecture is needed to help manage onboarding, release, production, and retirement of VNF applications from multiple sources and vendors. A flexible tenant architecture provides workflow tools for the onboarding, testing, production, and archive phases of a VNF. The architecture must allow multiple vendors to support end-to-end services (service chaining), and provide production-level verification for performance, scalability, security, availability, and manageability.

Considerations include the following:

- Isolation of tenant (vendor) code (VNFs)
- Tenant code storage
- Sharing of code requirements among vendors
- Validation of vendor specifications
- Quota policy
- Forecasting and elasticity policies
- Archive policy
- Microsoft Active Directory and login portal access control lists (ACLs)

- Authentication and authorization
- Operations testing and handover
- Tenant architecture support

Target Operation Support System Environment

Target operation support system (OSS) environment considerations include the following:

- Coexistence of multivendor bare-metal and virtualized platforms (with multiple versions)
- OSS and business support system (BSS) environments for automation and orchestration tasks
- Domain and cross-domain orchestration acting on different layers to provide scalability and support rapid service creation and deployment

Continuous Integration and Deployment: Agile Processes and DevOps

Considerations for continuous integration and continuous deployment in implementing agile processes and DevOps methodologies includes the following:

- The NFV target operating environment must be more dynamic than a traditional physical network function (PNF) environment, providing the capability to easily enable features and upgrade images.
- VNF version control is needed. You must decide whether to keep multiple versions or have the VNF manager (VNFM) or NFV orchestrator (NFVO) destroy and re-create VNF appliances with new releases.
- You need determine how to optimize processes for continuous integration and continuous deployment, shorter release cycles, and agile practices.
- Frequent software upgrades (VNFs, OpenStack, etc.) require code control, automated testing, and continuous deployment of new functions.
- Improved speed and ease is needed to add new NFV environments, quickly test and validate services, and quickly deliver new services and capabilities.
- Bimodal IT and SLAs and different classes of services (noncritical and mission critical) require different processes and policies.
- Processes need to be developed to operate sandboxes in the NFVI with various approval gates.

Strategy for an NFV Operating Model

The main goals of NFV are to promote rapid service introduction, efficient operations, and highly available elastic services. However, to achieve these goals, an organization must have an NFV operating model strategy for operational transformation. In this document, the phases of operational transformation are referred to as Day-0, Day-1, and Day-2 to help describe the planning, development, and production phases.

- **Day-0:** Understand the requirements; plan the strategy and roadmap; align business goals; and prioritize use cases, resources, and budget.
- **Day-1:** Develop a lab-ready platform environment with trained resources able to install applications and services and build operational capabilities. Define and design roles, responsibilities, processes, and tool interactions.
- **Day-2:** Roll out the production environment, with applications running in the production environment and ongoing service improvements, upgrades, and updates.

For each phase of operational transformation, a service provider needs to build a strategy based on new or transformed operational capabilities. The operational capabilities described here are some of the primary functions for which operators will need to plan and implement a new (or transformed) service model.

Note that this document uses multiple industry frameworks that best exemplify the operations transformation capabilities needed for NFV.

VNF Service Management

Most organizations have agreements, or at least a working set of assumptions, that describe how an underlying infrastructure service is managed in relation to a service or application. With VNF, Service Management includes a new type of underlying platform (NFVI) and migration from physical network functions to VNFs, as well as perhaps new people and a new organizational structure. These changes necessitate a new set of assumptions and agreements about how the NFVI platform and VNF service chains are managed.

The best way to address these changes is to build a Service Management process for platform and application owners in which all assumptions are verified to help ensure that groups agree on the service parameters. These parameters typically include goals for availability, mean time to repair (MTTR), change notification and approval expectations, VNF onboarding expectations, key performance indicators (KPIs), and governance. Parties included in this process may be infrastructure and platform owners, VNF owners, VNF engineering groups, and overall service owners.

In a Service Management arrangement, interested parties meet on a regular basis to share information about metrics, risks, concerns, and opportunities for the service. Initial meetings for VNF Service Management may include the following:

- A well-defined and agreed-upon onboarding process with defined KPIs
- Defined and documented service goals for fault, release, change, capacity, and quality for the VNF service
- KPIs for Service Management
- Ongoing Service Management processes for service review, governance, investment, and service improvement

Organizational Changes

Organizations typically change or build new groups to better support the new NFVI and NFV-based solution architectures. These groups may be unique or may initially contain virtual team members from existing infrastructure and service groups that have unique skill sets in areas such as OpenStack, programming, and scripting.

