Cloud Provider Platform: Architectural Guidelines Powered by VMware Cloud Foundation

Table of Contents

List of Figures

List of Tables

1. Executive Summary

To accelerate the customer journey to the SDDC, VMware has introduced VMware Cloud Foundation™, a new and unified SDDC platform for the private and public cloud. Cloud Foundation brings together VMware compute, storage, and network virtualization into a natively integrated stack that can be deployed on premises or run as a service from the public cloud. The Cloud Foundation stack consists of VMware vSphere®, VMware NSX®, and VMware vSAN™ along with VMware SDDC Manager™. The SDDC Manager component fully automates and orchestrates the deployment, operation, and life-cycle management of these underlying SDDC components.To maximize the benefits of this rapidly deployed hyper-converged SDDC infrastructure, VMware partners with several hardware vendors to deliver a validated and prebuilt SDDC stack using Cloud Foundation.

VMware partners with more than 4000 Cloud Providers through VMware Cloud Provider Program. Multiple VCPP partners offer VCF based dedicated cloud service to enable their customers to enjoy the full benefit of hybrid cloud with maximum consistency between cloud and on-premise. VMware Cloud on AWS is another example of cloud service built on VCF.

VMware Cloud Director® is a product available exclusively for cloud service providers via the VMware Cloud Provider Program (VCPP) to address their public and hybrid cloud infrastructure-as-a-service (IaaS) use cases. Originally released in 2010, it enables service providers to orchestrate the provisioning of Software-Defined Data Center (SDDC) services as complete virtual data centers that are ready for consumption in a very short time frame. Cloud Director applies the principles of pooling, abstraction, and automation to all data center services such as storage, networking, and security. Using Cloud Director, service providers can deliver a shared SDDC infrastructure to multiple customers—that is, tenants—while keeping the resources of these tenants isolated from each other. Essentially, Cloud Director enables a complete multitenancy SDDC platform as one of its key architectural principles.

Designing and implementing a vCloud Director environment has typically been a time-consuming and complicated task. Although VMware Cloud Foundation itself does not eliminate the design and implementation of vCloud Director components, it dramatically reduces the level of effort needed to deploy, scale, and maintain the underlying SDDC platform stack. VMware Cloud Foundation can fully automate the deployment of vSphere, vCenter, NSX-V and NSX-T, and vSAN, all of which can serve as the foundational platform for vCloud Director. Additionally, VMware Cloud Foundation orchestrates patch and upgrade management of these underlying SDDC software components in a fully automated way.

Tests internally conducted at VMware have clearly demonstrated the technical feasibility of implementing a pairing of VMware Cloud Director with VMware Cloud Foundation. This technical paper describes architectural options and provides guidance on how to design such an environment. With its unique capabilities to provide automated Software-Defined Data Center (SDDC) hardware and software, stack bring-up services, and automated lifecycle management, Cloud Foundation is an excellent choice for service providers who are using Cloud Director and want a way to quickly and efficiently spin up SDDC resources ready for consumption with Cloud Director. Cloud Foundation deploys these resources with a self-contained architecture that is verified and validated by VMware. Customers and service providers can avoid going through iterative, time-consuming design and architecture cycles to implement an SDDC stack consisting of VMware vSphere, VMware vSAN, and VMware NSX. Cloud Foundation “consumption-ready” units of SDDC deploy resources in an automated, repeatable, and quality-assured manner. Service providers can face the requirement of hosting VMware Cloud Director managed resources onsite at their customers’ data centers for data security and compliance reasons. Cloud Foundation enables them to achieve efficient integration into these data centers because it is a fully self-contained configuration that provides clear and easy interfaces with an existing data center, greatly simplifying deployment and reducing overall complexity.

This technical paper is aimed at architects and designers who want to assess the benefits of deploying VMware Cloud Director with Cloud Foundation and who lead the design and implementation of such an architecture on behalf of a service provider. It therefore offers a short overview of both products and their key terminologies as they pertain to this discussion, provides basic architectural guidance on how to deploy Cloud Director with VMware Cloud Foundation, and discusses the benefits of this approach.

2. Introduction to VMware Cloud Foundation

2.1Overview of VMware Cloud Foundation

VMware Cloud Foundation delivers a natively integrated SDDC stack consisting of vSphere, NSX, and vSAN as core technologies for compute, storage, and network virtualization. Additional software components including the VMware vRealize® Suite and the VMware Horizon® suite can be deployed through SDDC Manager automated workflows on top as another option. Cloud Foundation supports any vSAN Ready Node configuration for those customers who want to utilize vSAN, and any server supported in accordance with the VMware Compatibility Guide for those customers who want to choose NFS or VMFS on FC as opposed to vSAN. When deploying on-premises private clouds, customers can choose either to build their own hardware configuration ready for Cloud Foundation or to order a completely prebuilt and preconfigured rack configuration from one of several hardware vendors who are certified partners for delivering prebuilt Cloud Foundation configurations.

Cloud Foundation offers the following core features and benefits:

- Automated hardware and software bring-up – The Cloud Foundation Builder Appliance deploys a complete SDDC platform consisting of software and hardware in a fully automated way that requires only minimal user input and advance planning. The Cloud Foundation Builder Appliance orchestrates the deployment and configuration of an SDDC platform that adheres to the blueprint set out by VMware Validated Designs and typically takes less than 3-hours to deploy.
- Simplified resource provisioning – Cloud Foundation creates and maintains logical pools of compute, storage, and network resources from the underlying physical rack hardware resources. For this purpose, these physical resources can be allocated and decommissioned dynamically “on the fly.”
- Automated lifecycle management – SDDC Manager orchestrates patch and upgrade management of the underlying SDDC software components in a fully automated way.
- Scalability and performance – Cloud Foundation delivers a private cloud instance, which can easily be integrated into an existing network.
- Multi-Instance Management – Multiple Cloud Foundation instances can be managed together by grouping them into a federation, such that each member can view information about the entire federation and the individual instances within it. Federation members can view inventory across the Cloud Foundation instances in the federation as well as the available and used aggregate capacity (CPU, memory, and storage). This allows you to maintain control over the different sites and ensure that they are operating with the right degree of freedom and meeting compliance regulations for your industry.

2.2 Architecture

Cloud Foundation supports two architectures – Consolidated and Standard.

2.2.1 Consolidated Architecture

The consolidated architecture design targets smaller Cloud Foundation deployments and special use cases. In this design, the management and user workload domains run together on a shared management domain. The environment is managed from a single vCenter Server and vSphere resource pools provide isolation between management and user workloads. In a consolidated architecture model, care must be taken to ensure that resource pools are properly configured as the domain is shared by the management and compute workloads. The consolidated architecture does not support NSX-T or the automated deployment of Horizon and Enterprise PKS.

As you add additional hosts to a Cloud Foundation system deployed on a consolidated architecture, you can convert to the standard architecture by creating a VI workload domain and moving the user workload domain VMs from the compute resource pool to the newly created VI workload domain. After moving these VMs, you may need to update shares and reservations on the compute resource pool in the management domain.

Figure 1. Cloud Foundation Consolidated Architecture

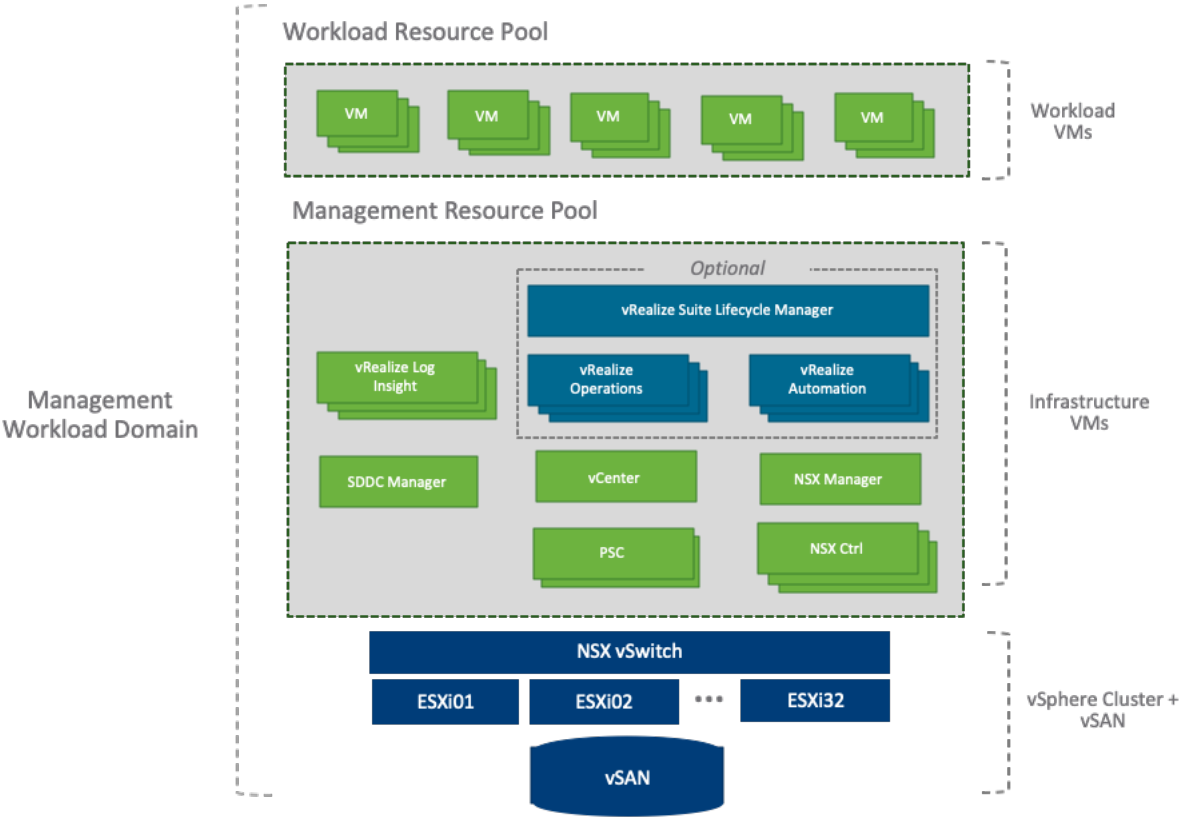

2.2.2 Standard Architecture

With the standard architecture model, management workloads run on a dedicated management domain and user workloads are deployed in separate virtual infrastructure (VI) workload domains. Each workload domain is managed by a separate vCenter Server instance which provides for scalability and allows for autonomous licensing and lifecycle management.

Standard architecture is the preferred model.

Figure 2. Cloud Foundation Standard Architecture

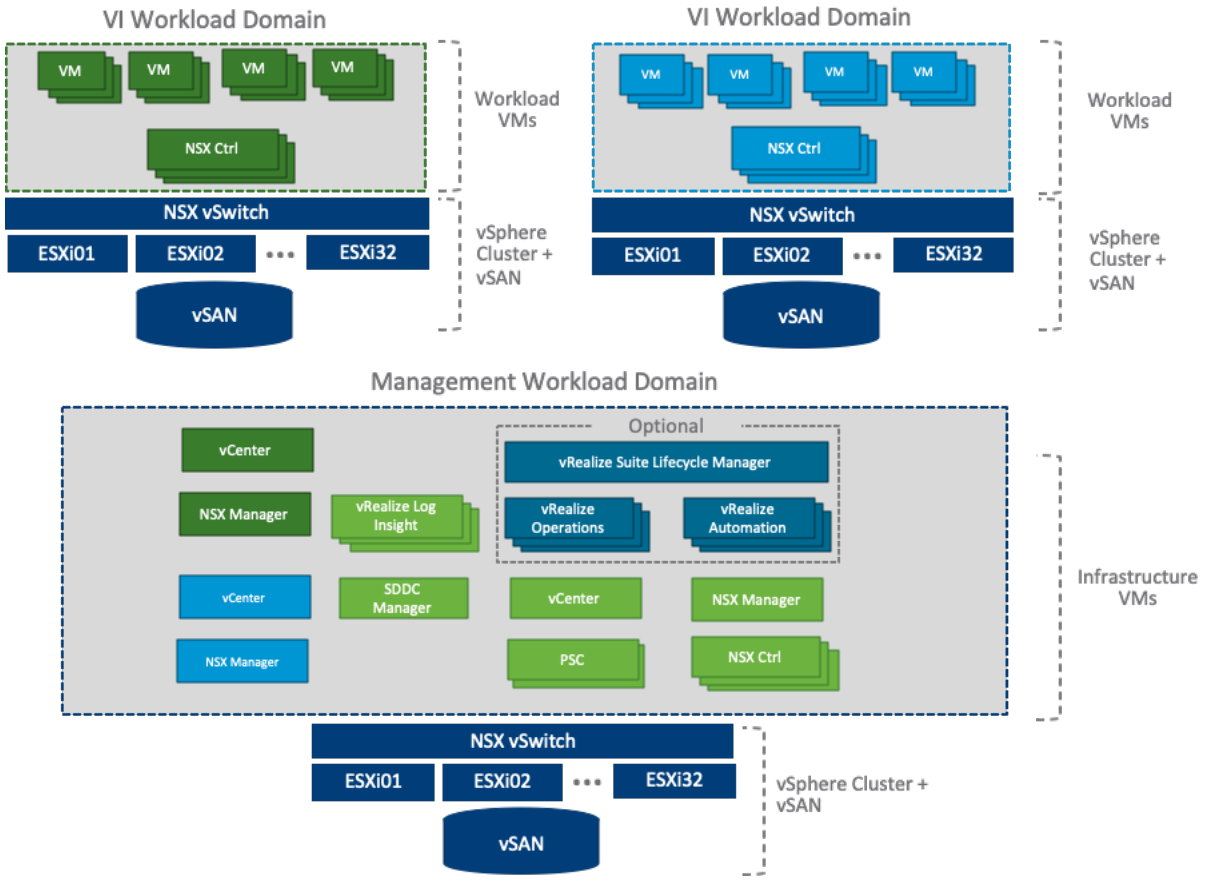

2.3 Imaging and Bring-Up

The automated installation and configuration of Cloud Foundation starts with deploying the Cloud Foundation Builder VM. The Cloud Foundation Builder VM includes the VMware Imaging Appliance, which you use to image your servers with ESXi software. After imaging your servers, you download and complete the deployment parameters sheet from the Cloud Foundation Builder VM to define your network information, host details, and other required information. You upload the completed spreadsheet back to the Cloud Foundation Builder VM where the provided information is validated, and the automated phase of the deployment process begins.

