Cisco NFVi and Unified Management with CMS

Use Cisco Network Functions Virtualization Infrastructure and Unified Management with Cisco Cloud and Managed Services Solutions to Monitor and Help Ensure Capacity for Virtual Network Services

This white paper presents at a high level the Cisco® Network Functions Virtualization Infrastructure (NFVi) solutions, discusses the operational challenges and proposes an approach to monitoring and keeping track of capacity: from cluster-level trending and planning to node and individual virtual services capacity and performance management through Cisco Cloud and Managed Services (CMS) Solutions practices and tools.

Overview

Virtualization has greatly increased the speed of server and application lifecycles. Today, servers and applications can be deployed in minutes, and network and services need to keep up the pace. With Cisco® Network Function Virtualization Infrastructure (NFVi) solutions, virtual network services can be deployed as easily as a new server.

But virtualization and automation also bring new operational complexity and challenges.

Cisco Cloud and Managed Services (CMS) Solutions maintain a holistic view of the NFVi elements and the solution components, including hardware, computing, and virtual network functions (VNFs), keeping track of performance and capacity and proactively managing incidents that can affect the health of the solution.

As part of the NFVi management process, CMS proactively monitors resource allocation on the NFVi nodes and keeps track of VNF resource load and utilization, providing comprehensive visibility for management and reporting.

Traditional Network Services Approach

In many data centers, network services run on dedicated appliances.

This approach locks the customer in to particular vendors and does not provide much flexibility. Operations teams often have to wait two to four weeks—sometimes even months—for new hardware and services activation.

Other challenges to the traditional network services approach include:

- The need to keep pace with the server team to deploy new network services and features quickly
- Complexity and overhead requirements
- Product and support costs and access to resources
- The need for familiarity with dedicated hardware appliances
- The need to order, cable, and rack-and-stack hardware appliances
- Underutilized dedicated appliances
- Limited and potentially costly options for high availability
- Unnecessary routing of traffic through one location

Flexible, Faster, and Easier Service Deployment with Cisco NFVi

Cisco NFVi provides a better approach (Figure 1).

Figure 1. Cisco NFVi Benefits

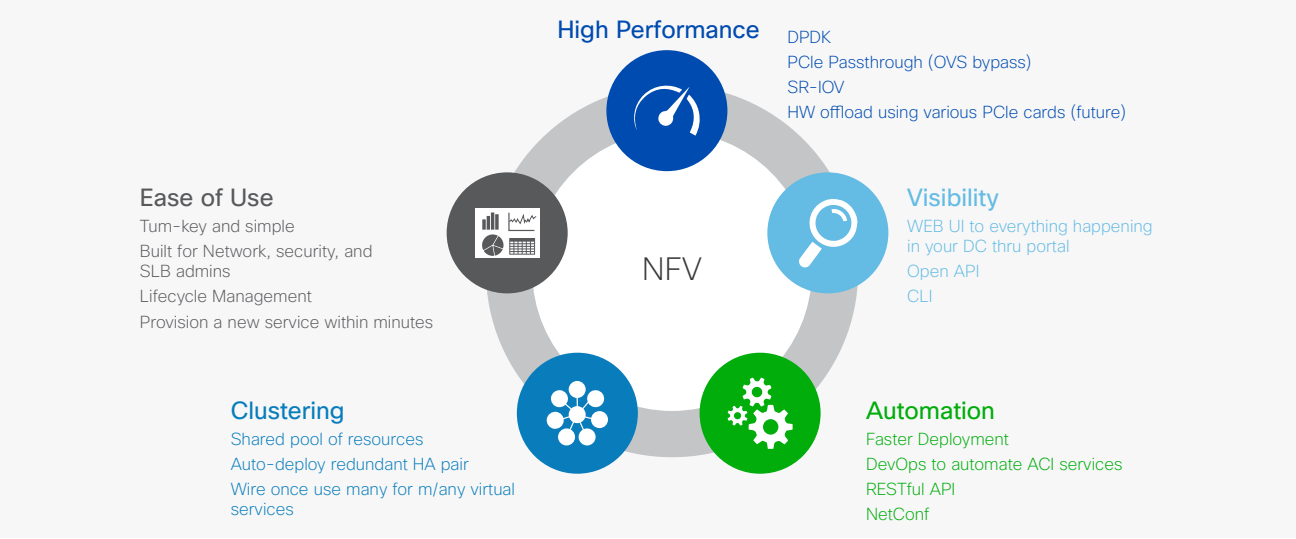

NFVi can quickly and easily deploy network virtual services with Cisco Secure Agile Exchange (SAE), built on the Cisco Cloud Services Platform (CSP) 2100, and Cisco Network Services Orchestrator (NSO). Teams can bring up these services at the pace that the DevOps and server teams need, usually in less than one hour—and even in minutes. This ready-to-use, open x86 Linux Kernel-based Virtual Machine (KVM) software and hardware platform is well suited for data center network functions virtualization (NFV). See Figure 2.