When defining organizational capabilities, several factors need to be considered for best results:

- **Definition of roles and responsibilities:** Make sure that each group or entity understands its roles and responsibilities to prevent gaps, confusion, and overlap between groups.
- **Skills enablement:** Understand the new skills that will be needed for success in areas such as OpenStack, automation, and orchestration.
- **Bimodal IT:** Establish a separation of duties, so that one new team manages the NFVI infrastructure, and other teams manage existing infrastructure until opportunities for integration are identified.

Sometimes the NFV project may introduce new services based on new VNFs. These are considered completely new, or greenfield, deployments, which are easier to implement with little or no impact on existing services. Greenfield deployment can be a good option for organizations introducing virtualization in their infrastructure. Organizations can create a new service and define the processes and tools to support it with less concern for existing infrastructure. However, in cases in which services rely on both physical and virtual appliances (for example, physical and virtual solutions in a service chain), a hybrid environment may be required, with additional planning needed for deployment and operation.

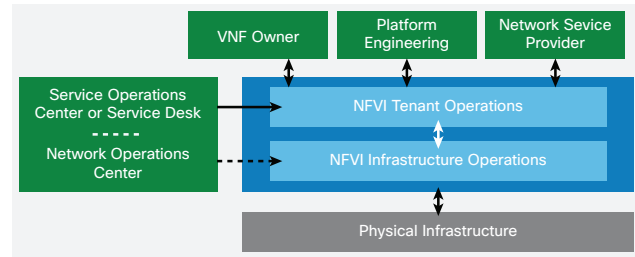
In cases in which VNFs are replacing PNFs, traffic may need to be migrated (offloaded) from physical to virtual services (or a mix of the two). Depending on the case, the strategy could be focused on freeing resources from a physical appliance to a virtual one (for example, moving the routing-reflection function from a physical router to a virtual appliance), while keeping the PNF responsible for heavy traffic.

In some cases, when the PNF is approaching end of life or retirement, organizations can consider a complete migration of the PNF to VNF. The most important aspect to be considered is the capacity of the VNF to handle the same load as the PNF, helping ensure that elasticity is available to provide horizontal scalability and thus maintaining the expected quality of service (QoS).

The rate of change in the new environment must also be considered for operations, even in bimodal IT organizations, because the organization will still have existing processing requirements. Identifying change windows and pre-approved changes for the new environment will be important in maintaining the desired agility and rate of change. Many organizations use new IT concepts to achieve the desired operational outcomes:

- **Fast IT:** This technology modifies lifecycle process to improve service agility, provisioning, and time to market.
- **NFVI as a service:** Organizations trying to get the most from a common NFVI platform for different business units may find it more effective to have a common NFVI operations team to support tenant (VNF owner) operations teams. Figure 3 shows a model.

Figure 3. NFVI-as-a-Service Proposed Model




Order Fulfillment

Order fulfillment is one of the most important steps in the service provider’s lead-to-cash process. The capability to autoprovision resources, as well as modify and terminate them, is one of the main benefits for which organizations adopt NFV.

This level of automation can require significant integration of OSS and BSS solutions, with some elements provisioned at a more abstract level with open APIs and enterprise service bus (ESB) technology. This level of integration applies to both order fulfillment and assurance, for which the orchestrator is the primary controller.

Figure 4 shows the abstraction between different layers of OSS and resource-facing components.

Figure 4. OSS Abstraction Concept



Abstraction helps reduce integration costs, support increased service definition changes, and improve operation results. In an NFV environment, order fulfillment capabilities typically interface with the orchestration layer so that communication can flow bidirectionally through the use of APIs.

Automation and Orchestration

Automation and orchestration are the primary functions for providing service agility, service quality, and rapid provisioning. However, there are complexities involved in integrating, managing, and operating these new tool sets. Additional expertise and changes in roles and responsibilities are also required.

Tool sets and processes may include the following:

- Enterprise orchestration
- Provisioning and activation
- Continuous deployment workflow
- Artifacts management
- Source-code management
- Continuous integration and testing
- Application lifecycle management
- Controller provisioning
- Service catalogs
- Monitoring (fault, service level, capacity, etc.)