Note You can use the VMware Imaging Appliance (VIA) included with the Cloud Foundation Builder VM to image servers for use in the management domain and VI workload domains.

During bring-up, the management domain is created on the ESXi hosts specified in the deployment configuration spreadsheet. The Cloud Foundation software components are automatically deployed, configured, and licensed using the information provided.

During the bring-up process, the following tasks are completed:

- PSC, vCenter Server, vSAN, vRealize Log Insight, and NSX components are deployed.
- The management domain is created, which contains the management components - SDDC Manager, all vCenter Servers, and NSX Managers and Controllers.

Please note that NSX-T components are not deployed during the initial bring-up process. For the first NSX-T VI workload domain in your environment, the workflow deploys a 3 NSX Manager/Controller cluster in the management domain. The workflow also configures an anti-affinity rule between these VMs to prevent them from being on the same host for High Availability. All subsequent NSX-T workload domains share these NSX-T Managers.

The details for executing these installation steps are described in the official documentation for Cloud Foundation and in demo videos on the Cloud Foundation YouTube channel.

2.4 Cloud Foundation Life Cycle Automation

VMware Cloud Foundation automated Lifecycle Management delivers simple management of your environment with built-in automation of day 0 to day 2 operations of the software platform.

- Automated deployment: Automates the bring-up process of the entire software platform, including deployment of infrastructure VMs, creation of the management cluster, configuration of storage, cluster creation, and provisioning
- Infrastructure cluster provisioning: Enables on-demand provisioning of isolated infrastructure clusters to enable workload separation
- Simplified patching and upgrades: Enables a simplified patching/upgrading process of the software platform (including VMware vCenter Server®)

Figure 3. Cloud Foundation Life Cycle Automation

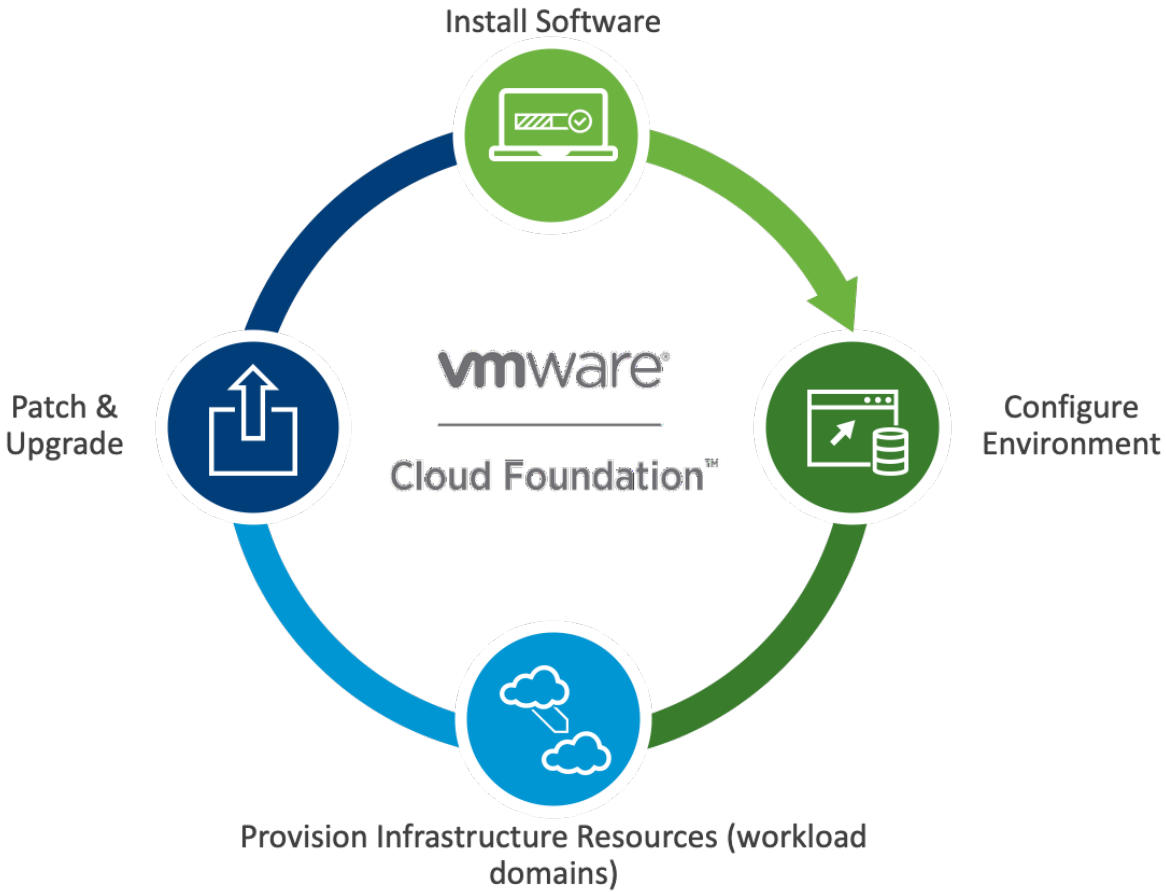

VMware Cloud Foundation provides automated lifecycle management on a per-workload domain basis. This way, available updates for all underlying components are validated for interoperability to consistently determine proper installation order and maintain compliance with best practices and compatibility matrices. The updates can also be scheduled for automatic installation on a per-workload domain basis to maximize flexibility without impacting system availability. This allows the infrastructure admin to target specific workloads or environments (development vs. test vs. production) to execute updates independently and maximize productivity.

2.4.1 VMware Cloud Director and Life Cycle Automation

VMware Cloud Director components (as specified in Appendix 8.3) are currently not part of Cloud Foundation life cycle automation. These components have to be manually installed after initial bring-up and have to be evaluated and updated manually accordingly before and/or after performing cloud foundation updates in accordance with the VMware Product Interoperability Matrix.
End-to-end Life cycle automation for Cloud Director is a roadmap item.

2.5 Domains – Logical Pooling of Physical Resources

Cloud Foundation introduces the concept of a *domain* to provision intelligent units of SDDC resources on the bare-metal Cloud Foundation hardware infrastructure described in section 2.2. At the simplest level, a domain describes an aggregation of a specified number of hosts commissioned to Cloud Foundation.

There are two types of domains within Cloud Foundation:

- Management domain – Only one management domain will exist per VCF instance. The management domain is primarily dedicated to hosting system components related to Cloud Foundation and further management components related to vSphere, NSX, and vRealize Suite.
- Workload domain – A workload domain in the Cloud Foundation context is dedicated to hosting end-user or customer workloads.

Note In the Consolidated Architecture model, you also run user workloads in the management domain. A workload domain only contains user workload resources.

Figure 4 shows an example of carving out several domains from within the Cloud Foundation rack configuration.

Figure 4. Cloud Foundation Domains

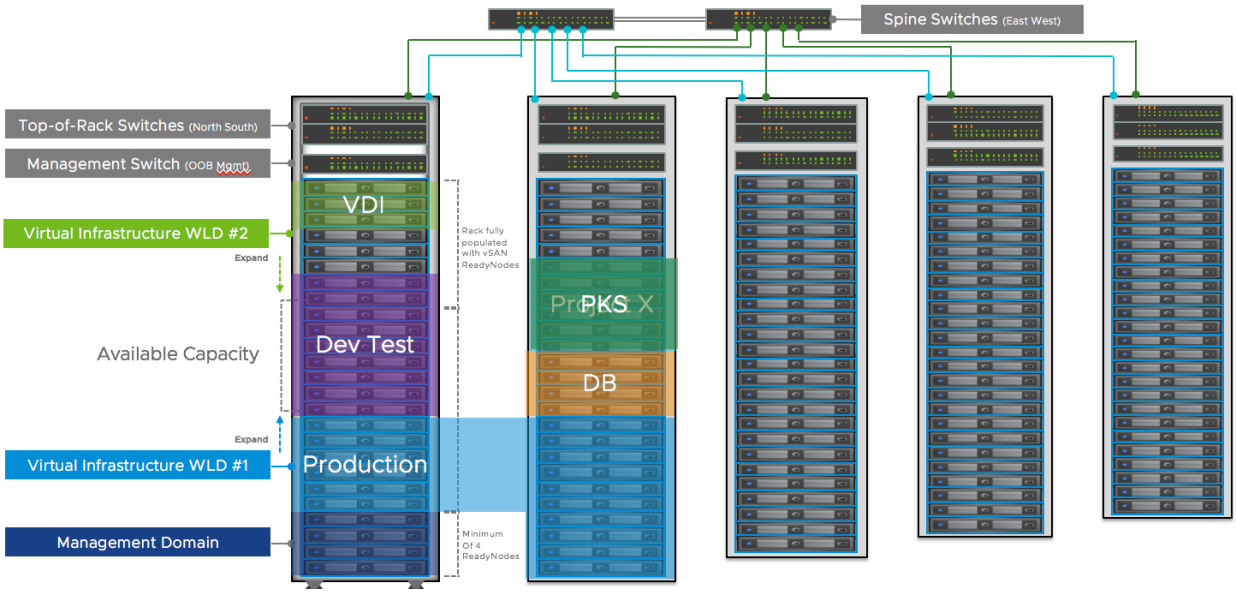

The concept of a domain has the following further characteristics or rules within Cloud Foundation:

- At its core, a domain is an aggregation of hosts with ESXi preinstalled and all necessary network infrastructure configured. This includes management network, VMware vSphere vMotion® network, and so on.

- The preinstalled ESXi hosts of the domain are configured production ready by the Cloud Foundation domain creation workflow. They also form a new vSphere cluster with a vSAN datastore in case of management domain and of those workload domains you chose VSAN as their datastore.

When creating a domain, a new vCenter Server Appliance instance is deployed. The vCenter Server Appliance is placed in the Cloud Foundation management domain. The vCenter Server Appliance instance is registered with one of the two Platform Services Controller instances within the management domain.

- A domain can be expanded to include multiple vSphere clusters.

When creating the management domain a NSX-V instance is deployed. The NSX-V Manager instance and the NSX-V Controllers of the NSX-V installation are placed in the Cloud Foundation management domain.

When creating a new NSX-V workload domain, a new NSX-V instance is deployed. The NSX-V Manager instance of the new NSX-V installation is placed in the Cloud Foundation management domain. The NSX-V Controller cluster will be deployed in the newly created workload domain.

- When deploying the first NSX-T workload domain, the NSX-T Manager instances of the new NSX-T installation are placed in the Cloud Foundation management domain. NSX Edges are needed to enable overlay VI networks and public networks for north-south traffic. NSX Edges are not deployed automatically for an NSX-T VI workload domain. You can deploy them manually after the VI workload domain is created. Subsequent NSX-T VI workload domains share the NSX-T Edges deployed for the first workload domain.

A domain which uses VSAN as its storage must have a minimum of 3 and a maximum of 64 hosts. These limits are basically defined by the configuration limits of vSphere 6.7, used within the current Cloud Foundation release, and the corresponding vSAN version.

2.5.1 Management Domain

A management domain within Cloud Foundation is primarily dedicated to hosting Cloud Foundation management components such as SDDC Manager as well as those related to VMware virtual infrastructure (VI) such as Platform Services Controller instances, VMware vCenter Server® instances, NSX-V Manager instances, NSX-T Manager instances, NSX-T Controller instances, and management components of vRealize Suite, all of which are part of the Cloud Foundation framework.

A management domain must use VSAN as its storage. On the contrary, a workload domain can use NFS or VMFS on FC as its storage in addition to VSAN. Due to this requirements, all the nodes of a management domain must be VSAN ready nodes.

Placing end-user or customer workloads into a Cloud Foundation management domain is allowable and sensible in certain scenarios:

- Relatively Small Cloud Foundation configurations that use the consolidated architecture model for which there are no resources to create a separate workload domain
- Management components of user workloads

However, consumers of a Cloud Foundation configuration must ensure proper sizing of the management domain to account for additional resource requirements imposed by their workloads. The standard and minimal size of a workload domain is four hosts; this can be expanded after initial deployment.

Figure 5 provides an overview of the Cloud Foundation management components within a Cloud Foundation management domain.

