



© 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com



Contents

Introduction	7
1.1 Overview	7
1.2 Document Purpose and Scope	7
1.3 Definitions Acronyms and Abbreviations	8
Customer Networking In a Service Provider Environment	10
2.1 Customer Network Topologies.....	10
2.2 Replicating a Managed Service Customer Topology	11
Multitenancy in a Cloud Service Provider	17
3.1 vCloud Director Multitenancy	18
3.2 Basic vCloud Director Tenant Topology	21
3.3 Multitenant Networking.....	22
Networking Layers Examined	24
4.1 Tenant Networking.....	24
4.2 vCloud Director Multitenant Data Center Networking in vSphere.....	26
4.3 Networking in a Multi-Cluster “Leaf-Spine” Infrastructure Topology.....	28
4.4 vCloud Director Multitenant Networking in NSX	30
IP Address Management and Routing	31
5.1 Tenant Address Management.....	31
5.2 Customer Address Assignment	33
5.3 Internet Address Management.....	36
5.4 External Network Address Sub-Allocation	38
5.5 Routing in a Multitenant Service Provider Environment	40
5.6 IPv6 Considerations	43
Commercial Considerations	43
6.1 Managed Service or Self Service.....	43
6.2 Additional Product Licensing.....	44
References	46
Appendix A: Provisioning an External Network in vCloud Director	47



List of Tables

Table 1. VMware vCloud Director Tenancy Terms	18
Table 2. vCloud Director Tenancy and Network Elements in NSX and vSphere	30
Table 3. Appendix A Configuration Parameters.....	47

List of Figures

Figure 1. Example Managed Service Customer Topology	12
Figure 2. Basic Cloud Service Provider Customer Topology.....	14
Figure 3. Micro-Segmentation with the Distributed Firewall	15
Figure 4. Basic Customer Topology with NSX Distributed Firewall.....	16
Figure 5. Example Managed Service Multitenant Data Center Topology.....	17
Figure 6. Layered Software Transition from VMware Cloud Provider Program MSP to CSP	18
Figure 7. Service Provider Data Center vSphere Clusters and Resource Pools.....	19
Figure 8. Mapping Provider VDCs to vSphere Resources	20
Figure 9. Basic vCloud Director Customer Topology.....	21
Figure 10. Managing Edge Gateways and Org VDC Networks in vCloud Director	22
Figure 11. Managing the NSX Distributed Firewall in vCloud Director	22
Figure 12. Example Cloud Service Provider Multitenant Data Center Topology.....	23
Figure 13. Cloud Service Provider Tenant Networking.....	24
Figure 14. vSphere Cloud Service Provider Tenant Networking	25
Figure 15. vSphere Cloud Service Provider Multitenant Networking.....	26
Figure 16. Configuring a dvSwitch Uplink Port Group VLAN Trunk	27
Figure 17. Trunking Multiple External Networks to a vCloud Director Environment.....	28
Figure 18. NAT on the Org VDC Edge Services Gateway	32
Figure 19. Error Message Caused by Provisioning an Overlapping Address Range.....	33
Figure 20. Static IP Pool Address Assignment	33
Figure 21. Org VDC Network Static IP Pool Configuration.....	34
Figure 22. Enabling the Enhanced Networking Workflow.....	34
Figure 23. Static – Manual Address Assignment.....	34
Figure 24. DHCP Address Assignment.....	35
Figure 25. Internet vCloud Director External Network	36
Figure 26. Outbound Address Translation (SNA/PAT)	38
Figure 27. External Network IP Address Assignment.....	39
Figure 28. Sub-Allocation of External Network Addresses	39
Figure 29. Provider Internet Routing.....	40

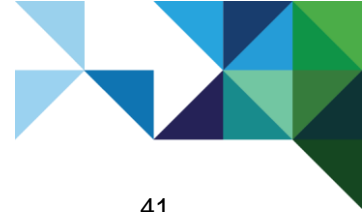


Figure 30. Routed DMZ Org VDC Network..... 41

Figure 31. Per-Tenant WAN Router Peering 42

Figure 32. Managing Certificates in the Edge Services Gateway..... 43

Figure 33. Cloud Service Provider Data Center with Per-Tenant Licensing Options 45

Figure 34. Adding a New External Network..... 47





Introduction

1.1 Overview

The VMware Cloud Provider Program is an ecosystem of over 4,000 service providers located in more than 100 countries offering VMware based cloud services. Local providers secure data sovereignty while providing a wide range of cloud services and vertical market expertise through specialized compliance and certifications.

VMware Cloud Providers are uniquely positioned to offer their services to the market and become a seamless extension of existing VMware enterprise customers' on-premises data centers. Managed service providers have traditionally operated data centers hosting multiple tenants and have developed reference topology models offering tenant separation while allowing the provider to offer value-add services that require access into each tenant environment from common management platforms.

One of the initial concerns about offering a multitenant platform using a customer-facing, overlaid management solution, such as VMware vCloud Director®, is understanding the capabilities and constraints for the customer solutions that can be delivered and for the service provider infrastructure which will support these solutions. With the inclusion of VMware NSX® in later releases of vCloud Director, the ability to create complex customer topologies has increased, but so have the concerns which can justifiably be raised over the deployment of both products technically and operationally.

1.2 Document Purpose and Scope

This document examines the way familiar customer topologies can be replicated within a VMware Cloud Provider's offering using the capabilities of NSX, controlled through vCloud Director. While this document focuses on the architecture of a service provider's multitenant environment, it starts with the evolution of a managed service customer's topology into that of a cloud service customer. Having established a simple, illustrative customer topology, this document then explores the way this topology is implemented on the underlying infrastructure, and then offers example topologies for separating multiple customers within both the virtualized world and that of the larger data center and its external connectivity.

While vCloud Director requires its own supporting network environment and configuration to provision compute workloads, those areas are outside the scope of this document. Because of this, a degree of familiarity is assumed with both vCloud Director and NSX. For more information on architecting and operating a vCloud Director solution, see Section 7, References. Similarly, while the infrastructure topology of access networks used by service providers to manage tenant workloads, applications and services follow similar models to those explored for customer networking, the specific challenges of providing management access to tenant networks is out of scope for this document.

This document is designed to assist service provider infrastructure architects tasked with the design of tenant networking within the data center environment. Its goal is to also assist service provider customer solution architects with the pre- or post-sale design of customer environments within a vCloud Director powered cloud service.



1.3 Definitions Acronyms and Abbreviations

1.3.1 Definitions

Customer	The service provider's customer; the organization who pays for the service, and the users who use the service.
Tenant	The portion of the infrastructure that is used by, and provides services to, the customer.
External network	A network that connects the vCloud Director tenant to the data center network infrastructure.

1.3.2 Acronyms and Abbreviations

AS	Autonomous System is a collection of routing nodes within the same administrative boundary (see BGP).
BGP	Border Gateway Protocol is a dynamic routing protocol used within the public Internet and private wide area networks. It can be used with a single autonomous system as an internal protocol (iBGP) and externally between autonomous systems (eBGP).
CE	Customer Edge (router) is a term used to describe the role of a service provider router that sits on the perimeter of the provider network and connects to the local network on a customer's premises.
CSP	Cloud Service Provider is the term used for a service provider capability that offers some or all of the key attributes necessary to be recognized as providing a cloud service. See <i>The NIST Definition of Cloud Computing</i> link in Section 7, References for more information.
DHCP	Dynamic Host Configuration Protocol (rfc2131 and updates) is a protocol that enables a host to initialize on an IP address based network without a configured IP address. The protocol sets out a process for the host to assume a temporary address and request an IP address from a local or remote authority on the network to which it is connected.
ECMP	Equal Cost Multi-Path is a routing term in which multiple next-hop choices all carry the same preference or "cost", allowing traffic to be distributed across several links or devices to increase network resilience and performance.
ESG	VMware NSX Edge™ services gateway, sometimes simply "Edge", or within vCloud Director "Edge Gateway" is a network and security virtual appliance that provides a number of services. See Section 2.2.2, Basic Cloud Service Provider Customer Topology for more information.
MPLS	Multi-Protocol Label Switching is a network packet forwarding technology often used by service providers in their high-speed core networks. It uses hop-by-hop labels instead of destination addressing to enable traffic engineering, or management, of network paths and traffic flows.
MSP	Managed Services Provider is the term used for a service provider capability that does not typically provide a customer facing, self-service portal which customer use to directly control their environments.
NAT	Network Address Translation is a technique in which the source and/or destination addresses within an IP packet's header are changed to hide the real address of a service. This is used for instance when devices on a private network connect to the public internet.
NIC	Network Interface Card (NIC) is the network adaptor which connects the ESX hosts to the external network infrastructure.



- OS** Operating system is the layer of software deployed onto the hardware layer of a physical or virtual computer. In VMware ESXi™ based solutions, the physical hardware uses VMware ESXi as its operating system and the virtual machines which ESXi supports run a variant of Microsoft Windows or a Linux distribution as the “guest OS”.
- PE** Provider Edge (router) is the term used to describe the role of a service provider router that connects multiple Customer Edge (CE) routers to the Provider’s core (or “P”) routers. In some provider platforms, the role of the CE router is virtualized, so the PE router is physically connected directly to the infrastructure supporting the customer solutions rather than through a physical per-tenant CE router.
- SDN** Software-defined networking is a technology that creates end-user network elements, defined in software which are then deployed on top of a physical “underlay” network.
- VCDNI** vCloud Director Network Isolation is a proprietary “MAC in MAC” encapsulation used in earlier versions of vCloud Director which allowed multiple, isolated customer networks to be “tunneled” between hosts over a single network.
- VDC** Virtual data center is a collection of resources that are managed by vCloud Director. See Table 1 for more information.
- VIP** Virtual IP is a term for an additional IP address that provides access to one or more devices without being assigned permanently to any of them. It is found in solutions where two or more devices provide high availability by presenting a single IP address for clients to connect to, without needing to know which device will service their request.
- VLAN** Virtual LAN is a network protocol which allows multiple, separate Layer 2 networks to be carried on the same physical (Layer 1) medium.
- VPN** Virtual private network is a technique for separating traffic within a shared infrastructure. It is often used when an encrypted overlay is added to an insecure shared network such as the Internet, or when customer separation is provided over a service provider WAN.
- VRF** Virtual routing and forwarding is a technique in which a single network device can manage multiple independent routing tables at once and apply the resulting forwarding rules to traffic associated with a specific instance.
- VXLAN** Virtual eXtensible LAN is an encapsulation protocol that allows multiple, separate Layer 2 networks to be carried over a common Layer 3 network.
- WAN** A wide area network is a telecommunications network that typically spans a large geographical area.