Figure 2. Cisco NFVi with Cisco SAE, CSP 2100, and NSO

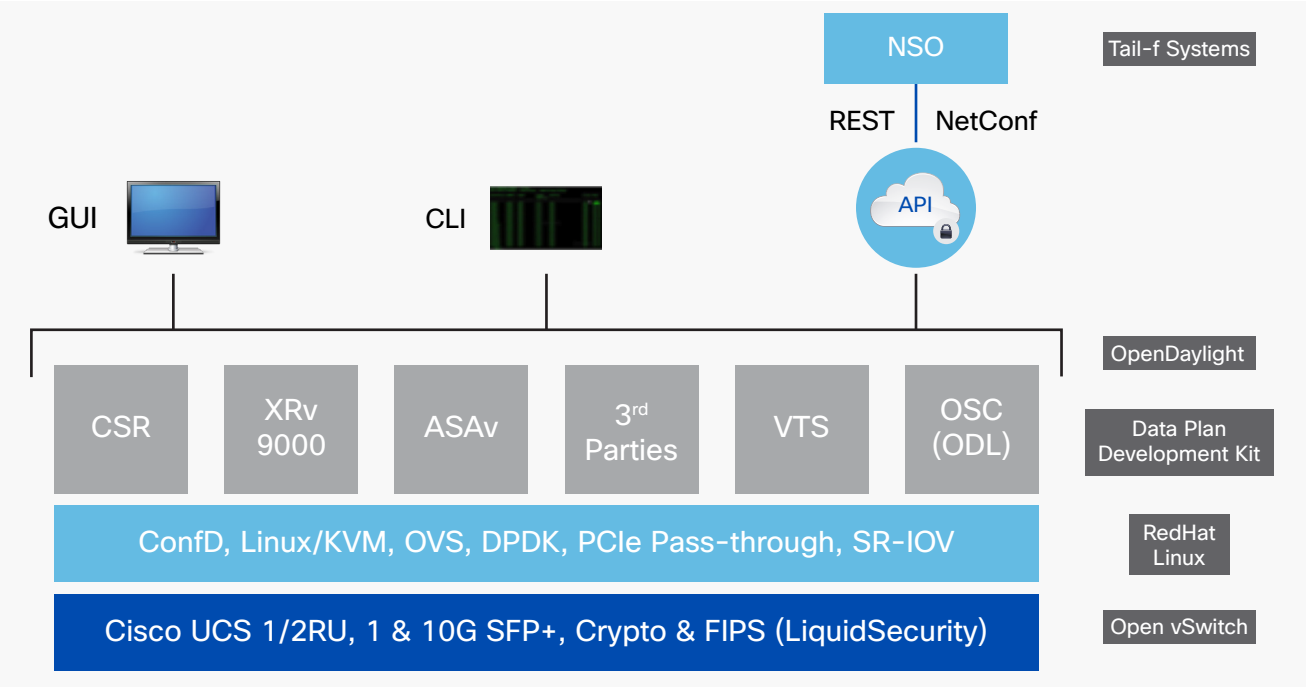

SAE is built on a CSP 2100 cluster. It is a purpose-built NFVi platform with open architecture (Red Hat Enterprise Linux [RHEL], KVM, Open vSwitch [OVS], and ConfD) and a management plane similar to that of network appliances. It is wired once and used many times for any network services (or VNFs). It supports Cisco and third-party VNFs, hardware offload (SSL, cryptography, and hardware security module [HSM]), and single-root I/O virtualization (SR-IOV) for high performance.

NSO provides end-to-end network orchestration through a service model uniformly for both physical and virtual network functions.

Virtualization and orchestration provide many benefits for the lifecycle of network services, but they also add complexity to the operation of the environment as new elements are introduced, such as:

- Mixed technologies
- Workflows and orchestration configurations
- Linux KVM
- Software elements and features such as OVS and SR-IOV
- Integration points (APIs)
- Infrastructure capacity management to accommodate all network services

In addition to the technology challenges, new processes and skill sets are required to manage and operate these highly virtualized services:

- If base elements break or are misconfigured, the impact on the network will be greater.
- New skill sets are needed to troubleshoot incidents and implement configurations.
- New configuration elements and items need to be taken into account.
- New processes are needed to plan and implement changes.