Figure 5. Cloud Foundation Management Domain

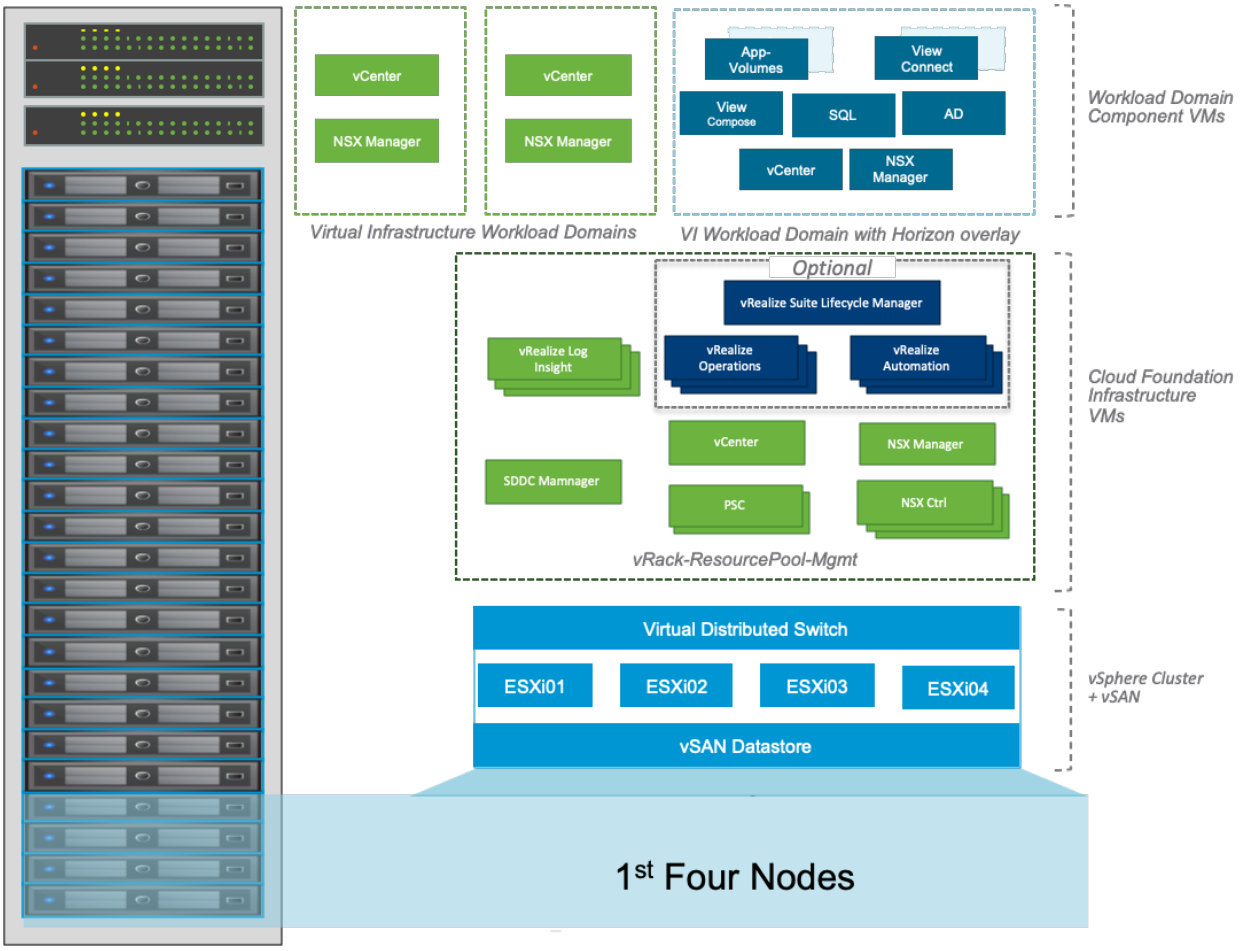

As depicted in Figure 5, the management domain consists of a vSphere cluster, initially with four hosts—that is, the first four hosts of the first rack. It hosts management VMs, which can be subdivided into the following areas:

1) Core Cloud Foundation management layer – The core Cloud Foundation management layer contains the minimal set of VMs available with every Cloud Foundation installation. The VMs in this layer in turn can be classified in the following way:

   a) Cloud Foundation system stack – This stack contains the basic system VMs hosting the Cloud Foundation configuration. Its components will be explained only briefly because they are not very relevant to a discussion related to Cloud Director architecture:

      - SDDC Manager –SDDC Manager is the central gateway to Cloud Foundation configuration management. SDDC Manager automates the entire system lifecycle (from initial bring-up, to configuration and provisioning, to upgrades and patching), and simplifies day-to-day management and operations.

   b) VI management stack – The VI management stack consists of standard management components related to vSphere and NSX:

      - Platform Services Controller – Platform Services Controller instances provide shared platform services—VMware vCenter® Single Sign-On, for example—to VMware components such as vCenter Server and SDDC Manager. A detailed description of all Platform Services Controller services is beyond the scope of this document. There are two Platform Services Controller instances configured in Enhanced Linked Mode that form a vSphere domain. All vCenter Server instances, which are subsequently

deployed when creating Cloud Foundation workload domains, are registered with one of the two Platform Services Controller instances.

- vCenter Server – This vCenter Server instance solely manages the vSphere cluster of the Cloud Foundation management domain.
- NSX-V Manager plus NSX-V Controller cluster (three VMs) – These VMs form the management or control plane of the NSX installation, which provides NSX networking services for the vSphere cluster in the management domain.
- NSX-T Cluster (three NSX Manager VMs) – These VMs form the management or control plane of the NSX installation, which provides NSX networking services for the vSphere cluster both in the management domain and the workload domains.

c) Operations management stack – Cloud Foundation is integrated with the vRealize Suite of products. The operations management stack consists of the following VMware vRealize Suite of products:

- vRealize Suite Lifecycle Manager – vRealize Suite Lifecycle Manager delivers complete lifecycle and content management capabilities for the VMware vRealize Suite. vRealize Suite Lifecycle Manager supports the deployment, upgrade, and patching of vRealize Log Insight, vRealize Automation, and vRealize Operations Manager.
- vRealize Log Insight – Log Insight delivers heterogeneous and highly scalable log management with intuitive and actionable dashboards, sophisticated analytics, and broad third-party extensibility. It provides deep operational visibility and faster troubleshooting across physical, virtual and cloud environments. Log Insight is installed by default for the management domain. You can add licenses to enable Log Insight for VI workload domains.
- vRealize Operations Manager – vRealize Operations Manager delivers intelligent operations management with application-to-storage visibility across physical, virtual, and cloud infrastructures. Using policy-based automation, operations teams automate key processes and improve IT efficiency. This is an optional component.
- vRealize Automation – vRealize Automation is a cloud automation tool that accelerates the delivery of IT services through automation and pre-defined policies, providing high level of agility and flexibility for developers, while enabling IT teams to maintain frictionless governance and control. This is an optional component.

Management layer for Cloud Foundation workload domains – With each instantiated Cloud Foundation workload domain, a set of certain management VMs are deployed. Because there are different types of Cloud Foundation workload domains, which will be explained in more detail in a following section, there are also different types of management layers for these workload domains:

d) Management layer for VI workload domains – The management layer for a VI workload domain consists of only two VMs, which are deployed with each VI workload domain:

- vCenter Server – This vCenter Server instance manages the vSphere cluster of the corresponding Cloud Foundation workload domain.
- NSX Manager – This NSX Manager instance connects to the vCenter Server instance of the Cloud Foundation workload domain to manage the NSX networking services for that workload domain.

e) Management layer for virtual desktop infrastructure (VDI) workload domains – VDI workload domains are not relevant to a discussion related to Cloud Director and are not described in detail here. The management layer also contains a vCenter Server instance and an NSX Manager instance, as is the case with the management layer of VI workload domains, but it additionally contains VMware Horizon management components.

f) Management layer for VMware Enterprise PKS – VMware Cloud Foundation enables automated deployment of VMware Enterprise PKS on an NSX-T workload domain. Enterprise PKS is a container services solution that simplifies the deployment and management of Kubernetes clusters. Enterprise PKS manages container deployment from the application layer all the way to the infrastructure layer according to the requirements for production-grade software. Enterprise PKS supports high availability, autoscaling,

health-checks and self-repairing of underlying virtual machines, and rolling upgrades for the Kubernetes clusters.

2.5.2 Workload Domains

Workload domains are dedicated to hosting consumer or customer workloads within the Cloud Foundation configuration. They are created via the SDDC Manager Web UI or SDDC Manager API by an administrator of the Cloud Foundation configuration.

As has already been mentioned, there are currently different types of workload domains available in the Cloud Foundation product. Although we speak about different type of workload domains, all workload domain types are instantiated from a default VI Workload domain:

a) VI workload domains – VI workload domains are reserved for VI workloads and will be discussed later in this document when the Cloud Director IaaS use case is described in more detail.

b) VDI infrastructure workload domains – VDI workload domains are created for hosting workloads related to VMware Horizon.

c) VMware Enterprise PKS workload domains – VMware Enterprise PKS workload domains are created for hosting container workloads on Kubernetes clusters.

Figure 6 highlights the components of a VI workload domain with four hosts, which are deployed within Cloud Foundation:

Figure 6. VI Workload Domain

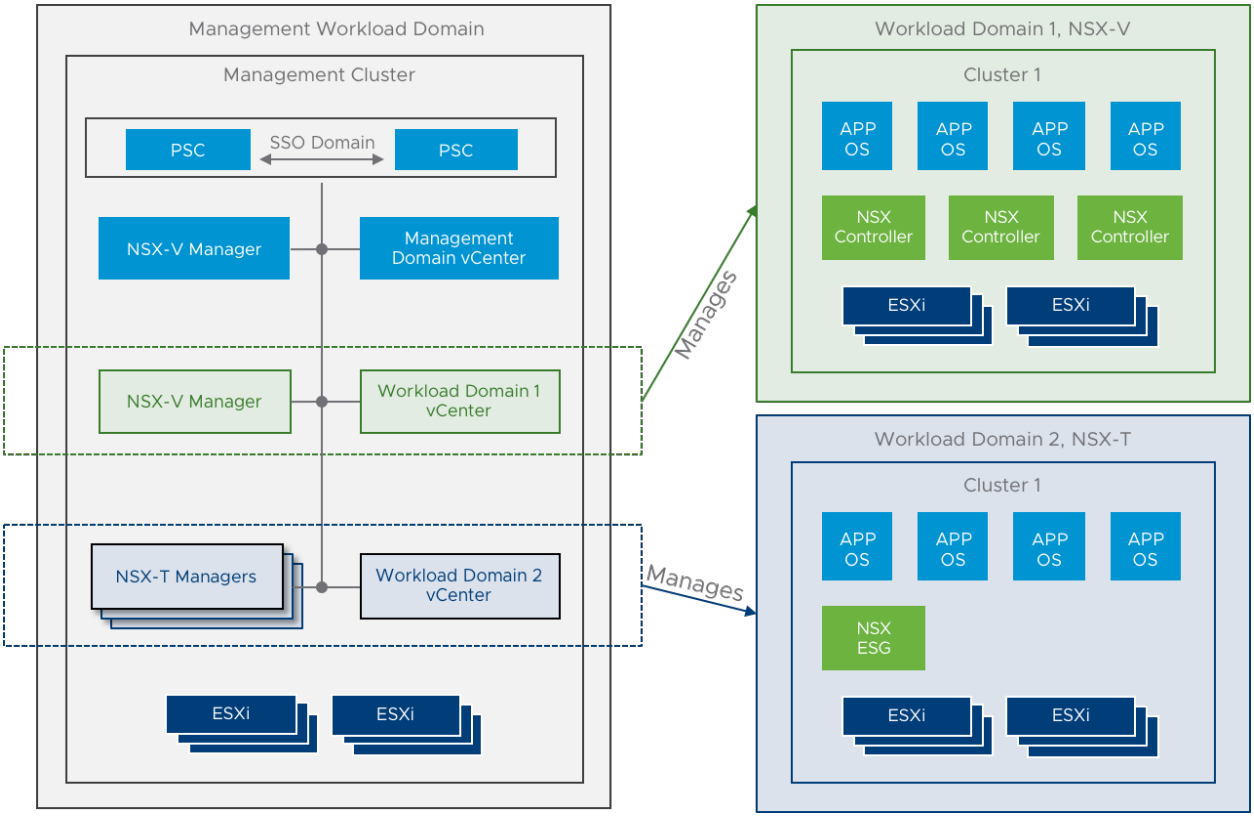

A workload domain can use either NSX-V or NSX-T as its networking stack.

A workload domain can use either VSAN, NFS or VMFS on FC as its storage. When a workload domain uses NFS or VMFS on FC, compatible hosts are not restricted to VSAN Ready nodes. Any vSphere compatible nodes in accordance with the VMware Compatibility Guide can be used.

Many workflow tasks are managed behind the scenes when deploying a VI workload domain. The following is a summary of the tasks related to this discussion:

- Components deployed in the Cloud Foundation management domain – Although a VI workload domain deployment is initiated, some related components for managing the newly created VI workload domain are deployed in the Cloud Foundation management domain:
  - A new NSX-V Manager instance and a vCenter Server Appliance instance are deployed into the management domain and are configured automatically.
  - The vCenter Server Appliance instance manages the VI workload domain. The NSX-V Manager instance is registered or connected with that corresponding vCenter Server instance.
- VI workload domain deployment – The VI workload domain is created by deploying and configuring the following components:
  - ESXi hosts (4)
  - VMware vSphere Distributed Switch™ (1)
  - vSphere cluster created with the four ESXi hosts
  - vSAN enabled within the vSphere cluster, creating a vSAN datastore if VSAN is selected as the storage type.
  - NSX-V Controller cluster corresponding to the newly deployed NSX installation or NSX Manager and consisting of three NSX Controller VMs

    Note The NSX Controller cluster, although not situated in the NSX data path, is deployed in the VI workload domain and **not** in the management domain. This is due to the close communication relationship between the managed ESXi hosts in the VI workload domain and the NSX Controller VMs.

As of VCF 3.0, you can now stretch a cluster in the management domain or in a VI workload domain across two availability zones for higher availability. You can perform a planned maintenance on an availability zone without any downtime and then migrate the applications after the maintenance is completed. You can prevent service outages before an impending disaster such as a hurricane or rising flood levels.