Customer Networking In a Service Provider Environment

As a Managed Service Provider, building complex customer topologies with VMware vSphere® environments is constrained only by the capabilities and topologies that can be built using the network configuration within vSphere together with that of the underlying data center network infrastructure. However, when customers want to change their network configuration, traditional services fall short in two distinct areas. First, while not all customers are capable of designing their own network configurations, those that are have previously had little, if any, access to be able to make changes within the service provider managed environments. Instead, the customer must raise an order, or request the engagement of a service provider solution architect to capture their requirements and provide them to the service provider's operational teams. The second shortfall in customer experience is that presented with the customer requirements, the operational teams must work out the necessary changes to implement the customer requirements, document the changes, raise appropriate change request or notifications, await approvals and scheduled outage windows before work can commence. After the work starts, coordinating changes across disparate teams and technologies is challenging, testing is complex and time consuming, and recovery planning must account for failures across all areas affected by the change.

Having simplified the delivery of complex customer compute solutions by virtualizing the server environments with vSphere, the service provider can use a similar approach to simplifying the networking layer of a customer solution by introducing NSX to provide software-defined networking (SDN). By making use of a common “underlay” network within the physical data center, NSX allows the service provider to configure complex, multi-customer network topologies and services without the need to reconfigure the underlay network each time. Further, because these services are defined in software, the second of the two shortfalls described earlier can be addressed in several ways. The first, and most straight forward, leverages the fact that a software-defined network can be managed from a single point. While it is possible to automate the configuration of multiple interconnected legacy network devices, this automation is often complex and error prone due in part to the fact that many of these devices were not originally designed with remote configuration management in mind. With NSX, a service provider network engineer can, from a single location, configure, build, deploy, modify, and decommission network connectivity and services for any tenant in the data center. While this central control is a huge help, efficiency and consistency gains are often made through the automation of repeatable tasks. NSX offers a feature-rich API that allows configuration of those same networks and services from external tools. The APIs can be used to create and manage changes, or to “read” the state of a customer environment for monitoring or compliance purposes.

vCloud Director allows the service provider to address the first of the shortfalls mentioned earlier, and the one which is arguably, most apparent to the customer. It provides a portal to allow customers to carry out their own provisioning and changes, store frequently used configurations as templates, and access templates that have been created and shared by the VMware Cloud Provider. vCloud Director is able to control the underlying vSphere environments it manages through the APIs exposed by VMware vCenter Server® and other vSphere components. The introduction of similar API access to the software-defined network layer as noted earlier means that in the latest releases (v8.20 at the time of writing), customers can now manage their environment's NSX provided networks and services through the same vCloud Director portal and vCloud Director API.

2.1 Customer Network Topologies

With complete access to the interfaces of all of the elements used to build a customer solution, service providers can use the entire range of capabilities exposed through those interfaces so that each customer receives a bespoke environment, perfectly suited to meet their requirements. However, this presents challenges across the service provider organization. Each bespoke design must be carefully “documented” so that operations and support teams can understand a solution when they are asked to work on it or for infrastructure teams when work is required on the data center's power or cooling systems. In the commercial side of the business, tariffs and billing systems must be able to cope with each customer's individual configuration, and the complexities of carrying out changes reflected in any associated charges.



Standardizing network topologies while still offering customers sufficient flexibility is key to increased speed and efficiency in the service delivery. The introduction of a Cloud Management Platform (CMP) by its nature abstracts the users, both service provider and customer, from the interfaces of the underlying infrastructure. In doing so it brings some constraints to the range of features it exposes from that underlying infrastructure. However, by adding cloud capabilities, a Managed Service Provider can begin to realize these speed and efficiency gains which the enforced “standardization” of the CMP brings.

2.2 Replicating a Managed Service Customer Topology in vCloud Director

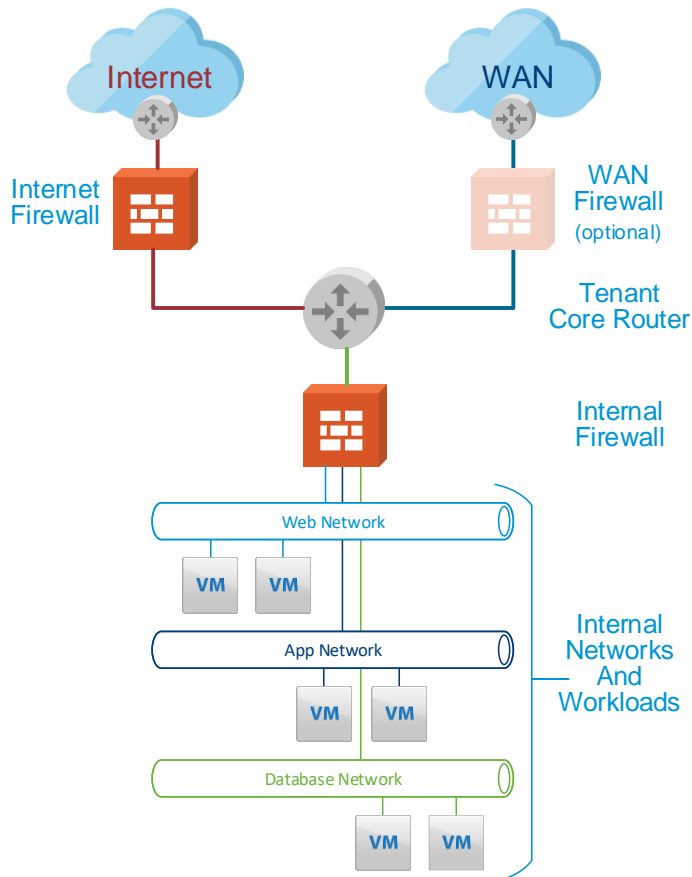
When approaching the adoption of vCloud Director and NSX from the perspective of a customer or the service provider Solution Architect, it is important to understand how the familiar customer topologies of the Managed Service Provider world can be replicated in the portal-driven Cloud Service Provider world. This section explores the traditional customer topology model, examines the new building blocks which vCloud Director and NSX provide and the equivalent customer topologies that are then built within the Cloud Service Provider platform.

2.2.1 Traditional Managed Service Customer Topology

Customer topologies, particularly in managed services bespoke deployments, can take many forms, but most share common traits. They typically have some or all of the following:

- External access from untrusted networks such as the internet.
- External access from trusted or semi-trusted, corporate wide area Networks.
- Perimeter security on some or all ingress paths.
- Separate networks within the solution for administrative separation of solution components.
- Internal security for controlled separation of solutions components.
- Routing/switching to allow the solution components to communicate with each other.

The following figure is an example “three tier” network that illustrates typical elements.

**Figure 1. Example Managed Service Customer Topology**

In the figure, the WAN connection is separated from the core router with a firewall. Depending upon the level of trust which the WAN and the workload warrant, as well as the customer's appetite for risk, this might not be required. When there is no firewall, a single network (usually a VLAN) connects the tenant core router directly to the WAN presentation, typically in the form of a dedicated, physical, per-tenant Customer Edge (CE) router. This simplified model will be used as the basis for illustrations throughout this document although it is acknowledged that this is far from the only topology that could be deployed.

2.2.2 The NSX Edge Services Gateway

vCloud Director allows the provisioning of compute workloads (and their associated storage) on the underlying vSphere environments which it controls as well the provisioning of networks to support them. With the ability to consume and manage NSX in recent versions of vCloud Director, networks are now provisioned using NSX managed VXLAN rather than vCloud Director Network Isolation (VCDNI) which was used in earlier versions. In addition to these NSX networks, vCloud Director now provides the ability to provision and manage an NSX network and security appliance called the Edge Services Gateway (ESG) rather than the VMware vCloud® Network and Security™ edge used in earlier versions.

Unlike virtual networks which are implemented within the VMware ESXi hypervisor, the Edge Services Gateway is a network appliance virtual machine (or virtual appliance) with interfaces that connect to the networks within the solution. In this position, the Edge Services Gateway can provide a number of network services, the main ones of which are the following:

- Routing – Using static routes and/or dynamic routing protocols.
- Firewalling – To provide filtering of “North/South” traffic entering or leaving the solution.



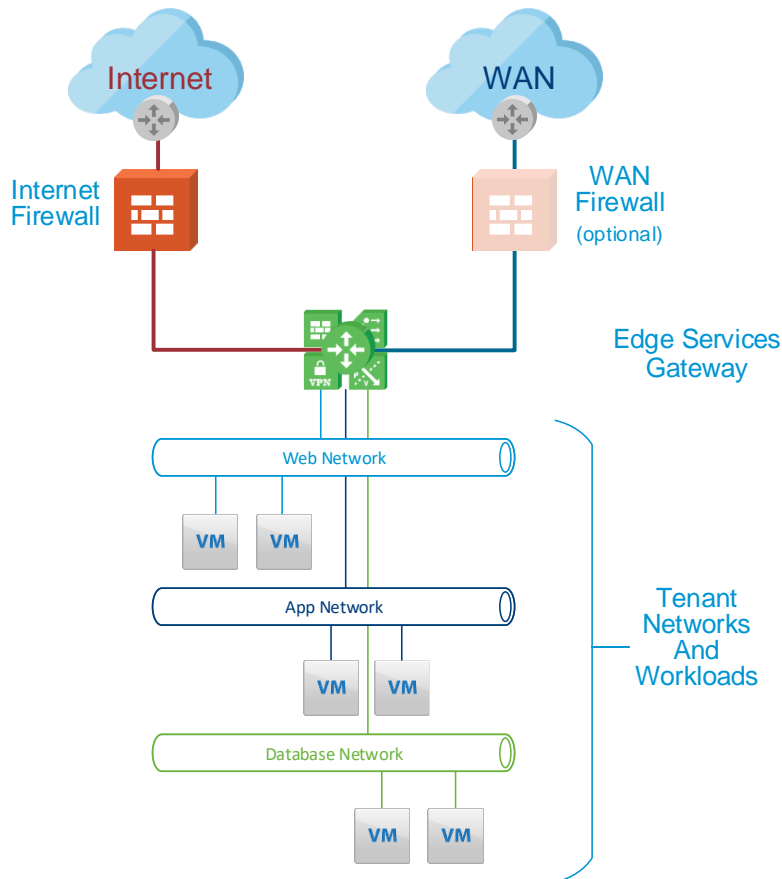
- Network Address Translation – Of either source, destination or both addresses.
- Load balancing – At either Layer 7 for greater feature capabilities or Layer 4 for greater throughput.
- VPN termination – Either Layer 3 site-to-site or client VPNs or Layer 2 VPN¹ to allow bridging of hybrid solutions where part of the solution sits outside of the vCloud Director environment or data center.
- Physical-to-virtual interconnection – Allows routing between external, physical networks (in the form of VLANs) and internal NSX virtual networks (in the form of VXLANS).
- DHCP/DNS – The Edge Services Gateway also supports DHCP (as a server or relay/helper) and a DNS forwarder.