All these factors require new expertise and processes and time and investment to hire and train staff and adopt the new technologies and learn to operate them. Many times they require from customers time they don't have and money they don't want to spend on operations.

Cisco CMS and Holistic Management of Cisco NFVi

CMS provides a complete portfolio of day-2 operate services that increases the overall availability of IT, allowing customers to focus on promoting innovation and meeting business demands.

CMS discovers all components and subcomponents of an NFVi solution through the use of the Cisco Unified Computing System™ (Cisco UCS®), Linux, KVM, CSP 2100, and NSO management interfaces such as the XML API, Representational State Transfer (REST) API, and Simple Network Management Protocol (SNMP). Monitoring includes clusters, nodes, and VNFs. It can also establish dynamic component mapping and relationships (DCM-R) between relevant solution elements. After components and subcomponents have been discovered, the system extracts configuration, capacity, and performance metrics for monitoring purposes, threshold-crossing alerts, events, and reporting as shown in Figure 3.

Figure 3. Cisco CMS Holistic Monitoring of Cisco CSP 2100 as a Solution

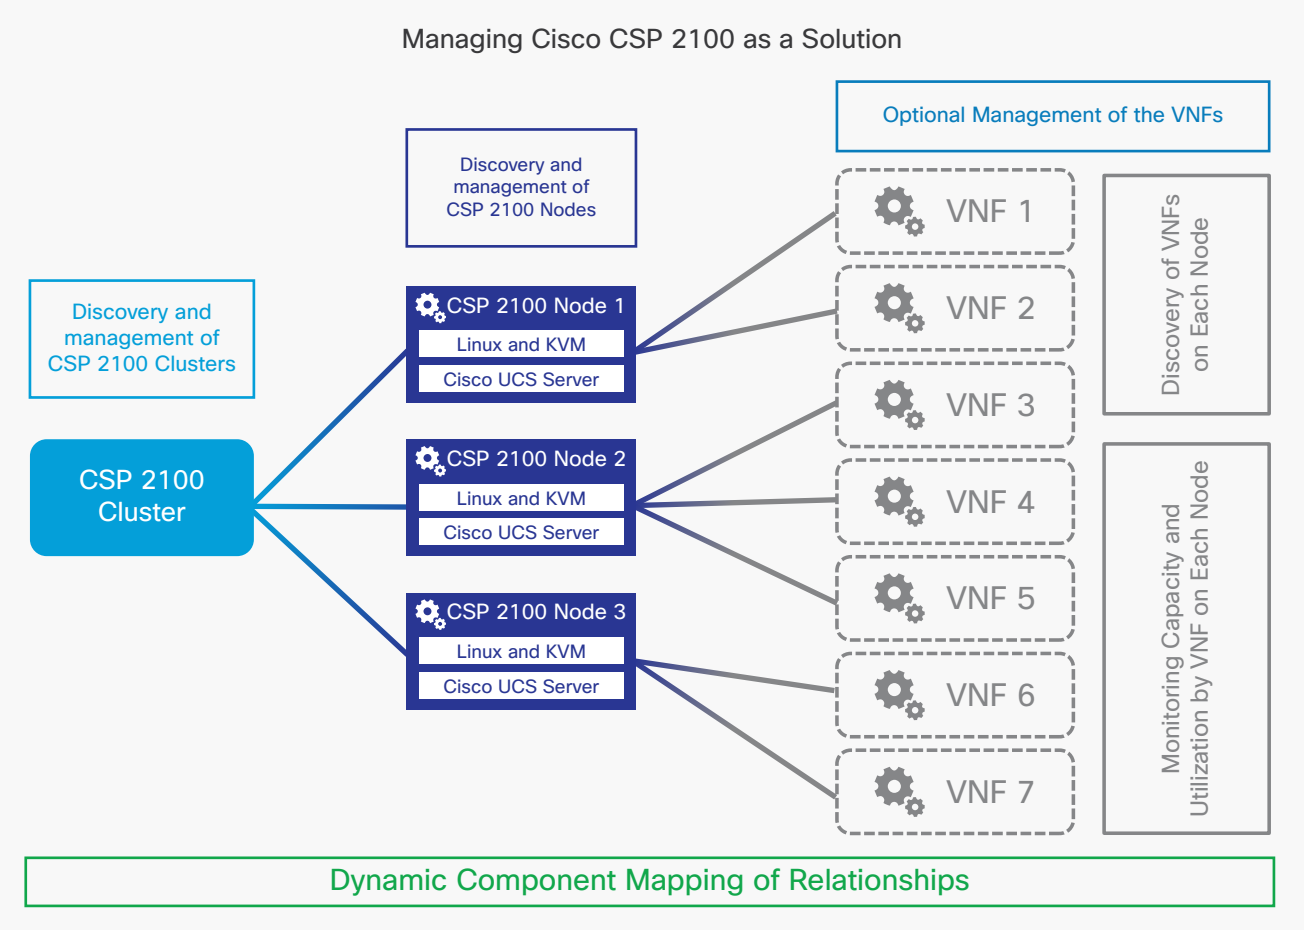

In addition to maintaining the health of the solution components, CMS tracks the allocation of resources such as CPU cores, memory, and disks from the perspectives of the clusters and nodes. CMS applies analytics to estimate the number of VNFs that can be accommodated based on the current allocation and capacity and trends over a three-month period, and it delivers this information in the form of a report published in the customer services portal.

Monitoring and Planning Cisco CSP 2100 Capacity

CMS uses the CSP 2100 REST API to collect resource capacity and allocation information at the cluster, node, and VNF levels. These resources include the number of CPU cores, memory capacity, and disk space available for allocation to VNFs at the cluster and node levels and the amounts allocated to each VNF out of the cluster and node capacities.