Figure 7. Dual Availability Zone Stretched VCF

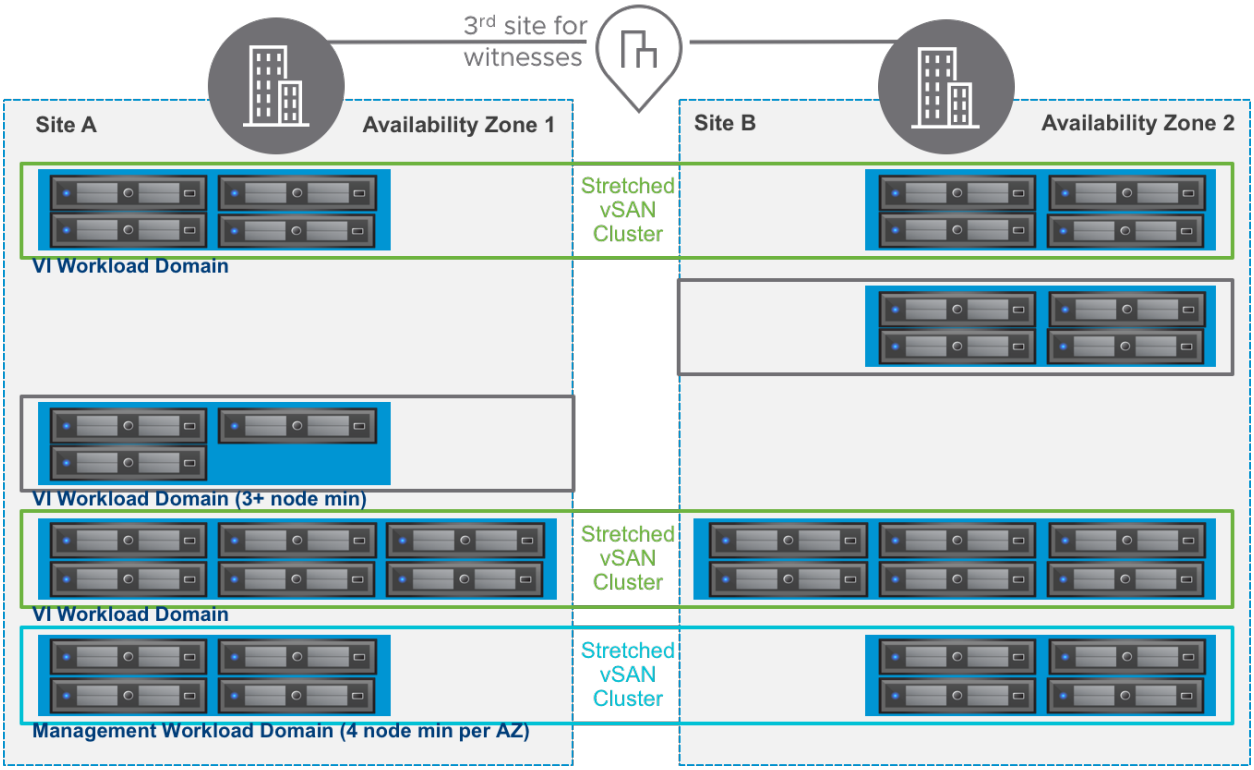

To summarize, a VI workload domain consists of an empty vSphere cluster prepared to provide virtualized compute, storage, and network resources. It is deployed and configured completely automatically.

3. VMware Cloud Director

3.1Overview of VMware Cloud Director

VMware Cloud Director is the central Cloud Management and Orchestration platform within the VMware Cloud Provider Platform, enabling Service Providers to deliver public and hybrid cloud data center services. From a provider perspective, it enables service providers to carve virtually isolated, multitenancy-enabled resources from a physical pool of compute, storage, and network resources. These virtual entities are also called *virtual data centers*. Within the construct of virtual data centers, Cloud Director currently is the only VMware SDDC automation product that implements secure multitenancy by isolating data center services at all virtual layers of data center resources.

The following are some of the most important additional features of Cloud Director:

- VMware vSphere vApp™ catalog – Tenants can create multi-VM configurations called vSphere vApps and store them as vApp templates in their own vApp catalog. This catalog is isolated from all catalogs of other tenants hosted on the same physical platform. The vApp templates are then used for repetitive deployment of vApps.
- Automated networking and security services based on NSX – Software-defined NSX networking and security services such as distributed switching based on VXLAN; distributed firewalling; and VMware NSX Edge™ services such as routing, NAT, VPN, and load balancing are consumed by Cloud Director—during vApp deployments, for example—in a fully automated way.
- Self-service Web portal – End users of the tenants using the Cloud Director platform have direct access to their vApp catalogs and deployed vApps through a modern, self-service Web portal.
- Cloud API – Tenants and providers can consume Cloud Director services via a comprehensive REST-based API. For example, providers can develop their own Web portal for consumption by their tenants.
- Multisite capabilities - Multisite capabilities allow an organization user to log in to the Cloud Director UI hosted at any of the sites where they have an Organization and Organization VDC. Upon login, the UI displays a sites icon that allows them to switch to other sites in which they have resources so that they can manage them from the same session.
- Central Point of Management - Cloud Director Central Point of Management (CPoM) allows tenants to access their dedicated private vCenter Server(s) through the Cloud Director UI. Service Providers can now deliver vCenters with a branded Cloud Portal experience, without the need to set up VPNs. VCD acts as proxy to API calls to the vCenters.

3.2 Architecture

3.2.1 Core Terminology

Cloud Director introduces a layer of abstraction to pool compute, storage, and network resources while enabling complete multitenancy. Table 1 summarizes the most important terms pertaining to this paper.

Table 1. Architectural Constructs of Cloud Director

Construct	Description
Tenant	The portion of the infrastructure that is used by, and provides services to, the customer.

Organization	An organization is the unit of multitenancy that represents a single logical security boundary. An organization contains users, virtual data centers, and networks.
Resource Group	A resource group is a set of compute, storage, and network resources dedicated to tenant workloads and managed by a single pairing of vCenter Server instance plus NSX Manager instance. Cloud Director manages the resources of all attached resource groups through API communication with vCenter Server and NSX Manager.
Provider Virtual Data Center	A provider virtual data center is a grouping of compute, storage, and network resources from a single vCenter Server instance. It consists of one or more resource pools with one primary—that is, initial—resource pool along with datastores connected to this resource pool. All resource pools must be from a single vCenter Server instance or resource group, so a single provider virtual data center cannot be spanned across vCenter Server instances. Provider virtual data centers typically are configured with root resource pools from a vSphere cluster. Therefore, cluster objects typically are assigned as resource pools for backing a provider virtual data center. Multiple organizations can share provider virtual data center resources.
Organization Virtual Data Center	An organization virtual data center is a subgrouping of compute and storage resources allocated from a provider virtual data center and assigned to a single organization. On the vSphere layer, a resource pool under each of the corresponding provider virtual data center resource pools is created and represents the organization virtual data center.
vSphere vApp	A vSphere vApp is a container for a distributed software solution and is the standard unit of deployment in Cloud Director. A Cloud Director vApp is different from a vSphere vApp in the manner it is instantiated and consumed in Cloud Director. It enables power operations to be defined and specifically ordered. It consists of one or more VMs and can be imported or exported as an Open Virtualization Format (OVF) package. A Cloud Director vApp can have additional constructs specific to Cloud Director such as vApp networks.
Site	A geographic location within which a Cloud Director instance runs. Usually a single physical location, but can also be spread across multiple, connected locations.

3.2.2 Management Cluster

In a similar manner to that previously described for Cloud Foundation management components, components related to Cloud Director are placed into a management cluster. Figure 5 shows a typical example management cluster for a basic Cloud Director implementation.

Figure 8. Cloud Director Management Cluster

Note Figure 8 depicts only the main, core Cloud Director management components required for establishing a basic Cloud Director deployment. Depending on the use case, more components and extensions might be deployed in the Cloud Director management cluster. However, these additional components are not relevant for a basic architectural discussion evaluating the feasibility of implementing Cloud Director with a Cloud Foundation installation and are therefore not shown here.

Note The Cloud Director DB as shown in Figure 8 is installed as an external database for a Cloud Director server group installed on Linux only. A Cloud Director server group that consists of appliance deployments uses the embedded database in the first member of the server group. You can configure a Cloud Director database high availability by deploying two instances of the appliance as standby cells in the same server group.

The Cloud Director management cluster, as depicted in Figure 8, consists of the following core components:

Cloud Management Components:

- Cloud Director cells – Cloud Director cells are the central management components of a Cloud Director environment. They provide the following core functionalities:
  - Enable access to Cloud Director managed cloud resources—virtualized compute, storage, and network resources—to tenants. Two access points are built into the Cloud Director cells to access these resources:
    - UI/API access point – The Cloud Director UI is implemented as a Web console and provides the standard portal for setting up and managing an IaaS solution based on Cloud Director. The Cloud API supports developers who want to build their own customized interactive clients apart from the standard UI to better satisfy additional UI requirements a service provider might have for its individual use cases.
    - Remote console proxy – The remote console of deployed VMs or vApps is proxied through the remote console proxy service of Cloud Director cells. Because there is a central point of access, this alleviates the remote console access in service provider environments from a security perspective.
  - Manage cloud resources based on vSphere and NSX; implement a logical abstraction layer by leveraging logical pooling of these resources; enable secure

multitenancy by fully isolating virtual data center resources between multiple organizations or tenants

- Cloud Director database – Although the Cloud Director cells are stateless from an application perspective, the Cloud Director database is the central repository for storing the configuration data of a Cloud Director installation. Depending on the The Cloud Director installation type, the Cloud Director database is embedded when the Cloud Director cells are deployed as appliances or must be deployed externally from the Cloud Director cells if the Cloud Director cells are installed on Linux. Only Postgres 10 databases are supported for use with Cloud Director 10.0 in accordance with the VMware Solution Interoperability Matrix.
- NFS transfer share – The NFS transfer share is a mandatory component of every Cloud Director installation. It provides temporary storage for concurrent OVF and ISO file transfers between resource groups. In Figure 5, it is depicted as a virtual appliance, but it can also be provided by an external NFS storage array.
- vCenter Server – vCenter Server as a cloud management component manages only the vSphere environment in which the cloud management components are hosted. This role is very similar to the vCenter Server role described in section 2.5.1 for the Cloud Foundation management domain.
- Load Balancer – A load balancer is required to provide a single point of access in a multicell Cloud Director installation. Cloud Director cells have no extraordinary requirements regarding the load balancer product. If an NSX installation is available—directly within the management cluster, for example—an NSX Edge device can easily be configured as a load balancer for Cloud Director cells. Otherwise, a third-party load balancer can be used.

Resource Group Management Components:

The management components of a Cloud Director resource group—one vCenter Server instance and one NSX Manager instance per resource group—are very similar to those already described for the management components of Cloud Foundation VI workload domains.

3.2.3 Resource Abstraction Layers

Figure 9 provides an example structure of the resource abstraction layers of Cloud Director.

Figure 9. Cloud Director Resource Abstraction Concepts

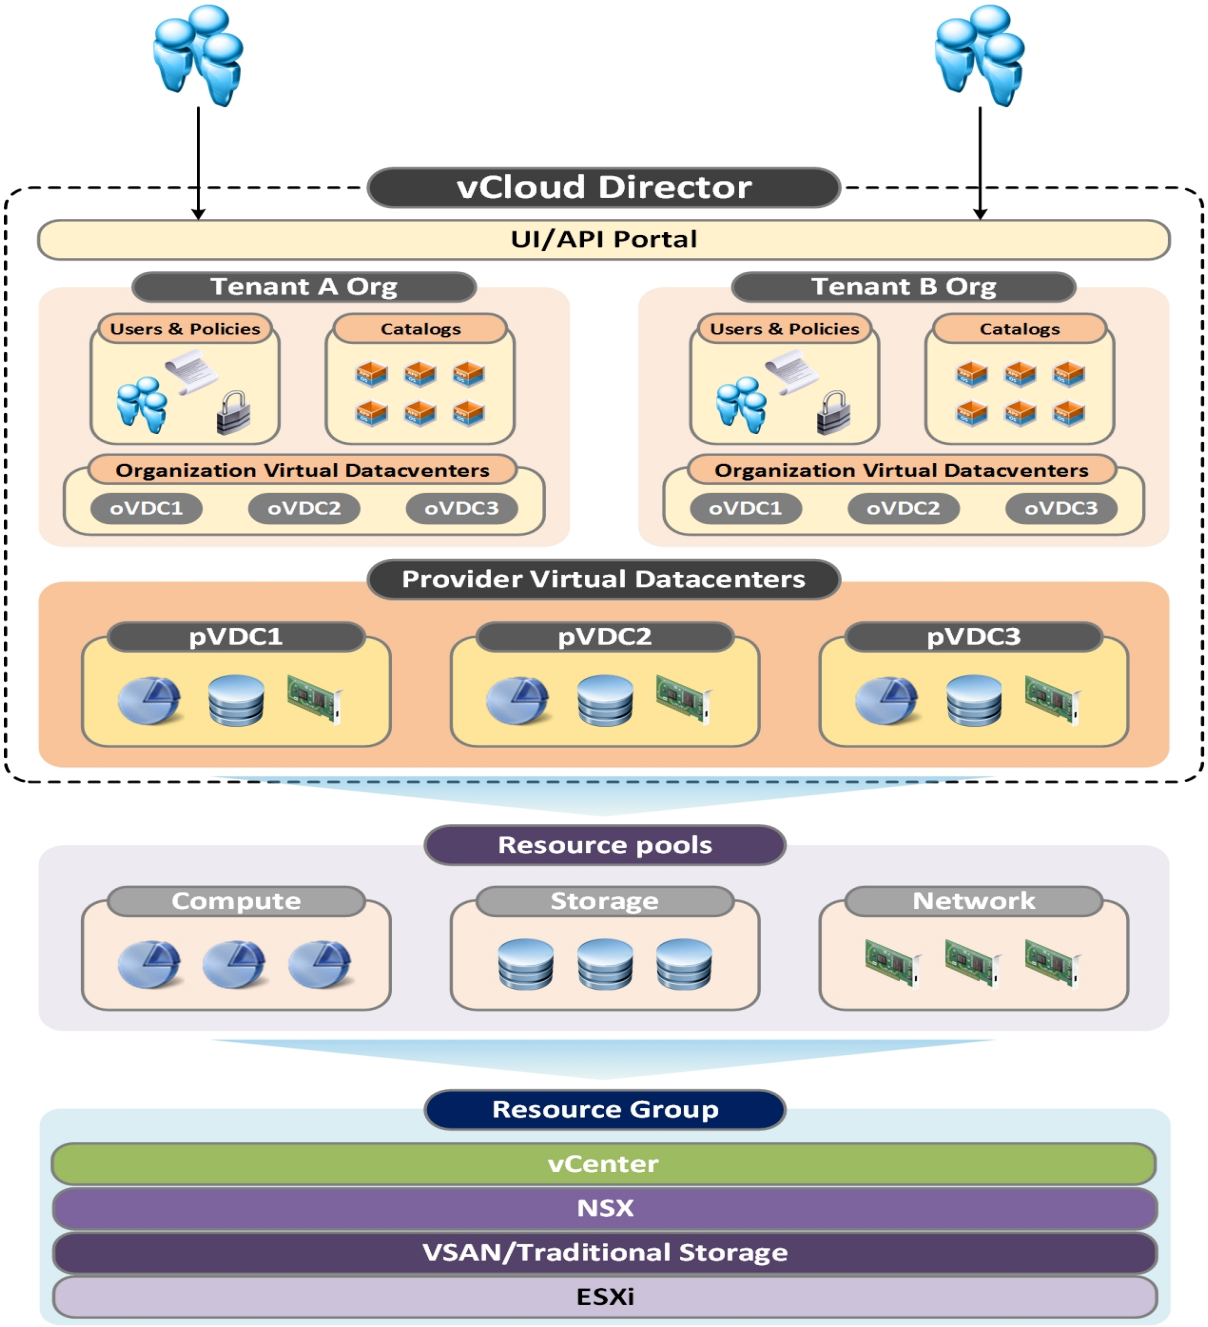

- The base layers compose a resource group consisting of a vCenter Server and NSX Manager pair along with managed ESXi hosts.