Providing both routing and firewalling capabilities, the Edge Services Gateway can perform the roles of both tenant core router and internal firewall as shown in Figure 1. The equivalent customer topology built using the Edge Service Gateway is described in the following section.

2.2.3 Basic Cloud Service Provider Customer Topology

Moving a customer from a traditional MSP environment to the service provider's cloud platform does not change the basic requirements on which the service was originally built. The attributes listed in Section 2.2.1, Traditional Managed Service Customer Topology still apply. Translating the topology to a vCloud Director managed service results in a very similar topology which is shown in the following figure.

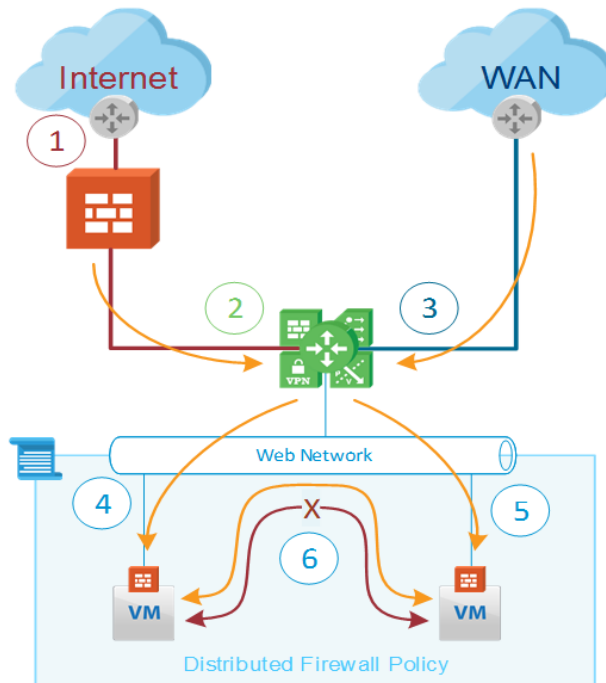
¹ See *Streamlining Customer On-Boarding with NSX L2 VPN Services* in the References section for more information.

**Figure 2. Basic Cloud Service Provider Customer Topology**

With the introduction of NSX, the roles of the physical tenant core router and internal firewall are replaced by the Edge Services Gateway. Because this device provides North/South firewall capabilities, it can also remove the need for a dedicated WAN firewall, saving the service provider cost and complexity so contributing to both CapEx and OpEx reductions. While the topology seems almost identical to that shown earlier, what is not visible in the figure is the customer's ability to deploy compute and network services from the Edge Services Gateway southbound through the vCloud Director portal. This addresses the two shortfalls identified with MSP services, the customer's ability to make changes to their own environments and the speed and reliability with which those changes can be implemented.

2.2.4 The NSX Distributed Firewall

In addition to VXLAN-backed networks and the Edge Services Gateway, the introduction of NSX features in vCloud Director brings another key benefit. The presence of the NSX components in each ESXi host, allows vCloud Director customers to utilize the NSX Distributed Firewall (DFW). The NSX Distributed Firewall implements a stateful packet filtering capability on each vNIC of every virtual machine (VM) under its management. In addition to the control of traffic arriving at the VM from outside the network to which the VM is connected, the Distributed Firewall allows control of traffic between VMs on the same network. This granular control of traffic within the same network is known as micro-segmentation. Micro-segmentation allows a degree of control over traffic that has already been allowed through the perimeter firewall which has not, previously, been possible. The following illustrates a portion of the Distributed Firewall "policy" applied to traffic to and from two "web server" VMs.

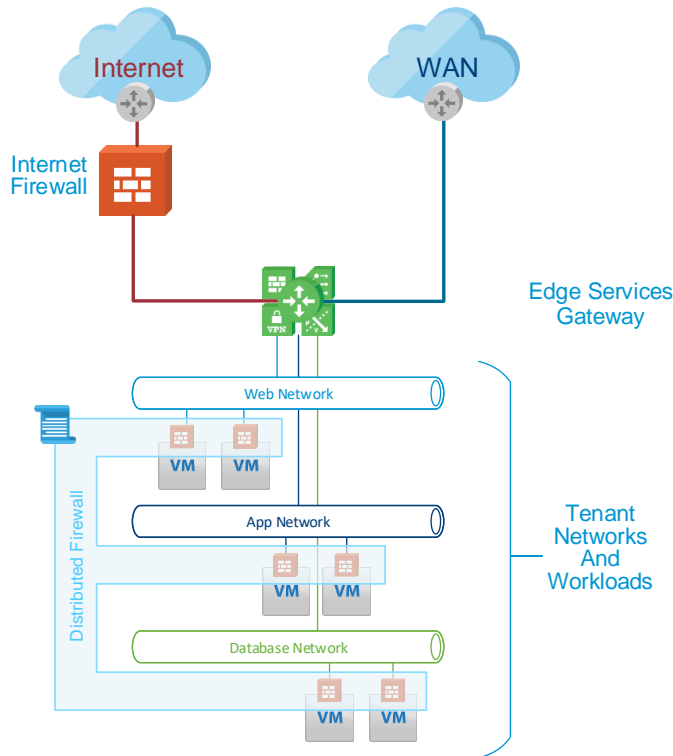
**Figure 3. Micro-Segmentation with the Distributed Firewall****DFW Policy Extract**

1. Incoming traffic from the internet is first filtered by, in this example, an external, physical firewall.
2. Traffic from the internet firewall is then processed by the Edge Services Gateway. This can include a second tier of packet filtering, Network Address Translation (NAT) or load balancing.
3. Traffic from the WAN is processed with the same services as the internet traffic. Because the traffic is potentially more trustworthy, the firewall policy for WAN traffic can be more open.
4. Traffic to the left web server from the internet, or the load balancing capability of the Edge Services Gateway is filtered by the DFW policy on the VM before it is allowed to reach the web server.
5. Traffic to the right web server is filtered independently and while it might have a ruleset that matches the other VM, it could also have different rules for different services.
6. Traffic between the VMs is filtered at the same point as traffic reaching the VMs from external sources. Traffic that must be allowed between the web servers, such as application clustering or monitoring/heartbeat traffic, is explicitly allowed by the DFW policy. In contrast, all other traffic between the web servers is explicitly denied. In addition, the DFW policy can log traffic between the web servers that is dropped by the policy. The presence of attempted connections between web-servers might indicate the presence of a security attack, and could be used to trigger an alert from a security event management system. In this simple example, the presence of micro-segmentation would help to prevent an attack from spreading and alert the provider to the attack in a way which was previously impractical to implement at scale.

The DFW policies that control the flow of packets to and from each VM's vNICs are configured centrally through the new vCloud Director HTML5-based firewall management interface and then distributed to the ESXi host to implement on the running VM. Should the VM be moved, the policy is reapplied on the destination host. This distributed policy management with per vNIC implementation when applied across the entire tenant network topology is shown in the following figure.



Figure 4. Basic Customer Topology with NSX Distributed Firewall



In the illustration, each vNIC is effectively separated from the network by a “firewall”. This allows the traffic between VMs on the same network to be controlled by a firewall policy in a way that is not practical using traditional networking infrastructure. See Section 7, References for more information on the benefits of micro-segmentation. Both the North/South policy on the Edge Services Gateway and the Distributed Firewall policy itself, can be managed from the vCloud Director user interface either by the service provider or, should the service provider choose to offer the facility, the customer themselves.

Because the Edge Services Gateway offers both firewall and VPN services, customers might question the presence of two firewalls in the internet connectivity, particularly as the benefits of micro-segmentation become accepted. At this point, the service provider can further reduce cost and complexity by removing the physical Internet firewall and moving its roles to the Edge Services Gateway. This level of risk acceptance is likely to be different for each customer, and it is possible that service providers might choose to offer solutions both with, and without, the physical internet firewall.

2.2.5 Further NSX Feature Support in vCloud Director

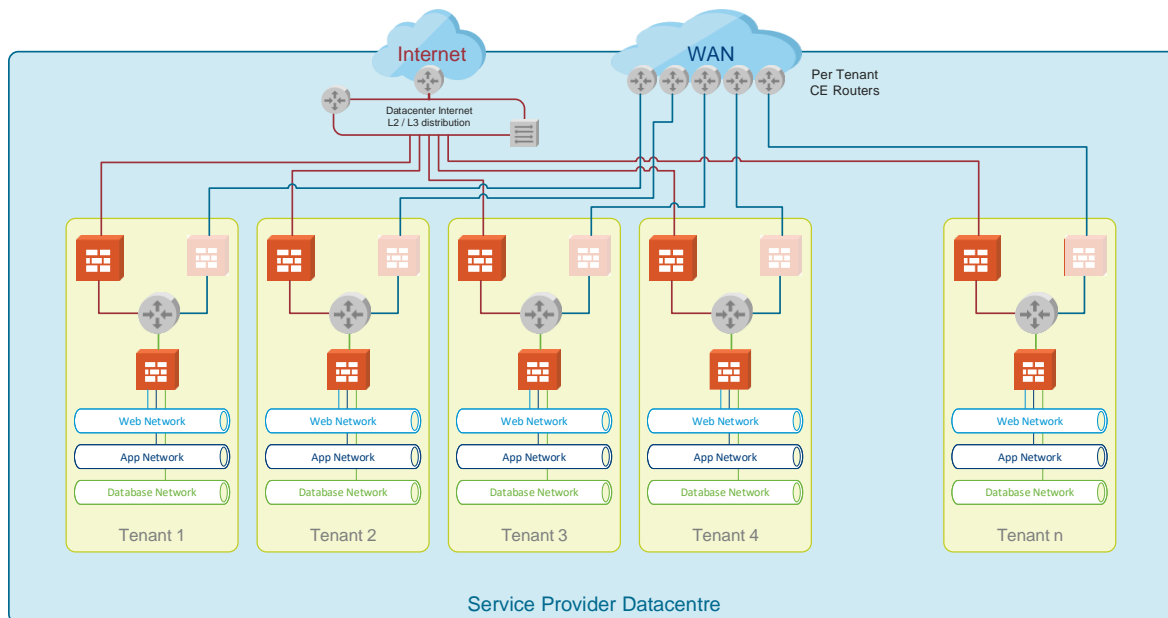
Each release of vCloud Director brings new features and functionality. While the current version does not support the full NSX feature set, recent releases of vCloud Director have each added further support for the consumption and management of additional NSX features. This trend is set to continue, so the reader is encouraged to check the most recent vCloud Director release notes to establish the level of support at the time of reading.