In the next-generation operations environment, new roles will be created in the existing service provider organization. One important role, the automation engineer, will add new capabilities to the organization with new skill and process requirements, including the following:

- Knowledge of Linux
- Experience with scripting, automation, and orchestration


- Source-code management
- Experience in using representational state transfer (REST) APIs
- Virtualization experience
- SDN experience
- Agile methods
- Test-based development
- Write unit and integration tests

Assurance

To achieve the desired levels of service quality and performance, assurance capabilities will need to be reassessed. For instance, in an automation and virtualization environment, the main actor is the orchestrator. The orchestrator is responsible for assuring the services being provisioned, as well as for performing remediation and optimization tasks.

In the new model, the orchestrator interfaces with both new and existing tools and processes. One newer layer of the assurance architecture is a data collection and aggregation layer. This layer is critical for managing information coming from inventory, events, alarms, logs, and metrics to help ensure that data can be stored and consumed for real-time service and inventory dashboards, service-impact analysis, fault and cause analysis, and log and performance analysis. Figure 5 shows an assurance architectural approach for managing next-generation service provider service platforms.

Figure 5. Proposed Architecture for NFV Assurance



Fault and Operations Management

Processes are similar for fault and operations management, but new tools, alarms, and expertise are needed to rapidly resolve fault conditions. Operators need to understand new outage types involving virtual entities and the linkage between virtual elements and physical devices. Operators need to plan for fault and operations capabilities by understanding new fault messages and required actions. In general, Day-0, Day-1, and Day-2 activities are needed that build and optimize fault capabilities over time.

To help ensure a smooth transition to Day-2 fault and operations management, operators should consider overall lifecycle strategies that address these changes:

- **Day-0:** Tools strategy and lab requirements for capability development
- **Day-1:** Fault and operations capability development for the minimum viable product
- **Day-2:** Operations and fault management training and service improvement
- **Day-2:** Metrics such as percentage of network privilege framework (NPF) or percentage of service impact identified proactively (for service improvement)

Security

An NFV environment can be susceptible to a variety of security risks that need to be considered as part of the solution design, and that will need to be monitored and managed on a daily basis. The ETSI has published some guidelines in ETSI GS NFV-SEC that are specific to NFV. In addition, many considerations that also apply to virtualization and multitenancy technologies will likely affect overall operations of such environments:

- **Loss of availability:** Denial-of-service (DoS) attacks on virtual appliances can affect operations in the data and control planes.
- **Loss of confidentiality:** In shared virtualized infrastructure, data leakage and eavesdropping are important operational challenges to be addressed.
- **Loss of integrity:** Multitenant environments can suffer from internal attacks such as Domain Name System (DNS) redirection and IP spoofing. Unauthorized changes to equipment parameters can also profoundly affect operations.

- **Loss of control:** Some techniques used by attackers can compromise control of resources, leading to long outages and significant service impact. Recovery can be difficult, with potential loss of data. Some techniques can allow attackers to control the network and compromised VNFs, enabling unauthorized configuration changes and theft of service.

Organizations need to apply efficient security practices that are relevant to physical, virtual, and shared environments, such as the following:

- Properly handle certificates and store private keys.
- Apply security patches to operating system, hypervisors, and VNFs to secure vulnerabilities and thwart exploitation mechanisms.
- Implement security within virtual machines using network overlays and firewalls.
- Secure API and user credentials.
- Secure VNF appliances with strong authentication methods.
- Block physical access to equipment and storage network devices.
- Implement security zones for virtual machines according to their functions.
- Harden the operating system and hypervisors by disabling unnecessary ports and applying local virtual firewalls.
- Use encryption to protect sensitive information and VNF images.
- Separate customer-facing VNFs and infrastructure VNFs in different OpenStack clusters to provide an additional level of operational security.
- Provide security event and incident management capabilities through a security operations center (SOC).

Authentication, Authorization, and Accounting for NFV

NFV has many infrastructure, virtualization, middleware, management, and orchestration components, and authentication, authorization, and accounting (AAA) requirements across these domains will differ. In many cases, third-party operators will have access to components to be able to use and administer portions of the deployment, including AAA components. These various domains and actors generally necessitate trust boundaries and appropriate administration with trust-certificate authority for individual actors for specific components.