  Note Although only one resource group is shown in Figure 9, a Cloud Director instance can manage up to 25 resource groups or vCenter Server instances. However, a resource group always marks a distinct border for the elasticity of a provider virtual data center.

  Note NSX-T deployment has a different maximums. One Cloud Director instance can support maximum of 11 NSX-T Managers and 200 ESXi hosts backed by NSX-T.

- The resource group presents resource pools—typically vSphere clusters or child resource pools—to the Cloud Director installation.
- Within a Cloud Director installation, one or more provider virtual data centers, pVDCs, are configured. As part of the configuration, each provider virtual data center can be assigned one or more resource pools or vSphere clusters. Each pVDC can have up to 64 resource pools.
- Within the next resource layer, an organization is configured for tenant access. In Cloud Director architectural discussions, the terms “tenant” and “organization“ are used interchangeably. The Cloud Director UI, however, is not familiar with the term “tenant” and solely uses “organization.”
- Within an organization, one or more organization virtual data centers are configured. Each one carves out resources from exactly one provider virtual data center. Multiple organization virtual data centers can carve out resources from a single provider virtual data center. But in contrast, a single organization virtual data center cannot carve out resources from multiple provider virtual data centers.
- Organization users access the resources of organization virtual data centers and can use catalog items to deploy vApps in these data centers.

3.2.4 Guidelines to Tenant Resource Capacity Clusters

Number of vCenter Server instances

- The vCenter Server sizing guide suggests using a medium profile of vCenter Server to support 4,000 VMs.
- Number of vCenter Server instances = number of VMs/4000 = 15,000/4000 = 4 (rounded)

Number of ESXi hosts

- ESXi host count is determined based on the number of powered-on VMs, using the formula below. The formula provides a rough estimate and the actual number of hosts required depends on many parameters and the type of workloads.
- Number of hosts = (number of powered on VMs * Avg # of vCPUs per VM) / (sockets*cores*hyper threading*vCPU-to-pCPU ratio) = (9000*1)/(2*8*2*16) = 17

Number of Cloud Director cells

- As with the Cloud Director design guide, the number of Cloud Director cells for this setup is calculated using following formula:
- Number of Cloud Director cells = (Number of VMs/4000) + 1 = (15,000/4000) + 1 = 5 (rounded)

4. Cloud Director with Cloud Foundation

4.1Benefit for Cloud Providers

Cloud Providers know that time is money. The faster customers onboard, the faster the time to revenue for the cloud provider is. That is why the combination of VMware Cloud Foundation and Cloud Director creates competitive differentiation for cloud providers. Leveraging the advantage of HCI and the value of VMware SDCC like automated day 0 deployment and easier day 2 operation, in combination with industrial-scale tenancy and management layer specifically designed for cloud providers, the combination of VMware Cloud Foundation and Cloud Director facilitates rapid onboarding, easier day2 operations, simplified consumption and faster monetization of services.

4.2 Typical Use Cases for Cloud Providers

There are three typical use cases of Cloud Foundation for cloud providers:

- As a flexible and simpler infrastructure for multi-tenant service.

  A cloud provider build their infrastructure using VCF and deploy entire VCD stack on VCF based infrastructure. Both VCD management resources and tenant workload resources are deployed on one or multiple VCF based SDDCs.

- As a flexible and simpler infrastructure for cloud management resources.

  A cloud provider build their infrastructure for management cluster using VCF and deploy entire VCD management resources on this VCF based infrastructure. Tenant workload resources are deployed on other infrastructures. These Tenant workload infrastructures can also be VCF based or otherwise.

- As a dedicated infrastructure for dedicated private cloud service.

  A cloud provider build one set of infrastructure per tenant using VCF. This dedicated SDDC contains VI Workload Domains that are presented to a single organization as cloud resource or vSphere resource.

4.3 Cloud Foundation and Cloud Director Bill of Materials

The Cloud Foundation and Cloud Director software setup used in this document is comprised of the following software Bill-of-Materials (BOM).

Table 2. Cloud Foundation and Cloud Director Bill of Materials (BOM)

Software Component	Version	Date	Build Number
Cloud Director	10.1	09 APR 2020	15967253
Cloud Builder VM	2.2.1.0	14 JAN 2020	15345960
SDDC Manager	3.9.1	14 JAN 2020	15345960
VMware vCenter Server Appliance	6.7 Update3b	05 DEC 2019	15132721
VMware ESXi	6.7 Update3b	05 DEC 2019	15160138
VMware vSAN	6.7 Update3b	05 DEC 2019	14840357
VMware NSX Data Center for vSphere	6.4.6	10 OCT 2019	14819921
VMware NSX-T Data Center	2.5	19 SEP 2019	14663974

4.4 Terminology Mapping

To design a Cloud Director environment on top of a Cloud Foundation installation, it is helpful to understand the design terminology regarding Cloud Foundation and Cloud Director. There are only a few mappings necessary between Cloud Director and Cloud Foundation constructs because there are only a few overlapping abstraction layers between the two products. These mappings are described in the following subsections.

4.4.1 Cloud Director Management Cluster and Cloud Foundation Management Domain

The Cloud Director management cluster and the Cloud Foundation management domain share the purpose of hosting the following management components:

- Product-related system components and VMs—for example, Cloud Foundation management components and Cloud Director cell servers
- vCenter Server management components
- vCenter Server or NSX Manager instances of managed resources

  Note Managed resources include “resource groups” in the case of Cloud Director and “VI workload domains” in the case of Cloud Foundation.

- Optional enhanced management components such as a syslog server (vRealize Log Insight) or a system management server (vRealize Operations)

When deploying Cloud Director with VCF there are 3 design options regarding the placement of the Cloud Director management components:

1. Place Cloud Director management components within the Cloud Foundation management domain.
2. Place Cloud Director management components within a Cloud Foundation workload domain.
3. Place Cloud Director management components outside Cloud Foundation.

Design Option 1: Place Cloud Director management components within the Cloud Foundation management domain.

Pros

- Easy and straightforward – The Cloud Foundation management domain already exists in a Cloud Foundation installation.
- Minimal planning required – The Cloud Foundation management cluster already exists as a mandatory component of Cloud Foundation.
- Starting VCF 3.0, VSAN Stretched Cluster is available for the VCF Management Domain and the VCF Workload Domain. Both the VCF Management Domain and the VCF Workload Domain can reside in two datacenters. If you need disaster resiliency for VCD depoyment, you can use VSAN Stretched Cluster to make the cluster disaster resilient.
- Better resource efficiency – The Cloud Foundation Management domain starts from 4 or more hosts and it is usually bigger than the resource required for VCF management instances. Placing VCD Management Cluster will consume the surplus resource of the VCF Management Domain and brings better resource efficiency.

  Note The assumption is made here that only the main, core Cloud Director management components required for establishing a basic Cloud Director deployment are installed in a default 4 node VCF Management Domain. It is

recommended to perform proper capacity analysis on the VCF Management Domain if additional components and extensions might be required, which could lead to extending the VCF Management Domain beyond the default 4 nodes. Also see Appendix 8.3 – VMware Cloud Director Footprint.

Cons

- Limited architectural flexibility regarding Platform Services Controller high availability – Cloud Foundation currently does not offer a load balancer option for this. If a Cloud Director instance integrates into a vCenter Single Sign-On or Platform Services Controller instance, a dedicated registration to one of the two Cloud Foundation Platform Services Controller instances is required with no automated failover.
- More complex Cloud Foundation management domain upgrades – Cloud Foundation upgrades today do not take into account interoperability and compatibility with Cloud Director components. Before initiating a Cloud Foundation upgrade, providers must first verify compatibility with Cloud Director and are most likely required upgrading the Cloud Director components prior to initiating the Cloud Foundation management domain upgrade.

Figure 10. Cloud Director Components Placement Design Option 1

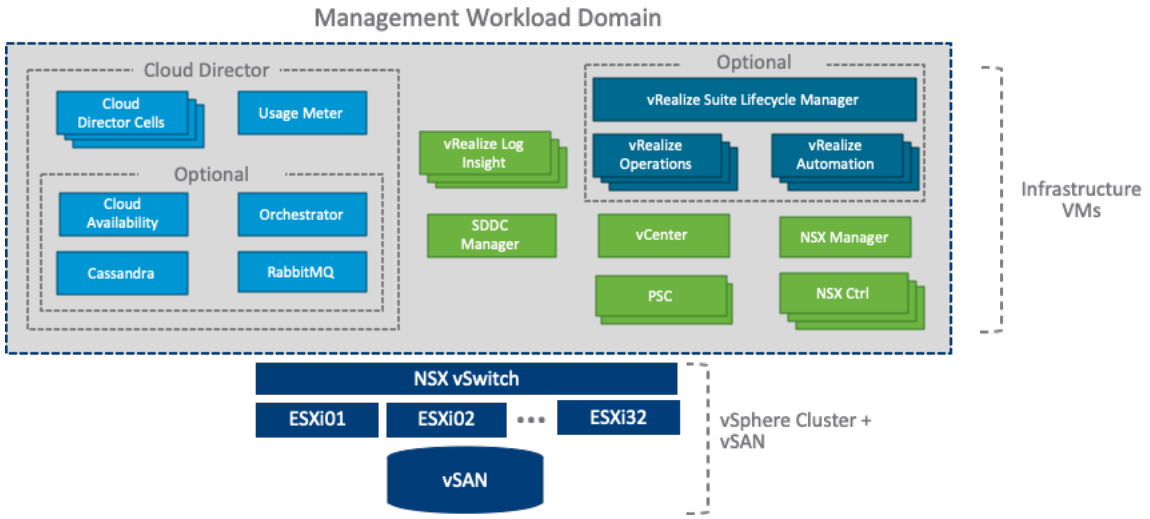

Design Option 2: Place Cloud Director management components within a Cloud Foundation workload domain.

Certain requirements or constraints in an environment might lead to the decision to separate the Cloud Director management components from the Cloud Foundation management domain.

Pros

- Leverage SDDC manager to provide automated deployment and LCM of the VCD Management Cluster.
- Direct adjacency to the underlying controlled and abstracted resources from Cloud Director.
- Starting VCF 3.0, VSAN Stretched Cluster is available for the VCF Management Domain and the VCF Workload Domain. Both the VCF Management Domain and the VCF Workload Domain can reside in two datacenters. If you need disaster resiliency for VCD depoyment, you can use VSAN Stretched Cluster to make the cluster disaster resilient.

- Deploy the VCD Management Cluster as a NSX-T based VI Workload Domain preventing downtime required for upgrading NSX-V to NSX-T in the future, especially when deploying the VCD management components on virtual networks.

Cons

- Limited architectural flexibility regarding Platform Services Controller high availability – Cloud Foundation currently does not offer a load balancer option for this. If a Cloud Director instance integrates into a vCenter Single Sign-On or Platform Services Controller instance, a dedicated registration to one of the two Cloud Foundation Platform Services Controller instances is required with no automated failover.
- More planning required – Greater effort must be made, especially when the externally placed Cloud Director management components must traverse several security zones before reaching the Cloud Foundation components to communicate with—that is, vCenter Server, NSX Manager, and so on.
- Added resource consumption – Placing Cloud Director components apart from the VCF Management components leads to additional resource consumption, which reduces ROI, especially when using a stretched cluster to provide disaster resiliency for Cloud Director components.

Figure 11. Cloud Director Components Placement Design Option 2

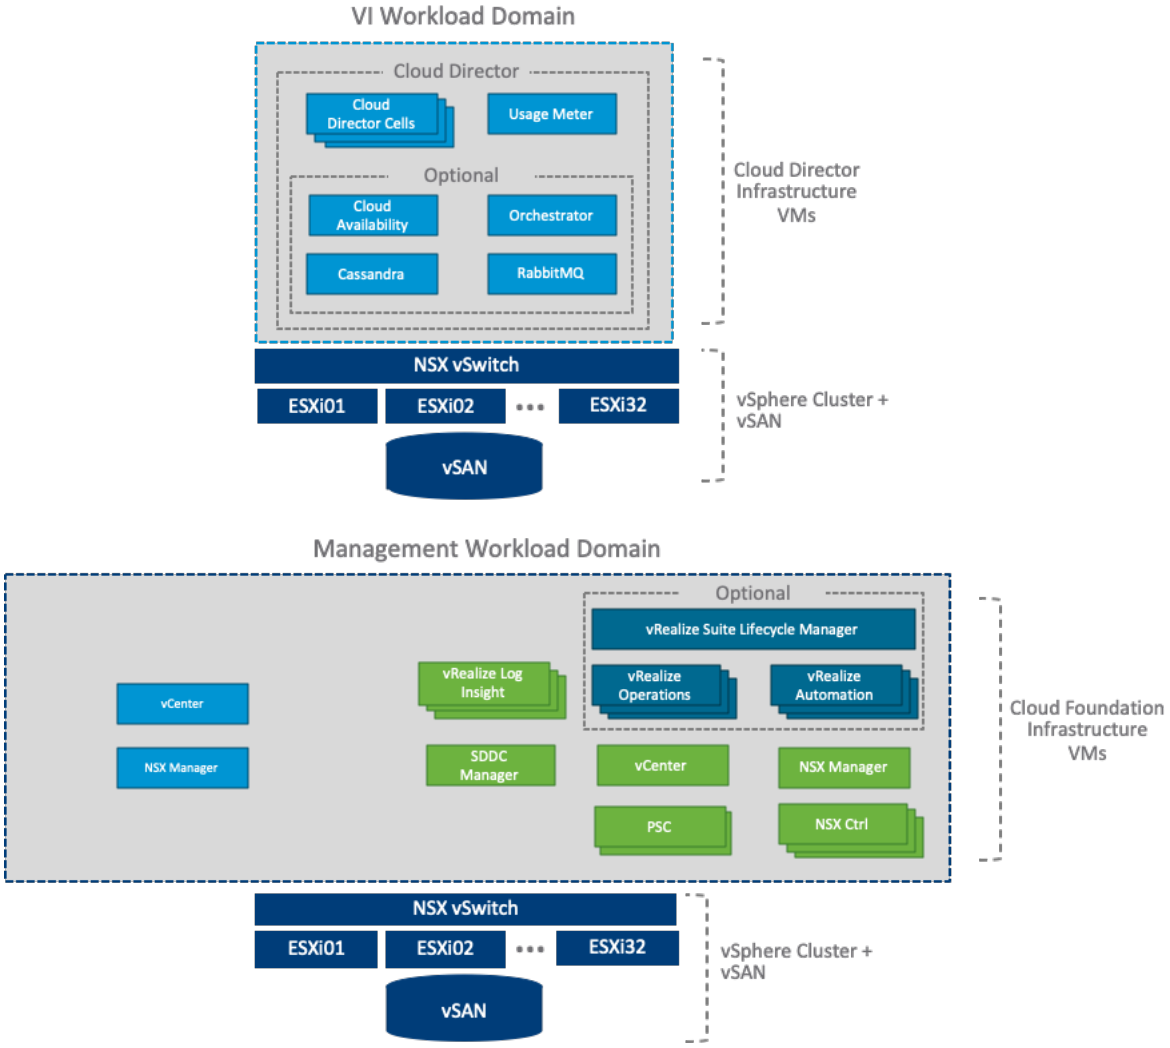

Design Option 3: Place Cloud Director management components outside of Cloud Foundation.