Multitenancy in a Cloud Service Provider

While Managed Service Providers are used to providing multitenancy within a single data center location, this is often created using discrete hardware for each customer's services, often with dedicated compute, storage, networking or security hardware with direct connectivity to dedicated, physical, WAN access CE routers. Similarly, the ESXi layer of each tenant is often managed from a dedicated VMware vCenter Server® instance. The following figure shows the basic tenant topology within a multitenant Managed Service Provider data center.

Figure 5. Example Managed Service Multitenant Data Center Topology



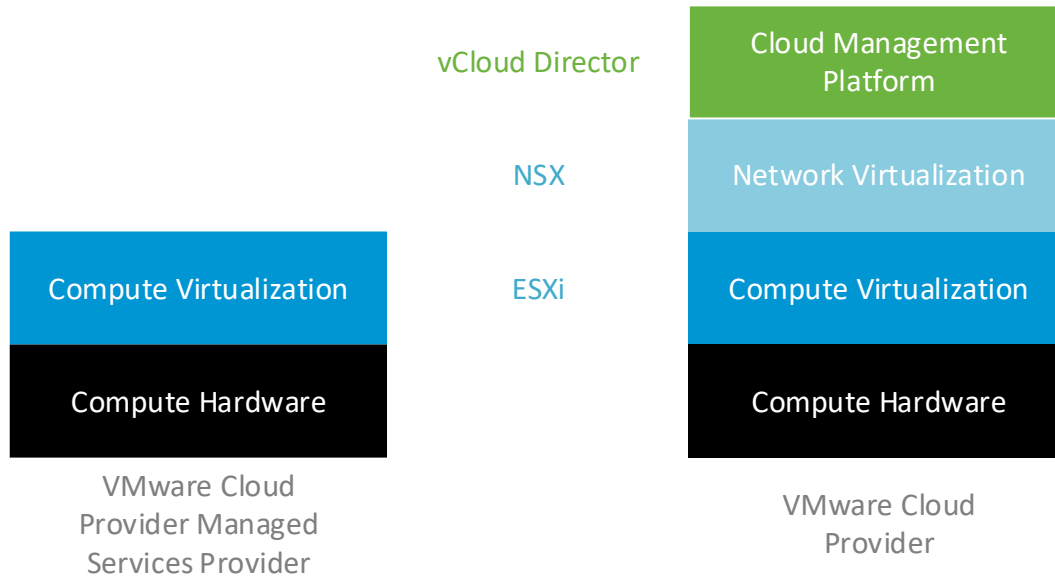
The introduction of cloud services allows a consolidation of that virtualization management layer into perhaps a single vCenter Server, managing compute hosts serving multiple customers. If the Cloud Service Provider is to offer rapid onboarding of new tenants, it is far quicker and more economical to share network hardware, too. In the past, sharing security hardware was possible, but required the provision of large, high capacity security appliances that could offer logical separation for each tenant, often in the form of virtual contexts or instances. However, the large capital outlay to provision such devices with capacity to scale to the number of tenants required, was a challenge for Cloud Service Provider business cases. VMware Cloud Service Providers now have the benefit of being able to provision dedicated security devices in the form of NSX Edge services gateways whenever they need to with the only proviso being that they capacity manage the underlying compute platform on which the services run, and the underlying network platform that connects the compute hosts together.

While the commercial benefits of sharing hardware are easy to see—increasing the utilization of hardware minimizes the amount of hardware that must be procured, deployed, and managed—it introduces the new challenge of providing logical tenant separation at all levels of the infrastructure stack. Creating separation at the vSphere layer might be straight forward, but extending that separation to the user interface or APIs is more complex. vCloud Director approaches this problem by abstracting the elements under its control away from the end user, instead providing them with a new graphical user interface for manual interactions and an API for machine interactions. It is important for the service provider architect to understand how the layers of a vCloud Director solution fit together and how actions at one layer are carried out in the other layers.



While the VMware software-defined data center (SDDC) model includes storage virtualization, the key layers that illustrate the move from Managed Services Provider to Cloud Services Provider and which are considered within this document are shown in the following figure.

Figure 6. Layered Software Transition from VMware Cloud Provider Program MSP to CSP



3.1 vCloud Director Multitenancy

In providing a multitenancy overlay, vCloud Director introduces some new terms for the constructs that describe that multitenancy model. The terms, which are used throughout the rest of this document are described briefly in the following table. For a more detailed explanation, see the documents linked in the Section 7, References.

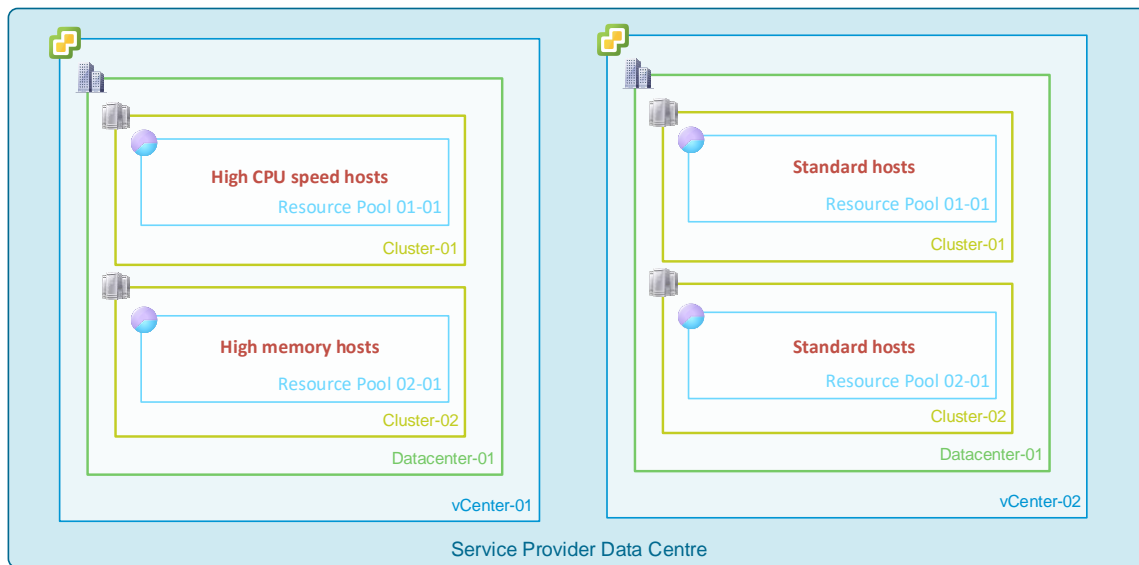
Table 1. VMware vCloud Director Tenancy Terms

Term	Description
Organization	An organization is a logical group of all users to whom resources will be presented.
Provider virtual data center	A provider VDC is a collection of vSphere resources (storage, CPU and memory) that vCloud Director can manage and use.
Organization virtual data center	An Organization VDC is a subset of a Provider VDC's resources that are available to an Organization.

The following figures illustrate these concepts.



Figure 7. Service Provider Data Center vSphere Clusters and Resource Pools

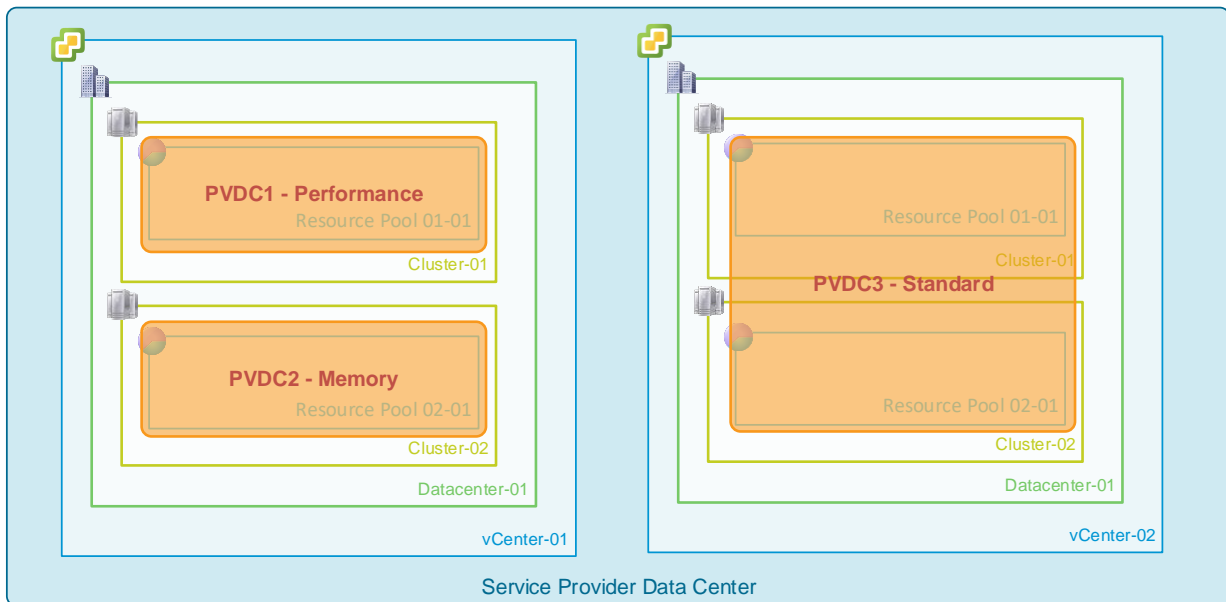


In Figure 7, the data center has two vCenter Server nodes, each managing two clusters. Under vCenter-01, the two clusters contain different types of hosts. Cluster-01 has hosts with higher CPU clock speeds for CPU intensive workloads, where Cluster-02 has hosts with extra memory allowing for hosts with higher memory demands. vCenter-02 on the other hand, contains two clusters with the same type of hosts, each with a balance of clock speed and memory designed for general workloads. Because it is not possible to place a resource pool anywhere other than across all hosts in a cluster, and vCloud Director places workloads within a resource pool, all hosts in a single resource pool, and therefore cluster, should be of the same type to ensure consistent performance.

Because a Provider VDC (PVDC) is connected to a single vCenter Server, the vCloud Director example in Figure 7 must have at least two vCenter Server nodes. However, to allow for the placement of workloads on either high CPU or high memory hosts, the resources under vCenter-01 should be split into two Provider VDCs, each mapped to a resource pool in one of the clusters. Because a Provider VDC can have multiple resource pools from the vCenter Server assigned to it, vCenter-02's resources could be presented as two separate Provider VDCs or, as shown in Figure 8, one PVDC with multiple resource pools.