With this information, CMS uses analytics to provide trending and planning reports. CMS calculates the average size of VNFs running on the clusters and nodes and the deviation for the largest and smallest VNFs. It then estimates the number of additional VNFs that can be added based on the remaining resources.

In the sample case shown in Figure 4, the CPU and disk resource allocations are under control and could, respectively, accommodate another seven and eight VNFs. However, memory allocation has greatly increased in the previous month and is now close to 85 percent. Note that even though the 32 GB of available memory might seem reasonable, the calculated average size of a single VNF is 11.4 GB, which means that there is enough memory only for another two VNFs with the calculated average size to be added to the cluster. It's time to get another CSP 2100 appliance and add capacity to this cluster.

Figure 4. Cisco CSP 2100 Capacity Trending and Planning Report

Cisco CSP 2100 Resources Capacity Planning and Trending

Resource	VNFs in Cluster	Resources Allocated	Average Allocation to VNFs	Resources Available	Number of VNFs That Can Be Created
CPU	14	30 cores	2.1 cores	15 cores	7
Memory	14	160 GB	11.4 GB	32 GB	2
Disk	14	2600 GB	185.7 GB	1496 GB	8

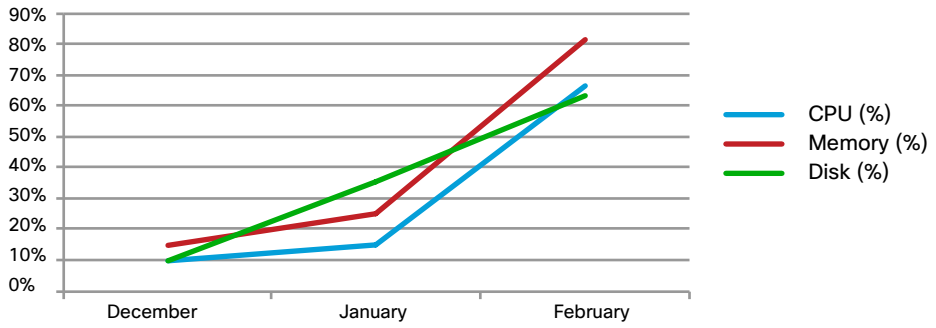

CMS also provides reports that allow a deeper look at cluster and node capacity and use. For the sample case described here, focusing on the memory constraint identified by the report shown in Figure 4, this additional reporting level provides guidance about how to allocate additional services in cluster nodes that deviate from the average in a balanced way (Figure 5).

Figure 5. Cisco CSP 2100 Cluster and Node Capacity Report (Memory Segment)

Cisco CSP 2100 Cluster and Node Capacity

Cluster	Total Memory Available for VNF Provisioning (GB)	Avgerage Memory Allocated (GB)	Peak Memory Allocated (GB)	Last Occurrence of Peak Memory Use
CSP_Cluster	32 GB	11.4 GB	16 GB	2017-01-16 23:00:00

Node	Total Memory Available for VNF Provisioning (GB)	Avgerage Memory Allocated (GB)	Peak Memory Allocated (GB)	Last Occurrence of Peak Memory Use
Node_01	6 GB	4 GB	16 GB	2017-01-16 23:00:00
Node_02	14 GB	3.2 GB	8 GB	2017-01-16 23:00:00
Node_03	12 GB	4.2 GB	8 GB	2017-01-16 23:00:00

Node_01 is the busiest node with only 6 GB remaining from the cluster total of 32 GB available.