Certain requirements or constraints in an environment might lead to the decision to separate the Cloud Director management components from the Cloud Foundation management domain. This option could be the case in a brownfield scenario where Cloud Director has already been deployed and VCF infrastructure is used for tenant workload resources or dedicated private cloud service.

Pros

- Architectural flexibility – Cloud Director can be registered with a Platform Services Controller instance, which is configured for high availability with a load balancer. On a higher level of the architecture design, the option is available to protect the Cloud Director management components by a stretched cluster—that is, vSphere Metro Storage Cluster or vSAN stretched cluster. The registration of Cloud Director with a Platform Services Controller instance enables authentication of Cloud Director system users against an existing vSphere domain. Because a Cloud Director system must access a VMware vSphere Web Client instance to delegate authentication, at least one vCenter Server instance must be registered with the same Platform Services Controller instance.
- Existing management cluster – A management cluster may already exists in a customer's environment and can be used for the Cloud Director installation.
- Clear role distinction of Cloud Foundation – In this model, Cloud Foundation is a pure "resource provider" for the Cloud Director installation, along with other potential resource providers that serve different purposes and different SLAs.

Cons

- More planning required – Greater effort must be made, especially when the externally placed Cloud Director management components must traverse several security zones before reaching the Cloud Foundation components to communicate with—that is, vCenter Server, NSX Manager, and so on.
- Added resource consumption – Placing Cloud Director components apart from a VCF Management components leads to additional resource consumption, which reduces ROI, especially when using a stretched cluster to provide disaster resiliency for Cloud Director components.

Figure 12. Cloud Director Components Placement Design Option 3

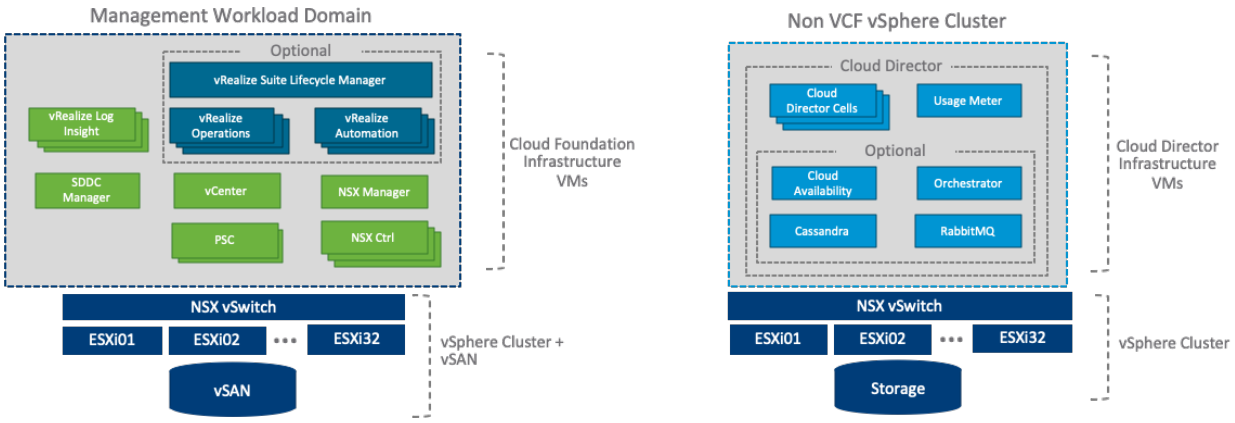

4.4.2 Cloud Director Resource Group and Cloud Foundation VI Workload Domain

The Cloud Director resource group and the Cloud Foundation VI workload domain have the common characteristic that both constructs describe a group of vSphere compute and storage resources as well as NSX networking resources managed by a single vCenter Server and NSX Manager pair.

This terminology mapping leads to the following design principle:

Design principle: Use a Cloud Foundation VI workload domain as a resource group within the Cloud Director installation.

This design principle must be followed when implementing Cloud Director with Cloud Foundation.

4.4.3 Cloud Director Central Point of Management and Cloud Foundation

Introduced in version 9.7, Cloud Director provides functionality that allows tenants to access dedicated private vCenters through Cloud Director. Service Providers can now deliver vCenters with a branded Cloud Portal experience, without the need to set up VPNs. VCD acts as proxy to API calls to the vCenters.

VCD as Central Point of Management offers the following Functionalities:

- Tenant Portal Access to Dedicated vCenters
- Inventory of vSphere estate
- Consolidated API End Point access to VMware estate for Cloud Providers
- Proxy vSphere API Access for Tenants

4.4.3.1. Use cases of Cloud Director Central Point of Management

As described in section 4.2, VCF instances can be used to provide multi-tenant cloud services or private cloud services.

In a multi-tenant service scenario, a VCF workload domain is usually shared between multiple tenants and therefore Cloud Director Central Point of Management cannot be used. However, when a tenant is assigned a dedicated VCF workload domain, that tenant has a dedicated vCenter Server and Cloud Director Central Point of Management can be used with this architecture to present this dedicated VCF workload domain as a vsphere resource.

Note All vCenter servers in a VCF instance share a single SSO domain and are placed in enhanced linked mode. Therefore a vCenter user account and proper RBAC configurations should be set up to prevent visibility into other vCenter servers and constructs that are shared between vCenter servers as part of enhanced linked mode.

In a private cloud service scenario, the entire VCF instance is dedicated to a single tenant and so is the SSO domain. Using Cloud Director Central Point of Management in this scenario makes the vCenter server RBAC configuration less complex and less of a concern. However it is still a good practice to implement a proper RBAC configuration to prevent tenant access to the VCF management domain vCenter server.

VMware recommends using Cloud Director Central Point of Management with dedicated VCF instances only. This is the only scenario in which there is true multi-tenancy across the whole stack.

4.4.4 Cloud Director Resource Consumption Examples on Cloud Foundation

4.4.4.1. Example 1 – Tenant Utilizing Multiple VCF Workload Domains

Figure 13 depicts a single customer utilizing multiple VCF Workload Domains. Each workload domain and subsequent resource pool and cluster is mapped to a provider virtual datacenter. While both have ESXi and vSAN for storage, they differ in NSX backing types with one being NSX-T and the other NSX-V. As the storage used is vSAN the solution administrator can further customize the storage offerings by utilizing different vSAN policies. This means the same storage backing can deliver different services to different customers even in the same cluster.

Figure 13. Tenant Utilizing Multiple VCF Workload Domains

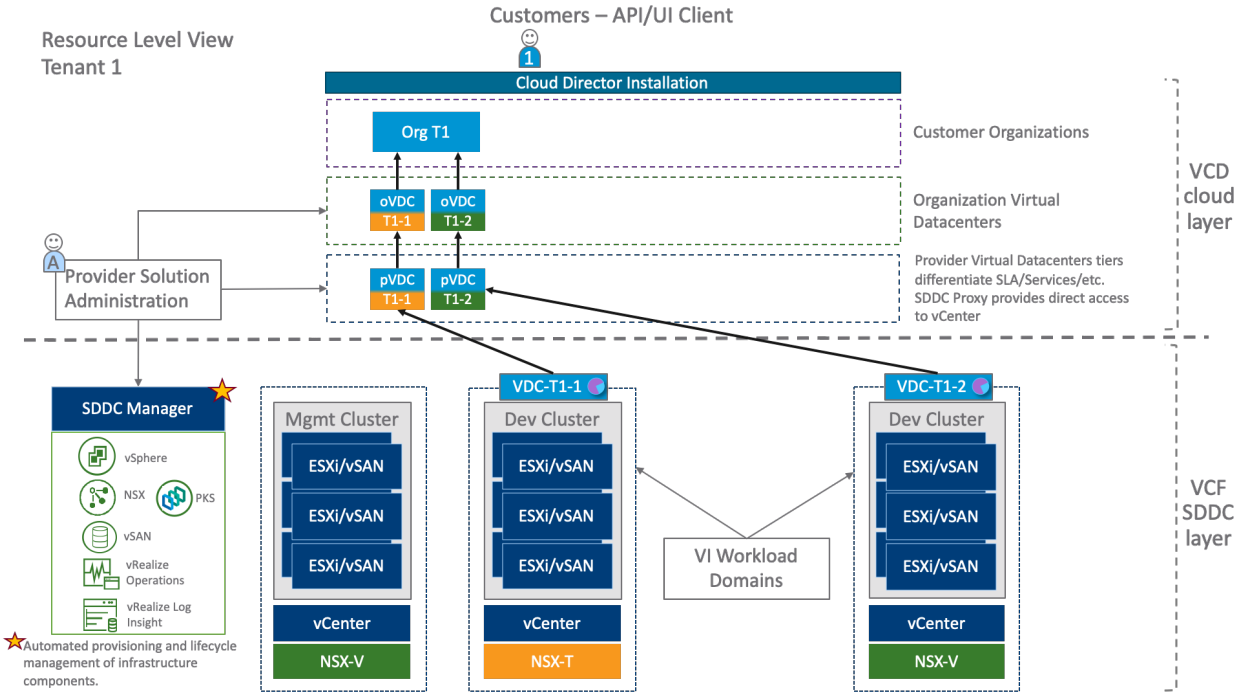

Here we also introduce some of the upstream management components and automation from VCF. The entire SDDC stack can be instrumented upon instantiation by being configured automatically to send logs to the already running and configured vRealize Log Insight cluster running in the management domain of VCF. Additionally, vRealize operations can also be deployed to provide operations management across the solution. Finally, VCF simplifies and automates patching and upgrading of the full SDDC stack with workload domain-level lifecycle management

4.4.4.2. Example 2 – Tenant Utilizing VCF Workload Domain and Dedicated Private Cloud using Central Point of Management (CPoM)

Figure 14 depicts a single customer utilizing a single VCF Workload Domain backed by NSX-T. You'll notice that there is also a PKS workload domain. VMware PKS is a container services solution to put Kubernetes in operation for multicloud enterprises and service providers. PKS simplifies the deployment and management of Kubernetes clusters with Day 1 and Day 2 operations support. PKS manages container deployment from the application layer all the way to the infrastructure layer according to the requirements for production-grade software. PKS supports high availability, autoscaling, health-checks and self-repairing of underlying virtual machines, and rolling upgrades for the Kubernetes clusters. VMware Cloud foundation automates the installation of the PKS control plane which makes it incredibly easy to deploy and manage VMware Enterprise PKS.

You can see that there is an "SDDC" box on top of the K8 Cluster vCenter, which is attached to the SDDC Proxy entity. Cloud Director can act as an HTTP/S proxy server between tenants and the underlying vSphere environment in VMware Cloud Foundation. A VCD "Software-Defined Data Center (SDDC)" encapsulates the infrastructure of an attached vCenter Server instance. An SDDC proxy is an access point to a component from an SDDC, for example, a vCenter Server instance, an ESXi host, or an NSX Manager instance.

Figure 14. Tenant Utilizing VCF Workload Domain and Central Point of Mangement (CPoM)

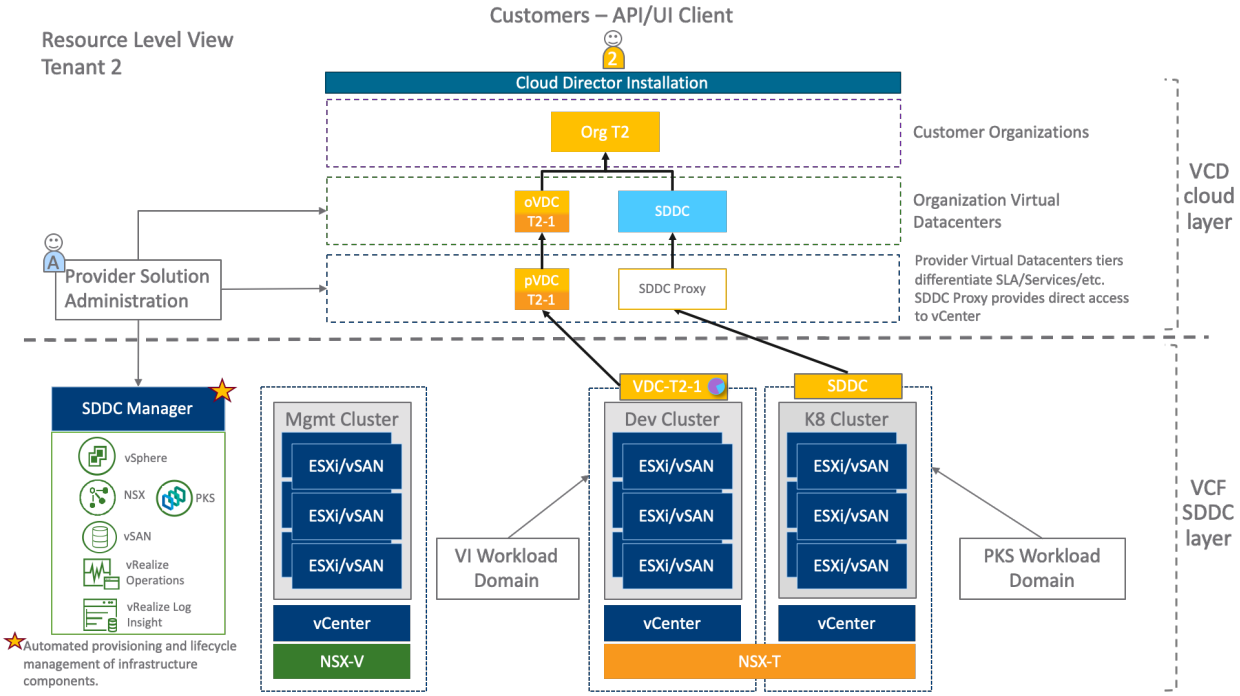

4.4.4.3. Example 3 – Tenant Utilizing Multi-Cluster VCF Workload Domain

Figure 15 depicts a single customer utilizing a single VCF Workload Domain. The workload domain is an NSX-V backed domain and would be representative of the majority of VCD deployments today. Each cluster has a single resource pool, mapped to a single pVDC.