Figure 8. Mapping Provider VDCs to vSphere Resources



The PVDC presents resources that can be consumed by the Organizations who subscribe to the Provider's Cloud service. To present the resources to those Organizations, vCloud Director uses an Organization VDC (Org VDC or simply OVDC) to represent a subset of a Provider VDC. A customer, represented by a vCloud Director Organization can access multiple Organization VDCs. In the example illustrated in Figure 8, a customer might require some high memory workloads and some standard ones, in which case they have two OVDCs, one in PVDC2 and one in PVDC3. A customer who only requires standard workloads might have a single OVDC, or might choose to have two to enforce different oversubscription ratios or workload placements for production and development workloads.

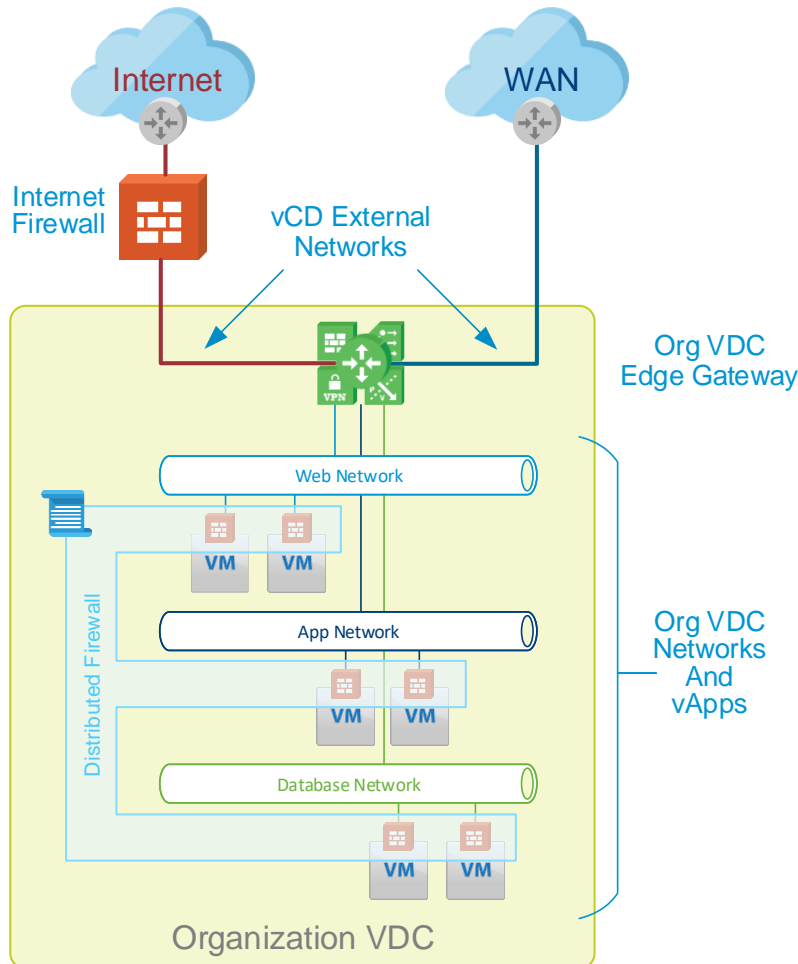
When accessing vCloud Director, an Organization's users with appropriate access privileges see all of their OVDCs listed.



3.2 Basic vCloud Director Tenant Topology

The following figure shows the vCloud Director Tenancy model overlaid onto a simple customer topology.

Figure 9. Basic vCloud Director Customer Topology

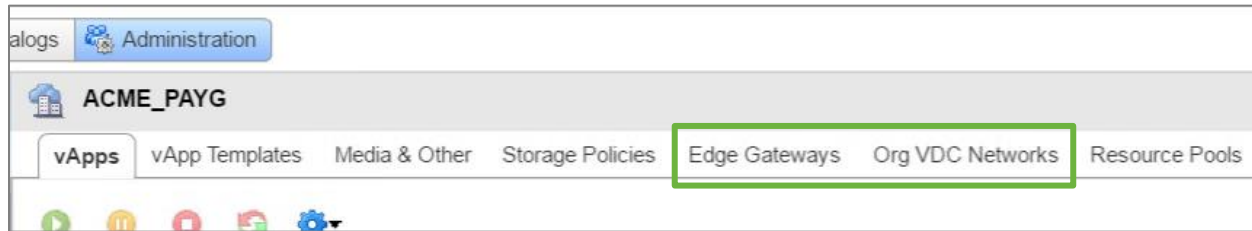


The Org VDC contains the customer's workloads which in vCloud Director, whether they are a single VM on a network, or multiple VMs on different networks, are known as vApps. Because vCloud Director does not manage all of the customer's resources, those in the physical data center outside must be managed by the service provider or, through a different customer-facing portal. Networks that connect Organization VDCs to external data center resources are described within vCloud Director as "external networks". They terminate on an Org VDC Edge Services Gateway to provide routed, Network Address Translated (NAT'ed) or directly connected access to and from the workloads inside the Org VDC. Networks that are confined to the vCloud Director environment are known as Org VDC Networks.

Both an Organization's Edge Services Gateways and Org VDC networks can be managed from the vCloud Director Org VDC management page as shown in the following figure.

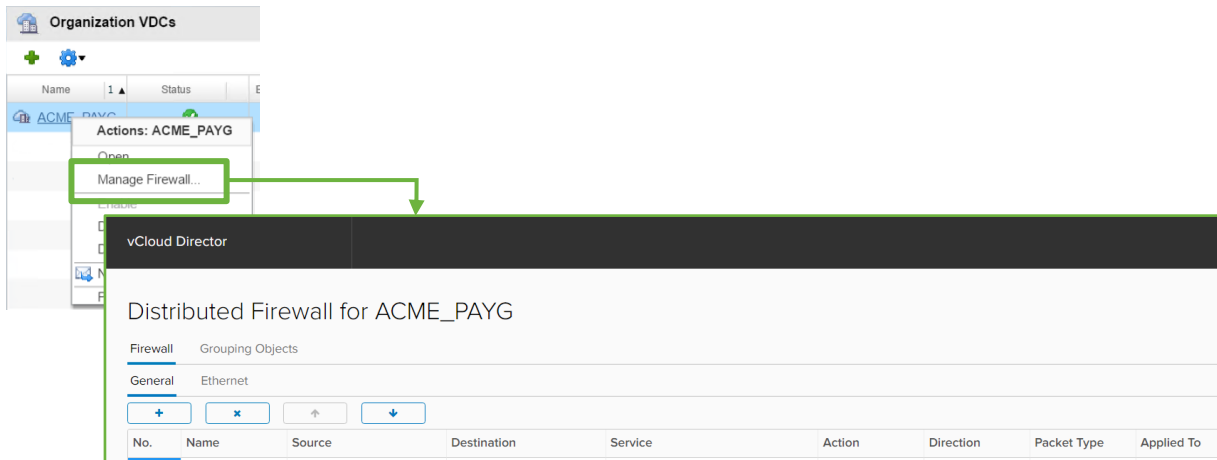


Figure 10. Managing Edge Gateways and Org VDC Networks in vCloud Director



In the current version of vCloud Director, the NSX Distributed Firewall uses the new HTML5 interface which is launched from the Organization VDC's **Actions** menu, as shown in the following figure.

Figure 11. Managing the NSX Distributed Firewall in vCloud Director



3.3 Multitenant Networking

In a VMware Managed Service Provider environment, networking is managed from a number of points. Physical network infrastructure is managed per device, from a vendor supplied central management platform, or from a bespoke automation capability. Virtual network infrastructure is managed from one or more vCenter Server nodes, each responsible for the connectivity between the physical network and the virtual workloads. In a VMware Cloud Service Provider environment, vCloud Director abstracts some of the virtual network management tasks and network management then falls into the layers described in the following section.

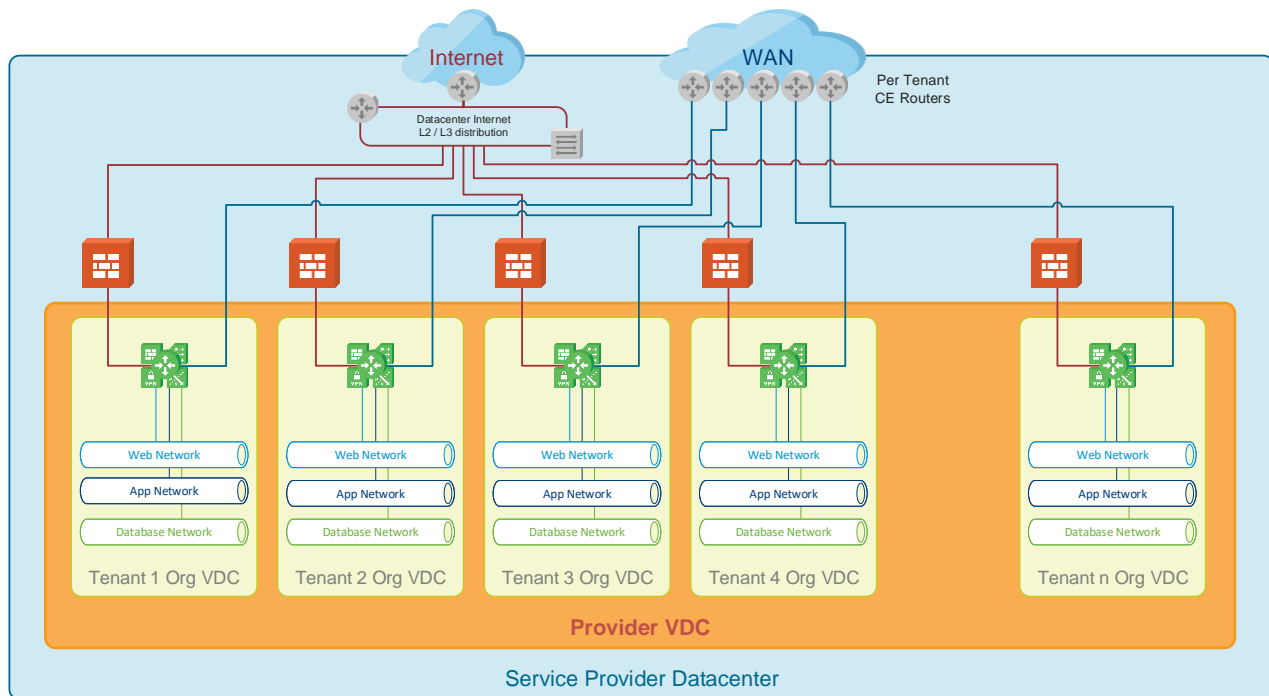
3.3.1 Network Layers in a Multitenant Cloud Platform

- **Data center and NSX underlay networking** – The layer of network configuration which remains the responsibility of the service provider. As well as management networks, this includes the NSX “Transport” network, which carries the VXLAN-encapsulated traffic between ESXi hosts, and the per-tenant networks which must be configured within the data center infrastructure when a new customer is onboarded.
- **vCloud Director networking** – The Org VDC networks which are created and managed entirely from within vCloud Director and use the preconfigured NSX Transport network for connectivity between hosts.
- **vCloud Director managed networking** – The external networks which are initially created in the relevant Provider VDC vCenter Server, but which are then “added” to vCloud Director and can subsequently be managed from the vCloud Director user interface or API.