VNF Security and Multitenancy

VNF security encompasses several security domains. Intra-VNF security is needed where security policy, authentication, key management, credentialing, and encryption are unique to a specific VNF. Inter-VNF security is needed for shared services used by the VNFs, including network services such as DNS, IP addressing, time services, certificate authorities, orchestration, and identity management. In addition, VNFs have their own trust boundaries to logically separate VNF service traffic from other VNF traffic. Concerns about legal intercepts and related complications need to be addressed. Use cases, such as backups, patching, software updates, and changes, should also be considered in relation to both intra- and inter-VNF security concerns.

Service-Level Agreements and Regulatory Compliance

Agreements with customers and other operational entities within the organization should include security policy. SLAs may also require analytics that require access to VNF and NFV resources. Regulatory compliance must also be considered to identify where physical resources may be needed.

Threat Vectors, Monitoring, and Detection

Security professionals need to investigate potential threat vectors posed by a more complex virtualized infrastructure and platform system with an increasing number of actors. Considerations for forensic analysis should be investigated to identify potential physical resources.

Billing

In the planning phase, an organization may need to implement a more complex chargeback model to meet the needs of the shared resource environment. In general, a combination of customer-facing services requires tenant metering and shared infrastructure consumption chargeback. Resources used by customers may need to be accounted for and sent to a billing system for service charging. Resources consumed by a tenant may also need to be reported for cost allocation or internal chargeback. Both cases may require tools to capture data from individual components of the NFV stack and send the data to service provider billing or accounting systems.

With the potential risk of dramatically increased resource use, organizations need to keep tight control over use to help ensure that operational objectives and budgets are met.

NFV DevOps

NFV operations need to focus on IT Information Library (ITIL) processes and the Enhanced Telecommunications Operations Map (eTOM) framework and align with DevOps needs. The main areas to consider here are the following:

- Asset and Configuration Management
- Change and Release Management
- Performance and Capacity Management

Asset and Configuration Management

Up to this point, operators have relied on mainly static asset and configuration management systems that rely on human input to extract configuration data. This data was used primarily for physical lifecycle management. In most cases these systems are ill-equipped to deal with real-time virtual elements that may be needed for fault or change management.

A strategy is needed to identify and understand specific requirements for real-time configuration items that address these types of issues:

- What are the uses and value of real-time data in operational processes such as fault and change management?
- How and where will virtual configuration items and linkages to physical elements be stored and exported to OSS and BSS layers? Typically, this processes happen within the orchestrator.
- What is the orchestrator roadmap for maintaining real-time configuration and asset management data?
- What are the integration roadmap requirements for existing asset and configuration systems with regards to multitenancy?
- What new or existing tool sets are needed for real-time linkage between elements in the NFV environment?

Change and Release Management

Overall, the rate of change and potential interactions will increase in the VNF environment. Operators will need to integrate with existing processes while building new ones for the release of platform updates (NFVI) or VNFs.

To integrate with existing change management processes, operators need to understand the risk. To avoid confusion and delay in production, Day-1 planning should include a complete understanding and categorization of VNF and NFV changes and how they can be performed using a continuous integration and continuous deployment methodology.

Typically, these changes include the following:

- Platform infrastructure upgrades
- Infrastructure and VNF platform software upgrades
- Infrastructure software patches and security fixes
- Addition of new VNFs
- VNF instance movement, addition, and deletion

VNF release management requires a new set of activities, tools, turnover requirements, and communications involving new groups, such as VNF owners, VNF engineers, tenant operators, and platform operators, in VNF lifecycle activities (onboarding, ongoing release management, and retirement). Table 1 later in this document lists next-generation OSS tools needed for change and release management.

Day-1 activities should then include the following goals to efficiently implement new, more rapid release management and deployment processes:

- Identification of Day-0, Day-1, and Day-2 processes and interactions
- Well-defined roles and responsibilities documented in a RACI chart
- Release process flows and tool interactions
- Turnover requirements among groups
- Service expectations (from the Service Management team)

Performance and Capacity Management

Understanding the performance and capacity of network functions running on commodity hardware is a challenge best met with capacity planning, continuous monitoring, and dynamic consumption management.

Commodity computing, storage, and network virtualization introduces several new performance and capacity variables, including vCPUs, memory, and vNICs. These new variables make it difficult to understand how applications will perform in the environment. Organizations will need to either rely heavily on testing and predictive analysis or ask vendors to provide validated designs and metrics that can be replicated.