Figure 15. Tenant Utilizing Multi-Cluster VCF Workload Domain

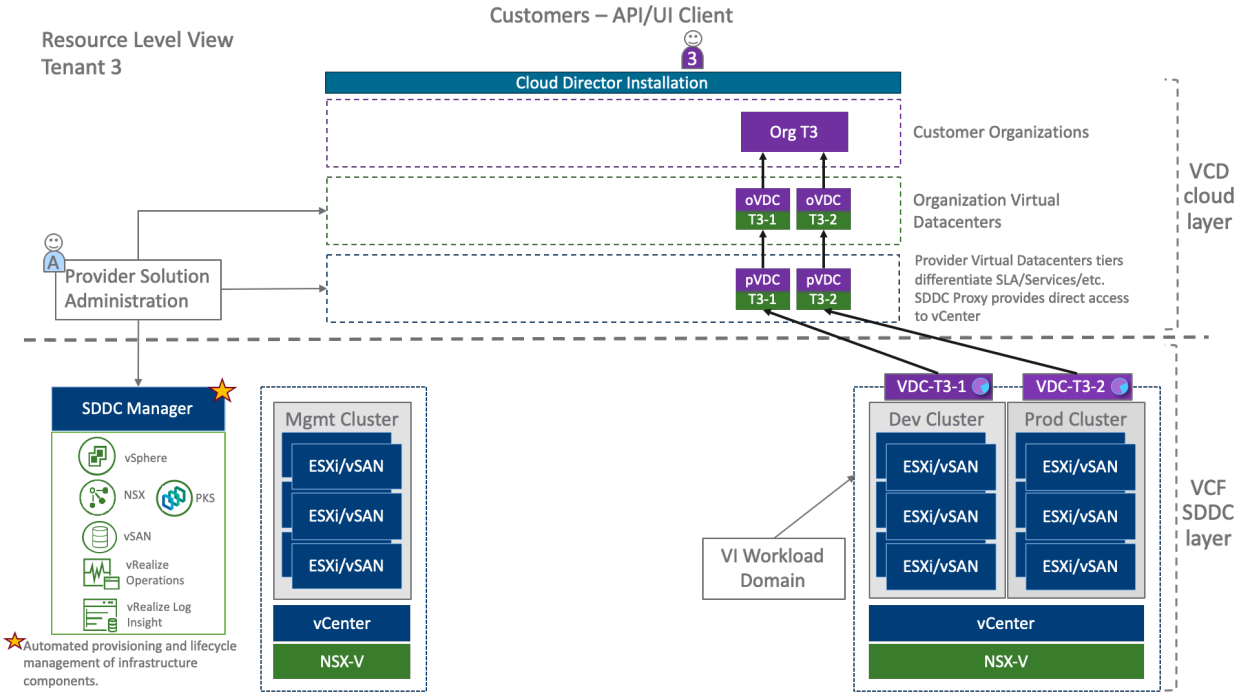

4.4.4.4. Putting it all Together

Finally we put it all together. In Figure 16 we can see that different customers (Orgs) are sharing resources from a single provider virtual datacenter. We can also see that resources from a single vCenter can be split across different provider virtual datacenters and that we can mix and match multi-tenants workload domains with different NSX backings and workload domains offering dedicated private cloud all together.

Figure 16. All 3 Tenants Together

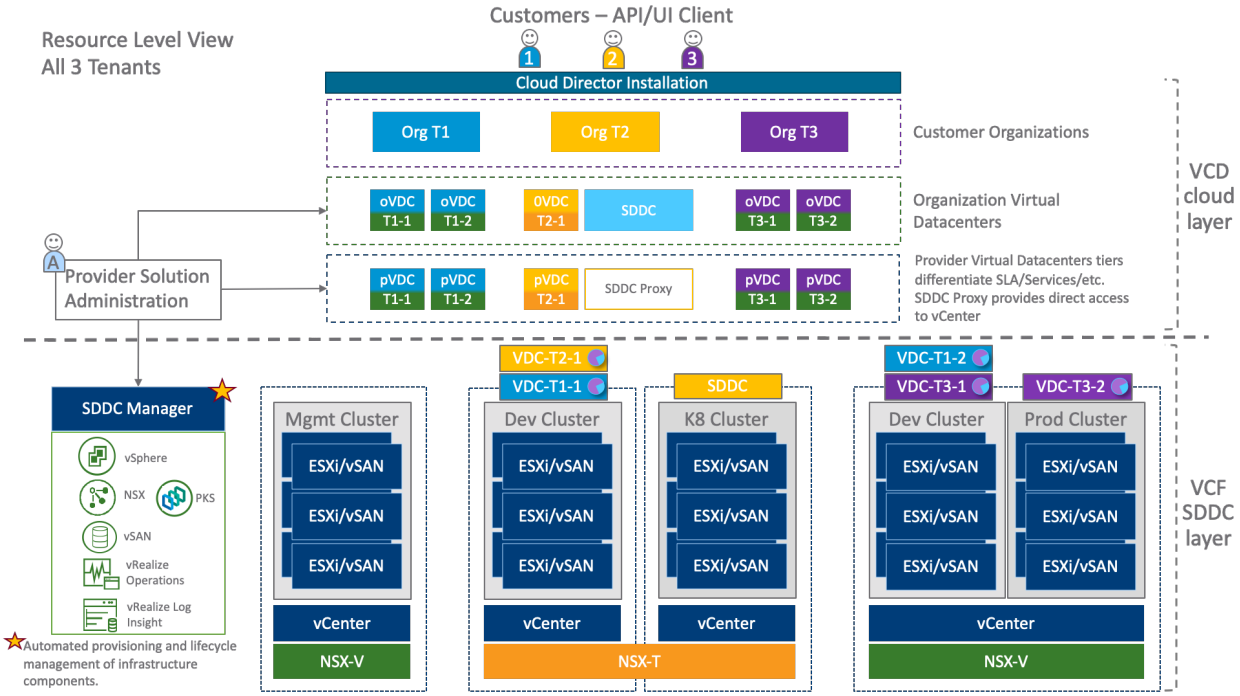

4.5 Example Architectures: Cloud Director with Cloud Foundation

4.5.1 Cloud Director Deployment Models in the Cloud Architecture Toolkit

The VMware Cloud Architecture Toolkit™ describes many generic design approaches for a Cloud Director environment. The following example architectures for deploying Cloud Director with Cloud Foundation are based on a subset of Cloud Director deployment models described in the toolkit.

4.5.2 Single Availability Zone

This deployment model is based on the following design principles:

- All Cloud Director and Cloud Foundation management components are hosted within the Cloud Foundation management domain.
- All Cloud Director resource groups are mapped to VI workload domains within the same single Cloud Foundation instance.

Following these design principles produces a fully self-contained Cloud Director environment in which all components—that is, management stack and resource groups—are hosted within a single Cloud Foundation installation. This approach enables easy configuration and operational procedures. As a downside, all components are hosted within a single availability zone. The loss of this availability zone implicates a loss of service for the Cloud Director and Cloud Foundation pair as well.

Figure 17 provides an overview of this deployment model.

Figure 17. Cloud Director with Cloud Foundation – Single Availability Zone

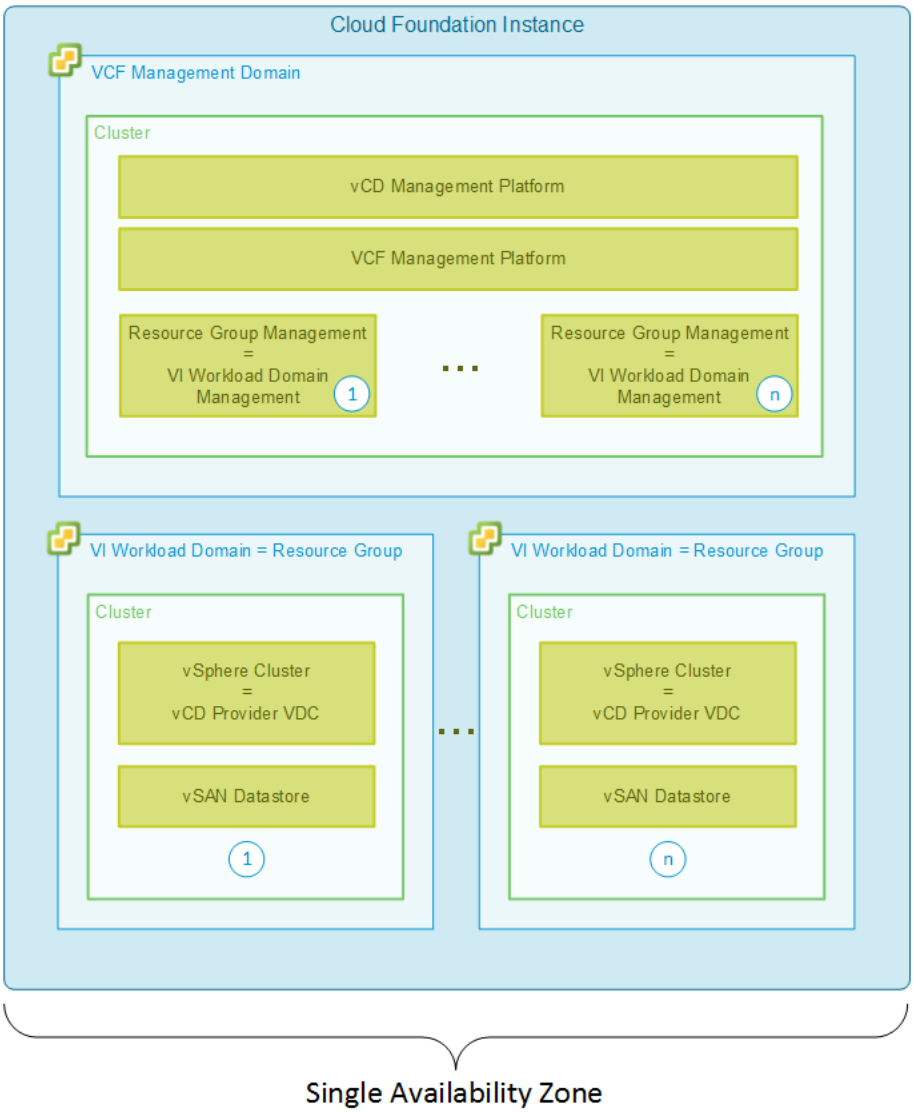

4.5.3 Dual Availability Zones

This deployment model is based on the following design principles:

- All Cloud Director and Cloud Foundation management components are hosted within the Cloud Foundation management domain.
- All Cloud Director resource groups are mapped to VI workload domains within the same single Cloud Foundation instance.
- The Cloud Foundation Instance is using VSAN stretched clusters across two availability zones.

Following these design principles produces a fully self-contained Cloud Director environment in which all components—that is, management stack and resource groups—are hosted within a single Cloud Foundation installation. In this deployment model the management domain is always stretched across availability zones. The VI workload domains can either be stretched across two availability zones or local to a single availability zone. As Cloud Director resource groups are

mapped to VI workload domains this scenario can be used for both distributed resource groups as stretched resource groups to protect workloads against a disaster or failure of an availability zone.

Figure 18 provides an overview of this deployment model.