Two additional VNFs with the average size of 11.4 GB can be accommodated in Node_02 and Node_03. The cluster has space for more than two additional VNFs if they are smaller than the average size.

The report provides the information needed to make the best use of the NFVi cluster resources.

CMS also provides a close look into the resource allocation and consumption of each VNF (Figure 6). Information such as resources allocated to each VNF, CPU load (percentage), memory use (percentage), and disk use (percentage) including average and peak values per VNF are presented in descending order of use by several sorting criteria (top VNFs by CPU, top VNFs by memory, and top VNFs by disk). This report allows users to compare the resource use with the amount of resources allocated and identify when a particular VNF cannot handle services by itself anymore and another instance should be provisioned.

As an example, the report depicted in Figure 6 presents the top-five VNFs by memory consumption. The two VNFs providing load-balancing service at Node_1 and Node_03 show very high memory use, and additional VNFs should be provisioned to handle the increasing traffic.

Figure 6. Cisco CSP 2100 Top VNF Utilization

Cisco CSP 2100 Top VNFs Utilization

VNF Name	Peak Memory	Allocated Memory	CSP 2100 Node	CSP 2100 Cluster
VNF_LB_A	92%	8 GB	**Node_01**	CSP_Cluster_01
VNF_LB_B	91%	8 GB	**Node_03**	CSP_Cluster_01
VNF_ASAv_01	76%	8 GB	Node_02	CSP_Cluster_01
VNF_ASAv_02	72%	8 GB	Node_03	CSP_Cluster_01
VNF_CSR1000v_A	65%	16 GB	Node_03	CSP_Cluster_01

With this approach, the goal of CMS is to mitigate the risk that business will be affected by resources that run out of capacity, by idle or unevenly balanced services, or by delays in critical deployments due to waiting times for the provisioning of additional capacity: for example, canceling a maintenance window that was already approved, planned, and scheduled to implement a new business application that depends on new critical network services.

Conclusion

CMS provides the information required for capacity management, allowing organizations to plan ahead and predict growth with trending for:

- Optimal allocation of resources and use of capacity
- Even balance of services loads
- Timely deployment of additional capacity

Cisco CMS helps customers expedite and simplify the adoption of Network Function Virtualization, and take full advantage of all the benefits of Cisco NFVi suites with CSP 2100 and NSO. They maximize their return on investment, and most importantly to effectively and successfully drive their business.

Figure 7. Cisco NFVi benefits in numbers

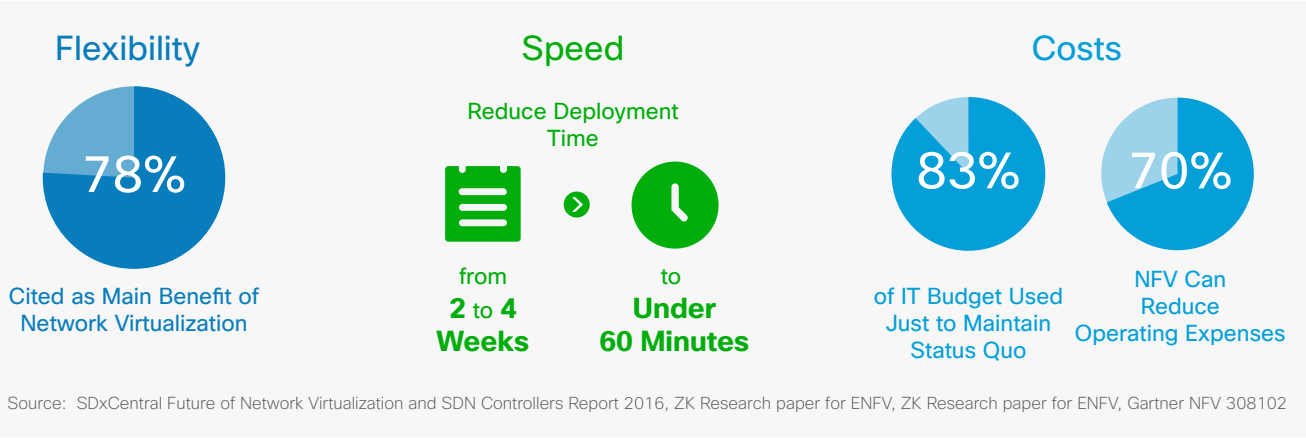

For More Information

Visit CMS on Cisco.com.

C11-738896-00 04/17