Best practices for performance and capacity management include a number of standard Day-1 and Day-2 processes such as baselining, trending, alarming, forecasting, and what-if analysis.

The major difference in the NFVI environment is in understanding and managing the performance of VNF applications, by monitoring errors, application problems, and capacity requirements. In general, traditional tools are still used to manage the performance and capacity of physical elements.

Organizations can start by understanding and baselining VNF performance in the NFVI environment and capturing NFVI and NFV logical performance thresholds in addition to physical thresholds for bandwidth, CPU, memory, and storage.

A recommended set of Day-1 and Day-2 processes include the following:

- **Day-1:**
 - Capture infrastructure and NFVI platform performance and capacity thresholds (can be physical or logical limitations).
 - Capture VNF performance and capacity capabilities, as well as limitations and thresholds in relation to overall service quality and performance.
- **Day-2:**
 - Collect, monitor, apply trend analysis, and report relevant capacity KPIs. Combine these processes with existing performance and capacity processes.
 - Use a forecast process to identify future requirements. Use forecasting for capacity management.
 - Perform capacity management using baselining, forecasting, trending, what-if analysis, and alarm analytics to make decisions about capacity additions and modifications.
 - Incorporate the concept of elasticity to right-size subscription levels.

DevOps Tool Capability

Hardware release cycles are measured in years, but VNFs and virtualization components (OpenStack, Kernel-based Virtual Machine [KVM], etc.) have upgrade cycles measured in days, weeks, or months. To keep up with changes in the releases, organizations need to develop or modify change management processes to allow these changes to occur rapidly while meeting the demands of the business. Organizations also need to adopt a new generation of OSS tools that can handle these challenges more easily.

Table 1 lists some of these new tools available and describes their importance for a DevOps model, including continuous integration and continuous deployment.

Table 1. Next-Generation OSS Tools

Tool	Description
Source control	This tool is responsible for keeping consistent the configurations and customizations performed by the vendor on top of the release changes implemented by the open-source community or by the VNF vendor.
Code repository	A repository manager that supports secure, clustered, high-availability is essential to helping ensure that code changes are stored, reused, and deployed properly in an automated and dynamic environment.
Automated testing	Automate the regression and functional testing of code builds overnight before applying them in production environments, and perform testing in the production environment for real-time validation and active production monitoring.
Bug and feature tracking	The capability to manage, report, and address problems is fundamental to providing confidence in and control over rapid deployment and release cycles. Automated integration between testing and bug tracking is needed to help ensure quick responses and error handling.

Conclusion

The process of deploying an NFV solution can differ for each organization. Every environment contains different variables—platforms, vendors, and equipment—and different teams are responsible for the various aspects of organization infrastructure management. The strategy an organization defines for deploying and expanding NFV is crucial to the organization’s perception of the technology adoption.

Regardless of the deployment approach used, to operate an NFV environment the service provider must align the operation processes, tools, and organization to deliver services quickly and efficiently to meet market demand and stay competitive.

Organizations also have a defined set of tools to manage and monitor the deployment of all the NFV components so that any impact on business services can be easily identified and resolved across all the elements of the solution: VNFs, virtualization resources, and underlying infrastructure.

Cisco Service Offerings

Cisco Advanced Services offerings include strategy, planning, and building services to help you navigate this new virtualized, automated landscape. Please contact us at itsmo-bdm@cisco.com.

For More Information

<http://www.cisco.com/c/en/us/solutions/service-provider/network-functions-virtualization-nfv/index.html>

A complete list of the acronyms is available at: [http://docwiki.cisco.com/wiki/Category:Internetworking_Terms_and_Acronyms \(ITA\)](http://docwiki.cisco.com/wiki/Category:Internetworking_Terms_and_Acronyms_(ITA))

Acknowledgements

This white paper was produced by:

- Renato Fichmann - rfichman@cisco.com
- Jake Hartinger - jharting@cisco.com
- Saurabh Jain - saujain@cisco.com

The following contributors helped review and align the content:

- Rajiv Asati - rajiva@cisco.com
- Fatih Ayvaz - faayvaz@cisco.com
- Carmelo Califano - ccalifan@cisco.com
- David Schofield - daschofi@cisco.com