Figure 18. Cloud Director with Cloud Foundation – Dual Availability Zones

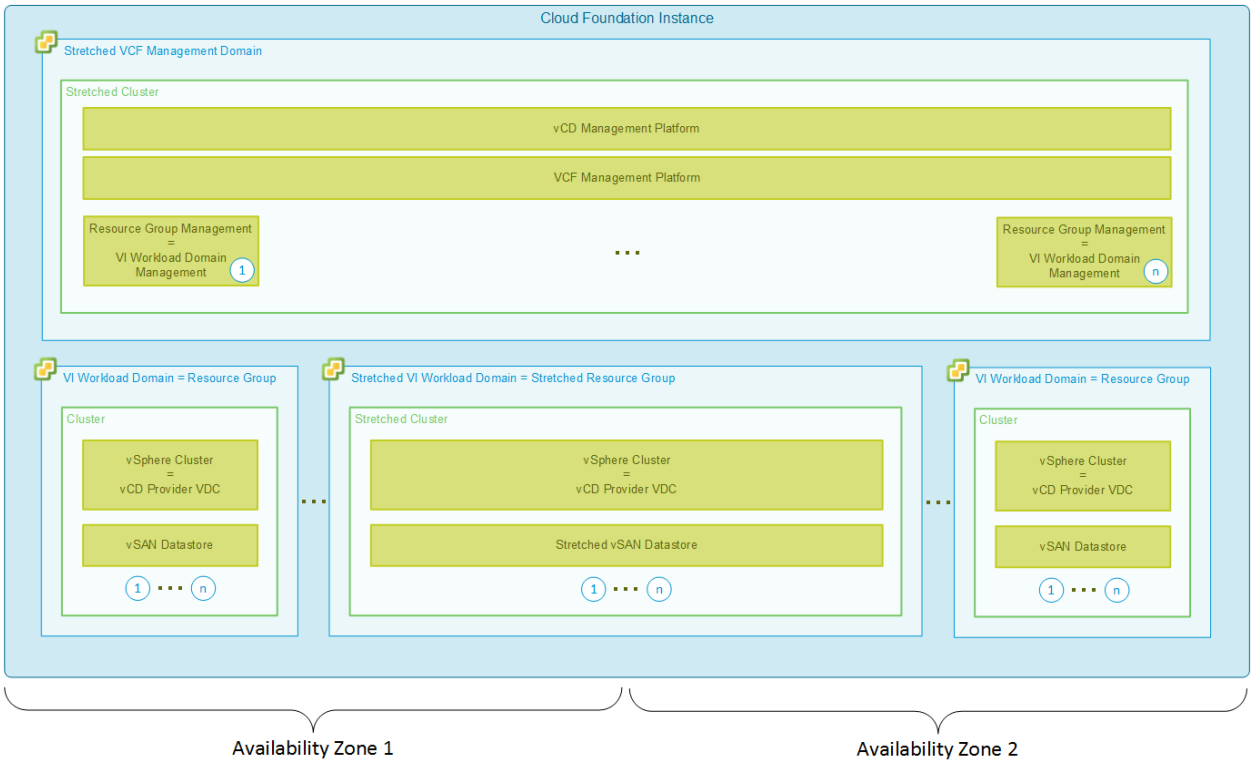

4.6 Networking Options

Cloud Foundation internal networking is deployed automatically by Cloud Foundation deployment workflows. It therefore offers a great simplicity at this layer. However, Cloud Foundation external networking options are mostly implemented manually. They therefore offer a high degree of flexibility when integrating a Cloud Foundation configuration into an existing data center.

4.6.1 Application Virtual Networks

Application Virtual Networks (AVN) are software defined overlay networks that abstract the hardware and realize the true value from a software-defined cloud computing model, laying the foundation for supporting workload mobility use-cases such as planned migration or disaster recovery. These networks can span a defined zone of clusters and traverse NSX Edge Service Gateways for their North-South ingress and egress and implements software-defined networking based on NSX in the management domain.

Application Virtual Networks provide the following key benefits:

- Provide a software-defined network topology for applications in VCF.
- NSX edge device for load balancing
- Simplified data mobility and future disaster recovery failover procedures
- Improved security and mobility of management applications

VMware Cloud Foundation provides the ability to deploy the full stack vSphere, vSAN, NSX and vRealize Suite as a single package along with SDDC manager. With each release of Cloud Foundation, there have been continual improvements to ensure our customers can take full advantage of what VCF offers for both new deployments and upgrades. As a design principle, VCF needs to be easy to install, easy to run, easy to upgrade, and easy to troubleshoot should something go wrong.

When VMware first introduced NSX, one of the most important features was the ability to stretch layer 2 networks over layer 3 segments without needing to change the underlying physical network. This is one of the core premises of "Software Defined Networking". NSX allows architects to build simple layer 2 networks that can span racks and geographically dispersed data center facilities.

Being able to stretch layer 2 networks has some obvious, and perhaps some not so obvious, advantages. The primary and arguably the most important advantage is workload mobility. Empowering admins with the ability to move workloads move between data centers without needing to re-IP, update DNS, and in some instances, change application configurations. This enables administrators to more easily perform disaster recovery with a fast RTO, perform non-disruptive data center migrations, and re-balancing workload placement.

Since Cloud Foundation allows different solutions to be implemented at any stage once the deployment is operational, no additional network requirements (such as adding VLAN-backed dvPortGroups) are needed when enabling solutions such as vRealize Suite, Horizon, and PKS. Finally, decoupling the VM networks from the underlying physical VLAN networks allows architects to fully integrate with external cloud providers and seamlessly join both on-premises and off-premises cloud networks.

Within the Cloud Foundation management domain, NSX Edges are deployed and peered (via Border Gateway Protocol – BGP) with the physical network to establish route redistribution to and from the SDDC. There are also new dvPortgroups in vSphere and new logical switches provisioned in NSX and connected to the north/south ECMP (Equal-Cost Multi Path) based on-ramp.

These AVNs have been pre-provisioned for all SDDC management VMs that are able to reside on the overlay network. Using the above as an example, the primary reason the VCF networking was designed this way is to enable future versions of the SDDC management components to failover to another region in the event of a disaster.

Preparing future versions of VCF to support seamless disaster recovery and application mobility are only a sample of the benefits delivered by Application Virtual Networking. Architects can also consider utilizing cross-region AVNs as an example implementation of how to configure mobility and provide disaster recovery for applications in VCF workload domains.

4.6.2 General Considerations

VMware NSX for vSphere is the core component within a Cloud Director implementation for providing automated software-defined networking services. In accordance with the VMware Product Interoperability Matrix, a version of NSX compatible with a corresponding Cloud Director version must be used.

As previously described, a separate NSX Manager and vCenter Server pair is deployed with a VI workload domain within Cloud Foundation. When the Cloud Director instance registers the vCenter Server installation of a VI workload domain as a resource group, it also actively registers and uses the NSX Manager instance as an API endpoint for providing software-defined networking services within this resource group.

The NSX instances that automatically spin off with every VI workload domain deployment are full-featured NSX installations with no limitations or special rules imposed by Cloud Foundation. Therefore, these instances can be consumed by the Cloud Director instance without special limitations.

The only rules to be evaluated when designing a Cloud Director solution with Cloud Foundation are those related to the network constructs that are introduced by Cloud Foundation. These constructs are briefly described in the following section.

The Cloud Architecture Toolkit also describes design best practices regarding external network connectivity for certain Cloud Director deployment models. When designing Cloud Director environments with Cloud Foundation, it is essential to understand the following external network connectivity options offered by Cloud Foundation.

Choice of NSX-V and NSX-T

For years, VMware NSX for vSphere had been our choice of products to virtualize an enterprise network with a software-defined, application-first approach. As the application landscape was changing with the arrival of public clouds and containers, NSX-T was being designed to address the evolving needs of organizations to support cloud-native applications, bare metal workloads, multi-hypervisor environments, public clouds, and now, even multiple clouds.

NSX-T and VMware Cloud Foundation

VMware Cloud Foundation supports NSX-T for Workload Domain. You can choose NSX-T when you create a new workload domain and if you use this VCF Workload Domain as VCD's Resource Group, a user can use NSX-T backed orgVDC.

NSX-T and Cloud Director

VMware Cloud Director now supports NSX-T. If you want to provide a orgVDC backed by NSX-T, you must connect NSX-T Manager to VCD. Currently with VCD 10.0, some important features including DFW, Load Balancer and VPN are not supported.

Best Practice Today

Since NSX-T is relatively a new comer to Cloud Director deployment, NSX-T does not offer all the features which NSX-V offers within VCD environment. Currently with VCD 10.0, some important features including DFW, Load Balancer and VPN are not supported. NSX-V will be staying as our choice solution until NSX-T acquires a full set of NSX-V features within VCD environment, except when a user needs container capability or compabililty with non-vSphere hypervisors.

4.7 Storage Options

Before VCF 3.5, vSAN was the only storage layer available and natively embedded in Cloud Foundation. Because vSAN is inherently built into the vSphere hypervisor, it is the natural choice as a storage layer for hyper-converged infrastructure (HCI) platforms based on vSphere. With VCF 3.9 and later, both NFS and VMFS on FC or additional storage options for VCF Workload Domains.

Cloud Director can consume vSAN storage service offerings and integrates with the Storage Policy–Based Management (SPBM) implementation of vSAN. In this context, Cloud Foundation does not introduce any special rules for consuming vSAN storage.

Certain aspects must be considered that are not recommended to be realized with vSAN—for example, regarding public catalog offerings within Cloud Director. But these are limitations specific to vSAN that are not tied to Cloud Foundation. Even when relying on vSAN as a primary storage option within Cloud Director, there are alternative, NFS-based solutions that can be additionally deployed.

A detailed discussion on how to develop a storage strategy for Cloud Director with vSAN is presented in the VMware technical paper *Developing a Hyper-Converged Storage Strategy for VMware vCloud Director with VMware vSAN*.

5. VVD for Cloud Providers - Cloud Foundation

VVD for Cloud Provider is the fastest and most reliable way for a Cloud Provider to build a multi-tenant SDDC. VVD for Cloud Provider deploys a SDDC stack powered by VMware Cloud Foundation. Users can deploy the SDDC in a few hours & clicks. This SDDC includes the latest VMware Cloud Foundation 3.9.x and Cloud Director 10.0 appliance and other components.

6. Conclusion

Tests internally conducted at VMware have clearly demonstrated the technical feasibility of implementing a pairing of VMware Cloud Director with VMware Cloud Foundation. This technical paper describes architectural options and provides guidance on how to design such an environment.

With its unique capabilities to provide automated Software-Defined Data Center (SDDC) hardware and software, stack bring-up services, and automated lifecycle management, Cloud Foundation is an excellent choice for service providers who are using Cloud Director and want a way to quickly and efficiently spin up SDDC resources ready for consumption with Cloud Director. Cloud Foundation deploys these resources with a self-contained architecture that is verified and validated by VMware.

Customers and service providers can avoid going through iterative, time-consuming design and architecture cycles to implement an SDDC stack consisting of VMware vSphere, VMware vSAN, and VMware NSX. Cloud Foundation "consumption-ready" units of SDDC deploy resources in an automated, repeatable, and quality-assured manner.

Service providers can face the requirement of hosting VMware Cloud Director managed resources onsite at their customers' data centers for data security and compliance reasons. Cloud Foundation enables them to achieve efficient integration into these data centers because it is a fully self-contained configuration that provides clear and easy interfaces with an existing data center, greatly simplifying deployment and reducing overall complexity.

7. Acknowledgment

7.1.1.1. Authors

Girish Manmadkar – Senior SDDC Architect – VMware Cloud Foundation

Arnim van Lisehout - Staff Consulting Architect – VMware PSO

Daniel Paluszek - Staff Cloud Provider Solutions Engineer

7.1.1.2. Reviewers

Jorge Lew – Senior Consultant – Cloud Provider Software

Ken Lamoreaux - Director Technical Product Manager, Cloud Provider Software

Luis Ayuso - Senior Product Marketing Manager, Cloud Provider Software

7.1.1.3. Executive Sponsors

Peter Wei – Vice President – VMware Cloud Foundation

Rajeev Bhardwaj – Vice President – Product Management , Cloud Provider Software

8. Appendix

8.1 References

VMware Cloud Architecture Toolkit for Service Providers Documentation Center

VMware Cloud Director Product Page

VMware Cloud Foundation Product Page

VMware Cloud Foundation Blog Page

VMware Cloud Foundation Quick Reference Page

8.2 Software Versions

This paper relies on an evaluation's being conducted with the following software versions, which are the current versions of the corresponding products as of the time of writing this document:

- VMware Cloud Foundation 3.9.2 and the corresponding versions for VMware vSphere, VMware NSX, and VMware vSAN.
- VMware Cloud Director 10.0

When planning a Cloud Director environment with Cloud Foundation or an upgrade of such an environment, ensure that the resulting software bill of materials complies with the VMware Product Interoperability Matrix. Concerning the versions of vSphere, NSX, and vSAN, the Cloud Foundation architecture ensures interoperability of these product versions shipped within a particular Cloud Foundation release. Cloud Director, however, is an external component from the perspective of Cloud Foundation. Interoperability between particular versions of Cloud Director and those of vSphere, NSX, and vSAN within a Cloud Foundation release must be verified when planning an installation or upgrade of such an environment.

Design rules and configuration limits presented in this document are applicable for the software versions mentioned. Regarding upcoming Cloud Foundation releases, these design rules and configuration limits might be subject to change, and additional options might become available.

8.3 VMware Cloud Director Footprint

a Cloud Directory deployment can consist out of many components, especially in a production environment. In Table 3 below the Cloud Director primary components are listed. Table 4 lists optional components that can be installed alongside Cloud Director for additional functionality. These tables are provided to give you initial guidance on sizing requirements for your Cloud Director deployment. Please consult the product documentation on the VMware's website for version specific sizing requirements.

Table 3. Cloud Director Primary Components

Component	vCPU	Memory (GB)	Disk (GB)	Version
NFS Server	2	8	1024	
Cloud Usage Meter	2	4	60	3.6.1 Hot Patch 3 (Build 14877528)
Cloud Director Load Balancer	2	1	2	
Cloud Director Cell 01	2	8	56	10.0.0.0 (Build 14638910)
Cloud Director Cell 02	2	8	56	10.0.0.0 (Build 14638910)
Cloud Director Cell 03	2	8	56	10.0.0.0 (Build 14638910)

Table 4. Cloud Director Optional Components

Component	vCPU	Memory (GB)	Disk (GB)	Version
RabbitMQ Load Balancer	2	1	2	
RabbitMQ Node 01 – for Cloud Director	2	6	54	version 3.7
RabbitMQ Node 02 – for Cloud Director	2	6	54	version 3.7
Cloud Director Cassandra Perf DB – Node 01	4	16	1016	Cassandra versions 3.x
Cloud Director Cassandra Perf DB – Node 02	4	16	1016	Cassandra versions 3.x
Cloud Director Cassandra Perf DB – Node 03	4	16	1016	Cassandra versions 3.x
Cloud Director Cassandra Perf DB – Node 04	4	16	1016	Cassandra versions 3.x
vRealize Network Insight	12	48	1024	5.0.0 (Build 1568279774)
vRealize Network Insight - Proxy	8	16	150	5.0.0 (Build 1568279774)
vRealize Operations Manager Load Balancer	2	1	2	
vRealize Operations Manager – Node 01	8	32	275	8.0.0 (Build 14857692)
vRealize Operations Manager – Node 02	8	32	275	8.0.0 (Build 14857692)
vRealize Operations Manager – Node 03	8	32	275	8.0.0 (Build 14857692)
vRealize Operations Manager - Tenant UI	2	8	90	2.3.0 (Build 14826907)
vRealize Orchestrator Cloud Director	2	6	27	7.6.0 (Build 13020602)