These three types of networks are shown in the following figure. The diagram shows the per-tenant networks required to connect each customer's WAN access to their vCloud Director Organization VDC.

Figure 12. Example Cloud Service Provider Multitenant Data Center Topology



In this graphic, the networks from the internet distribution and WAN routers are managed within the data center network infrastructure, typically at customer onboarding. The “Web”, “App”, and “DB” networks in each tenant Org VDC are created and managed from vCloud Director either by the customer or the service provider. The networks from the internet firewalls (if provided) and WAN routers, once configured, appear in vCloud Director as external networks and are subsequently managed from the vCloud Director user interface. See Appendix A: Provisioning an External Network in vCloud Director for more details.

Data center external connectivity in a Cloud Service Provider environment follows the same models in the physical infrastructure as it would in a Managed Services Provider environment.

- **Per-tenant** networks are used where Layer 2 separation is required across shared data center infrastructure. For example in cases where there is overlapping customer addressing, or a need to manage traffic flows without resorting to Layer 3 routing. Customer access from their WAN to their vCloud Director environments, or from co-located services within the physical provider data center, are examples of per-tenant networking.
- **Shared** networks can be used when there is no risk of overlapping addresses (such as public internet access) and where Layer 3 routing can be used to steer traffic to the correct destination.

Shared networking between multiple customers within a single Layer 2 broadcast domain raises the risk of a network problem affecting multiple customers. To mitigate this, VMware Cloud Providers can choose a hybrid approach in which common networks (again, such as internet access) are terminated on high-performance Layer 3 devices which forward traffic to multiple, smaller downstream networks that offer separate broadcast domains, thereby reducing the effect one customer can have on others.



Networking Layers Examined

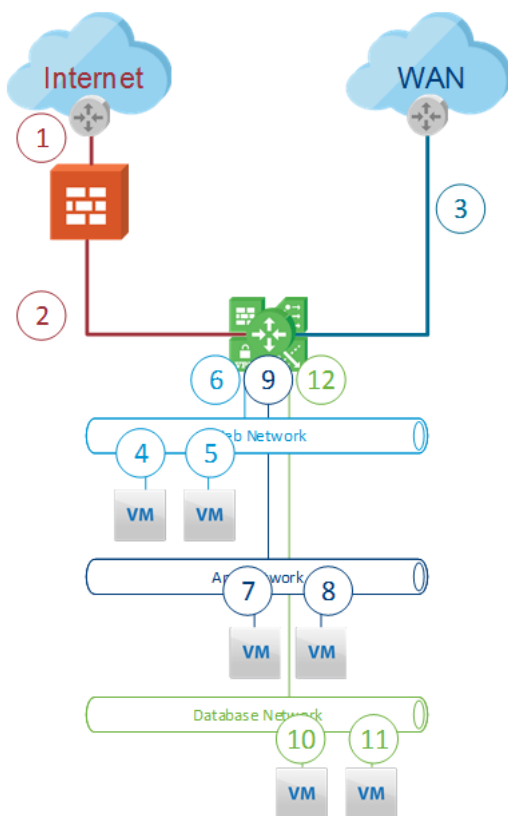
In a vCloud Director Cloud Service Provider, network configuration changes can be instigated and managed from within the vCloud Director user interface or API by either the customers themselves, or the VMware Cloud Provider on their customers' behalf. While some changes, such as IP address pool management or gateway addresses, remain within the vCloud Director configuration database, most impact the layers below. For example, in the case of a new Org VDC network, the request is passed from vCloud Director to the VMware NSX Manager™ connected to the vCenter Server that hosts the resource pool for the provider VDC in the Organization VDC where the new network is to be created. The NSX Manager configures the network within its internal state model, and then the underlying vSphere layer through the NSX Manager's link to its parent vCenter Server.

The following sections examine this layered model from the perspective of the previous single tenant topology as well as the multitenant view at both the vSphere and NSX configuration layers.

4.1 Tenant Networking

The basic Cloud Service Provider tenant topology in Figure 2 (and shown in the following figure for convenience) was constructed from twelve separate network segments.

Figure 13. Cloud Service Provider Tenant Networking

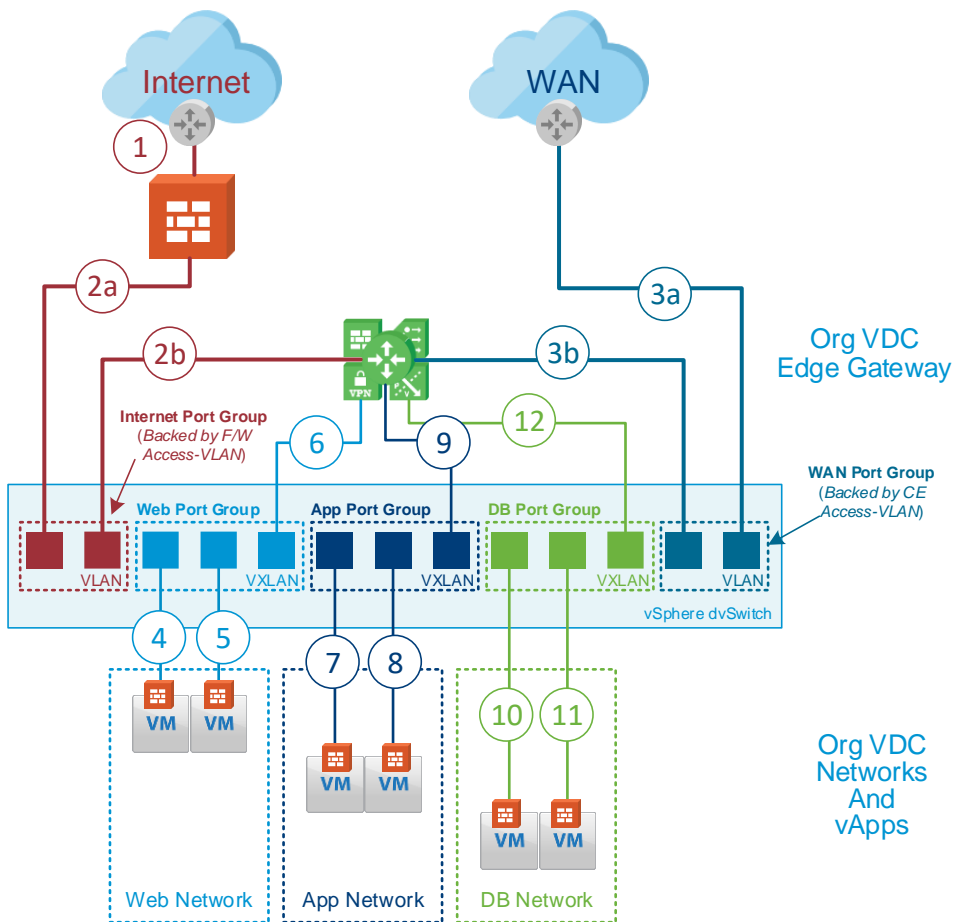


1. Internet access to internet firewall
2. Internet firewall to Edge Services Gateway
3. WAN access to Edge Services Gateway
4. Web VM1 to vSphere dvSwitch port
5. Web VM2 to vSphere dvSwitch port
6. Edge Services Gateway web interface to vSphere dvSwitch port
7. App VM1 to vSphere dvSwitch port
8. App VM2 to vSphere dvSwitch port
9. Edge Services Gateway app interface to vSphere dvSwitch port
10. Database VM1 to vSphere dvSwitch port
11. Database VM2 to vSphere dvSwitch port
12. Edge Services Gateway database interface to vSphere dvSwitch port



The following figure represents the same tenant topology as it appears within the NSX and vSphere layers.

Figure 14. vSphere Cloud Service Provider Tenant Networking



- Connection 1 is made in the physical data center network infrastructure, and does not reach the NSX or vSphere layers of the solution.
- Connection 2 is made up of two parts, a physical connection between the customer’s internet firewall’s “inside” interface, which is presented to a port in the VLAN-backed Internet Port Group in the vSphere dvSwitch, and the Edge Services Gateway’s Internet interface connection to a second port in the same port group.
- Connection 3 is also made up of two parts, a physical connection between the customer’s WAN router “LAN” interface, which is presented to a port in the VLAN-backed WAN Port Group in the vSphere dvSwitch, and the Edge Services Gateway’s WAN interface connection to a second port in the same port group.
- Connections 4 and 5 present the Web VMs to ports in the Web Port Group which, because the Web Network is a vCloud Director Org VDC Network, is created within NSX as a “virtual wire”, so appears in the dvSwitch as a VXLAN-backed port group.
- Connection 6 presents the Edge Services Gateway web interface to the Web Port Group.
- Connections 7 to 9 follow the same pattern as 4, 5, and 6 except for the App Network/Port Group.
- Connections 10 to 12 also follow the same pattern as 4, 5, and 6 but this time for the DB Port Group.

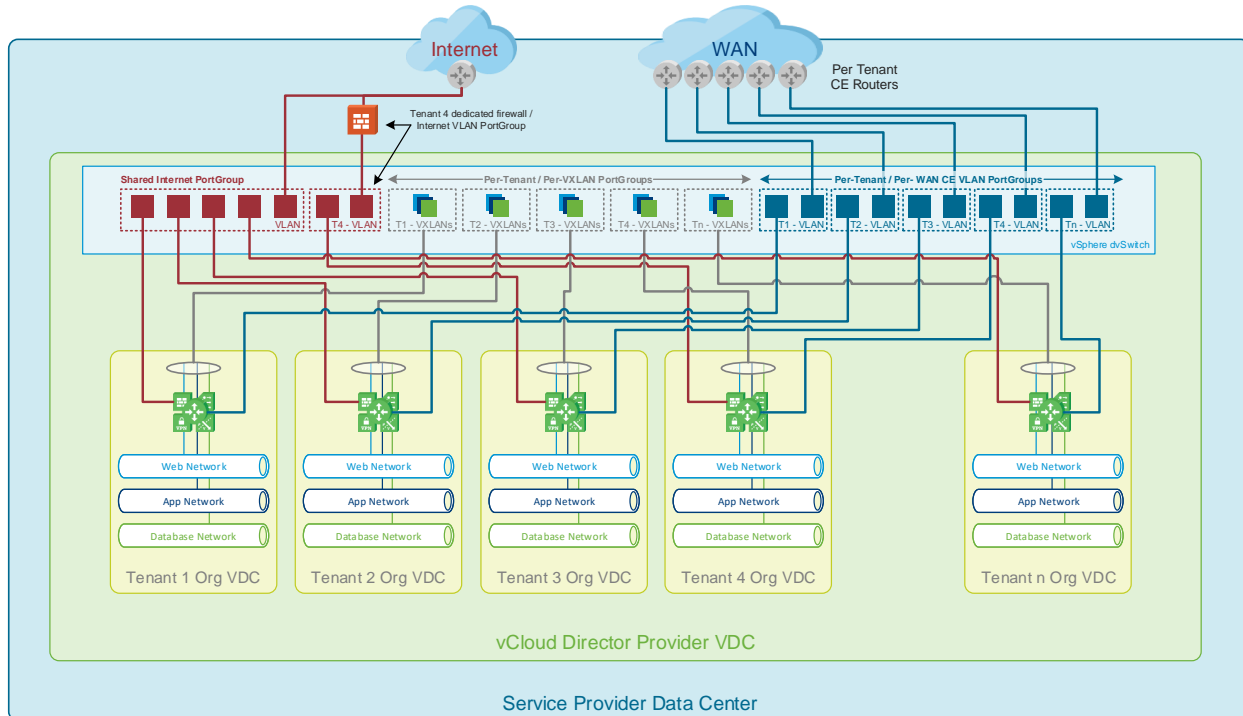


- The logical position of the NSX Distributed Firewall on each virtual machine interface is also shown to represent the point at which a DFW policy is applied to the traffic flow into or out of a vNIC on a virtual machine.

4.2 vCloud Director Multitenant Data Center Networking in vSphere

The detailed analysis the previous section illustrates the representation of a single tenant within the vSphere layer. The following represents the same analysis but applied to a multitenant vCloud Director data center.

Figure 15. vSphere Cloud Service Provider Multitenant Networking



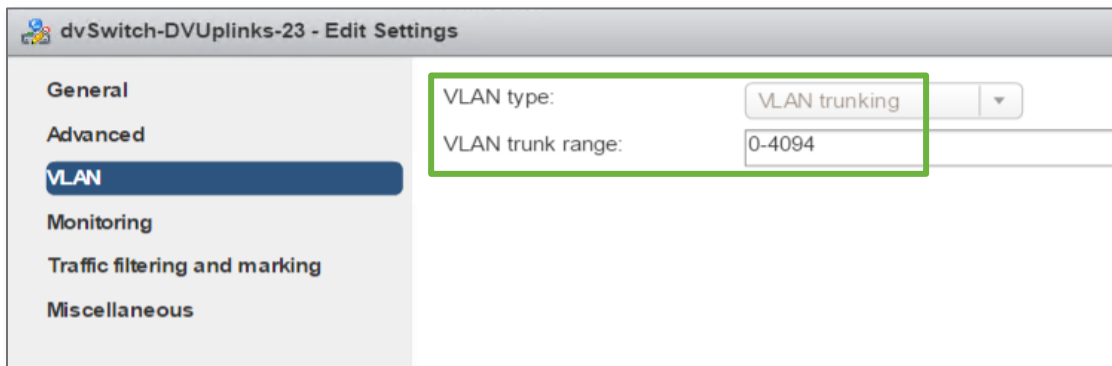
- The Internet access shown in this graphic is shared across all tenants who do not require a physical firewall. The shared internet access connection is presented as a port in a VLAN-backed port group in the vSphere dvSwitch. Each vCloud Director tenant has an external network connection between a port in this group and the internet interface of the Edge Services Gateway in their Org VDC.
- Tenant 4 has elected to retain a physical internet firewall. In their case, within the physical data center infrastructure, the shared internet access is presented to the “outside” of their firewall, and a separate VLAN connects the “inside” of the firewall to a separate VLAN-backed port group within the vSphere dvSwitch. An external network connects a second port in this port group to the internet interface of the Edge Services Gateway in their Org VDC.
- Each customer has a separate WAN CE router. Because the connection to their tenant environments could, therefore have overlapping addresses (from within their vCloud Director organization, or from their WAN), each must be separated through the data center and into the vCloud Director managed environment. This typically means that each tenant’s WAN connection is presented as a separate external network with a separate VLAN ID, and therefore requires a separate VLAN-backed port group in the vSphere dvSwitch to connect to the WAN interface of their respective Edge Services Gateways.



- While simplified for clarity in this graphic, in the same way as the Edge Services Gateway interfaces and virtual machine interfaces in Figure 14, each tenant's Org VDC networks appears as a VXLAN-backed port group within the dvSwitch, with ports for the Edge Services Gateway interface and any vApp virtual machines connected to that network.

While this example illustrates the separation of VLANs behind the per-tenant WAN access, in a service provider data center it is likely that at some point in the infrastructure, several VLANs of a similar type and security level will be “trunked” on a single link. When that is the case, the per-VLAN presentation between the vSphere networking layer (beneath the Provider VDC) and the data center infrastructure shown in Figure 15 is not necessary. In a similar way to a physical switch, a dvSwitch Uplink port can carry multiple VLANs encapsulated on the single connected link using IEEE 802.1q. To do this, when a dvSwitch Uplink Port Group is being created, its VLAN Type is set to “VLAN trunking” as shown in the following figure. The range of VLANs allowed on the trunk can also be configured here.

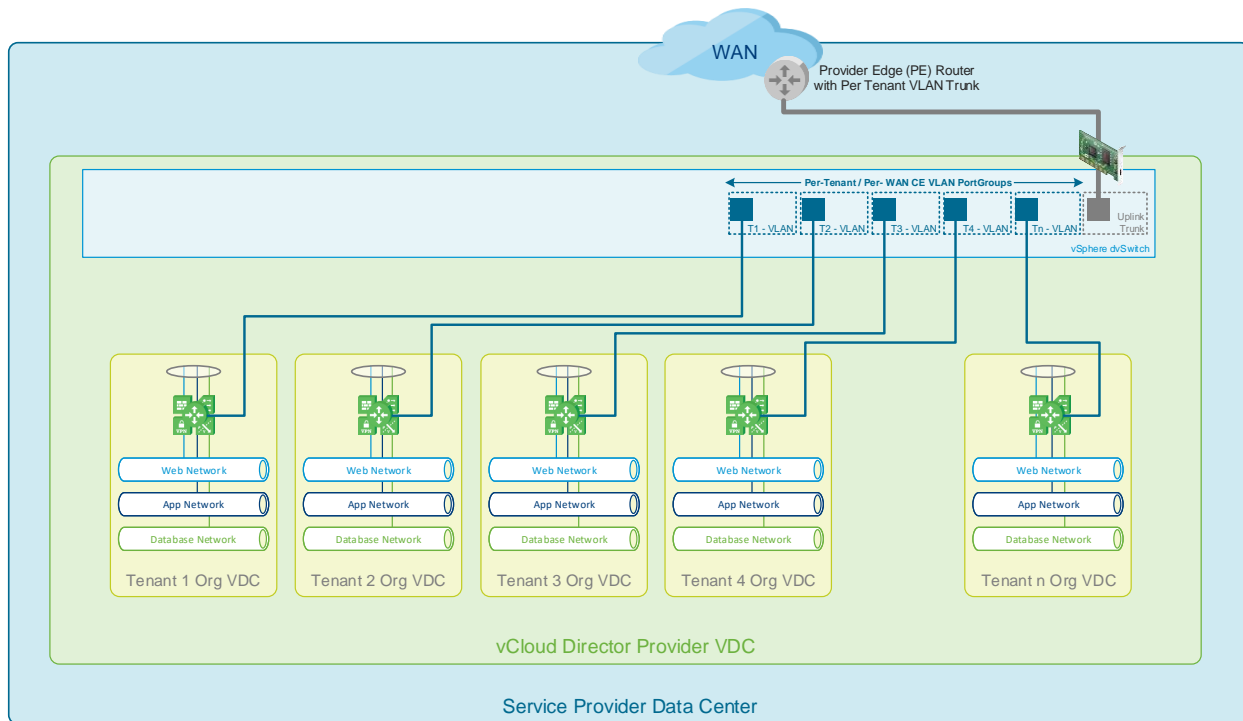
Figure 16. Configuring a dvSwitch Uplink Port Group VLAN Trunk



An example of where this can be applied is shown in the following figure. In this example, the service provider is presenting WAN access by connecting the MPLS Provider Edge (PE) router to the vCloud Director platform. Each customer's WAN VPN VRF is presented, by the PE router, to a sub-interface on a trunk connection, the other end of which connects to the customer's vCloud Director tenant, terminating on its Edge Services Gateway. Similarly, multiple customer WAN CE router connections can be terminated on VLAN-tagged access ports of an “aggregation” switch, whose uplink then delivers the trunked connections to each tenant's Edge Services Gateway.



Figure 17. Trunking Multiple External Networks to a vCloud Director Environment



This technique can equally be applied to multiple, separate internet VLANs where each is presented to a separate customer firewall in the data center, or, where multiple customers' co-located services are "trunked" into the vCloud Director environment over shared, high-bandwidth connections.

The number of dvSwitches used to deliver these VLAN-backed port groups depends upon a number of design considerations, one of which is the number of physical network adapters in the ESXi host. For hosts with only a single pair of adapters, all VLANs must be trunked over the same uplink port group (as shown in the range in Figure 16). In the case of a host with multiple adapters, or with adapters that can simulate multiple adapters to vSphere (such as some blade servers), dvSwitches can be created to separate, for example, management, internet and WAN traffic. For ease of configuration when this is possible, a range of VLAN IDs can be pre-allocated for "WAN uplinks" and, as each tenant is onboarded, they are allocated the next ID in the range. The Uplink Port Group is created with the specific range configured, and, the port group for each new tenant's "External Network" configured with their specific VLAN ID from the range.

4.3 Networking in a Multi-Cluster "Leaf-Spine" Infrastructure Topology

The illustrations in this section assume it is possible to connect (trunk) the same VLAN between the data center's external connectivity environments and the workload compute environments. In some infrastructure topologies, this is not possible because multiple routing hops might ordinarily exist between different parts of the infrastructure. Solutions employing discrete rack-mount compute nodes (as opposed to blade-based solutions with integrated networking) are often configured in a "leaf-spine" topology where VLANs are constrained within a rack, and inter-rack connectivity is achieved over a routed Layer 3 core.

When there is no ubiquitous Layer 2 coverage across the entire compute estate, the placement of tenant Edge Services Gateways onto compute resources with access to the required external connectivity becomes a challenge. vCloud Director offers some assistance because its placement engine deploys edge VMs to clusters based upon VLAN connectivity. Presenting the VLANs directly to a subset of racks (typically at least two for resiliency) provides that Edge Services Gateways will only be deployed to those clusters because of the VLAN availability there. However, vCloud Director will also be able to provision ordinary workloads to those same clusters because of the availability of tenant networking required for the



southbound connectivity from those same Edge Services Gateways. Because of this, clusters with external connectivity in this model are known as combined Edge/Compute clusters.

There are a number of alternative models for the placement of Edge Services Gateways on separate “Edge Clusters” without tenant workloads, each of which brings its own complexity and operational overhead. The principles outlined earlier in this section still apply, but must be applied on top of the Provider’s chosen Edge cluster option. The options are discussed in more detail in *Architecting a VMware vCloud Director Solution for VMware Cloud Providers (Section 6.3.2)*. See the Section 7, References for more information.



4.4 vCloud Director Multitenant Networking in NSX

For vCloud Director to take advantage of NSX, the vSphere environment which vCloud Director is to manage must be prepared for NSX first. The design considerations for an NSX underlay (VTEP Transport) network are the same whether the consumer is to be vCloud Director or, another CMP or orchestration tool, the specifics of which are beyond this document. However, during the design of a standalone NSX deployment, a number of choices are presented to the architect. When vCloud Director is managing the NSX environment, some of those decisions are made by vCloud Director, both at installation and during subsequent management of the environment.

The relationship between the network elements in vCloud Director and those of the underlying NSX and vSphere platforms, is examined in the following table.

Table 2. vCloud Director Tenancy and Network Elements in NSX and vSphere

vCloud Director	NSX or vSphere	Notes
Provider VDC	NSX Transport Zone	The configuration of a PVDC causes the creation of a Transport Zone within NSX. The new Transport Zone's Control Plane Mode is initially set to multicast and should be changed from the NSX user interface or API immediately if this is not the preferred setting.
External Network	vSphere Port Group	An external network is created in the data center infrastructure, vCenter Server networking, and vCloud Director Cloud resources dialogues. See Appendix A: Provisioning an External Network in vCloud Director for more details.
Edge Services Gateway	NSX Edge	An Edge Services Gateway created with the Org VDC configuration results in the creation of an NSX Edge. The initial name given to the Edge Services Gateway in vCloud Director is encoded into the NSX Edge name together with a UUID field to ensure unique names across tenants.
Org VDC Networks	NSX Logical Switch / vSphere Port Group	The creation of an Org VDC Network in vCloud Director results in the creation of an NSX logical switch with the initial Org VDC network name and a UUID in its name. This then results in a dvPortGroup being created in vSphere with the logical switch name and VXLAN network (segment) ID encoded into its name.



IP Address Management and Routing

5.1 Tenant Address Management

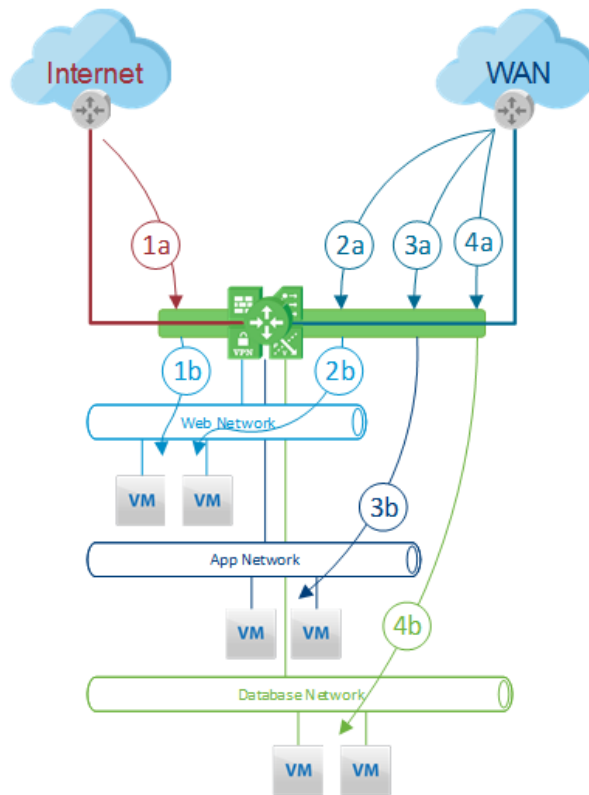
There are a number of external connectivity options to enable customers to reach the workloads on the networks within their tenant environment. Each requires that the IP addresses used externally to access the tenant workloads be known by the customers' connectivity method, WAN, VPN or other. The addresses that are used to access the solution can be provided by either the service provider or the customer. The details of the addresses used to access the tenant networks can be statically configured on routing devices used for access, but is often distributed using dynamic routing protocols. The configuration of static routing is largely dependent upon the specific routing hardware in question and is therefore beyond the scope of this document. The distribution of routing information using dynamic protocols is, in contrast, a key element of the efficiency gains which the Cloud Service Provider seeks to deliver.

The different methods of sourcing and managing addresses for tenant networks and their advertisement to customer access networks is discussed in the following section.

5.1.1 Service Provider Managed Addressing

Some service providers choose to manage the address space within their tenant environments and allocate appropriately sized ranges of addresses to their tenants. The advantage for the service provider is that they do not have to deal with multiple customers using the same "overlapping" addresses, which greatly simplifies access from the provider's management platforms into multiple tenant environments. However, a downside of this approach is that it is quite likely that the addresses allocated to a customer could be in use elsewhere within the customer's wider network. To prevent this duplication of addresses creating a problem, service providers enforce a layer of NAT at the boundary of the service. In cases such as this, the customer's Org VDC networks are addressed from the service provider's coordinated address space, typically using "private" addresses sourced from the ranges defined in RFC1918 (see Section 7, References). So that the customer can reach these addresses, ranges of mutually agreed upon, often public Internet, addresses are assigned and translated to the internal addresses used within the tenant networks.

The address translation (NAT) can be configured and carried out on an external, provider-managed device within the data center typically dedicated to each tenant, or, it can be carried out on the Edge Services Gateway and managed through vCloud Director. When NAT is carried out on inbound connections and the destination IP address is changed from the one on the boundary network to the real IP address of the target, NAT is more specifically known as Destination NAT or "DNAT". When NAT is carried out on the Edge Services Gateway, the external network that connects the tenant Edge Services Gateway to the Customer Edge router is allocated the subnet that contains the range of NAT addresses. The following figure illustrates this on the basic tenant topology used earlier.

**Figure 18. NAT on the Org VDC Edge Services Gateway****1. Internet NAT**

- a. The connection from the public internet terminates on an address on the internet external network assigned to a NAT on the Edge Services Gateway.
- b. The edge translates the destination address to that of a web server on the Web network.

2. WAN-to-Web Network NAT

- a. The connection from the WAN terminates on an address on the WAN access external network assigned to a NAT on the Edge Services Gateway.
- b. The edge translates the destination address to that of a web server on the Web network.

3. WAN-to-App Network NAT

- a. As per (2a).
- b. The edge translates the destination address to that of an app server on the App network.

4. WAN-to-Database Network NAT

- a. As per (2a).
- b. The edge translates the destination address to that of a database server on the Database network.

The interfaces of Edge Services Gateway and upstream devices are allocated addresses from the subnet assigned to the external network, and the remainder are made available to be used for NAT to/from addresses assigned to VMs within the customer's Org VDCs. This process is examined further in Section 5.4, External Network Address Sub-Allocation.

While this model simplifies the service provider's network configuration, it can create problems for the customer, because some applications are intolerant of NAT and either will not work, or will require additional steps to overcome the issues caused by NAT. Because the ranges allocated to the Org VDC networks are hidden behind the NAT addresses, the real addresses of VMs and VIPs do not need to be exchanged with the upstream WAN devices. Because the addresses that are used for NAT are from the ranges that are allocated to the networks that directly connect the WAN devices to the Edge Services Gateway, the WAN devices will learn the NAT ranges as "Connected" networks and can then distribute those addresses to their upstream connections as needed.

5.1.2 Bring Your Own IPs

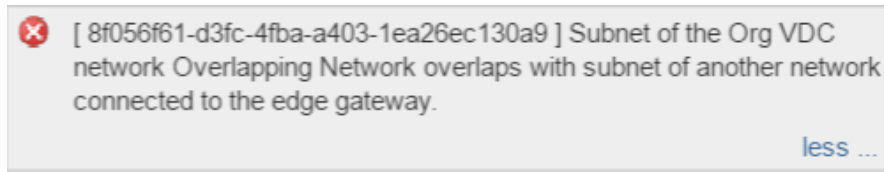
While forgoing the simplicity which service provider managed addresses afford the provider, allowing the customer to bring their own addressing can greatly simplify the customer's network configuration. The addresses assigned to workload VMs or load balancing virtual IPs (VIPs) are used directly by remote clients on the customer's WAN to connect to services within the Org VDC, removing the need for NAT. Some customers consider a third-party data center to be a less trusted environment and might insist on a layer of NAT between the workloads within the data center and their WAN environment. When this is the case, even if the customer provides both the Org VDC network addresses and the NAT range, their configuration is effectively identical to that described in Section 5.1.1, Service Provider Managed



Addressing with the exception that they might overlap with addresses used by another tenant or the provider’s management environment.

Workflows which, either at tenant onboarding, or as a “Day 2” action, create new Org VDC networks must assign new address subnets to those networks. To do so, the customer requires an IP addressing schema from which to allocate addresses for use within vCloud Director managed networks. vCloud Director provides an error message if a network is requested with an overlapping address range as shown in the following figure, but the customer or provider creating the network must have access to an appropriate subnet allocation to overcome the error message.

Figure 19. Error Message Caused by Provisioning an Overlapping Address Range



A number of techniques exist for allocating addresses from a larger schema. In some cases, each network subnet must be requested from the schema owner at the point of provisioning. In others, a larger “supernet” range of addresses are allocated by the customer schema owner for use within the provider environment. This allows delegation of management to either the service provider if they are providing a management service, or to the customer solution owner responsible for the workloads in the vCloud Director Org VDCs.

5.2 Customer Address Assignment

The assignment of addresses for consumption by workloads in the Org VDC networks “customer” (provided by either the customer or the provider) takes place in one of three main ways. Two of these involve vCloud Director configuring the address onto the VM during its creation, while the third allows the VM’s guest operating system (OS) to be built without an IP address and to request one at boot time. Each method is described in the following sections.

5.2.1 Static IP Pool Assignment

With Static IP pool assignment, vCloud Director chooses an address from the IP pool assigned to the Org VDC network. On supported guest OS, with VMware Tools™ installed and with “Guest Customization” enabled, the chosen address is then statically configured on the new VM during provisioning. This is the default option within vCloud Director and will result in an error if a VM provision is attempted without an IP pool with at least one spare address assigned to the network.

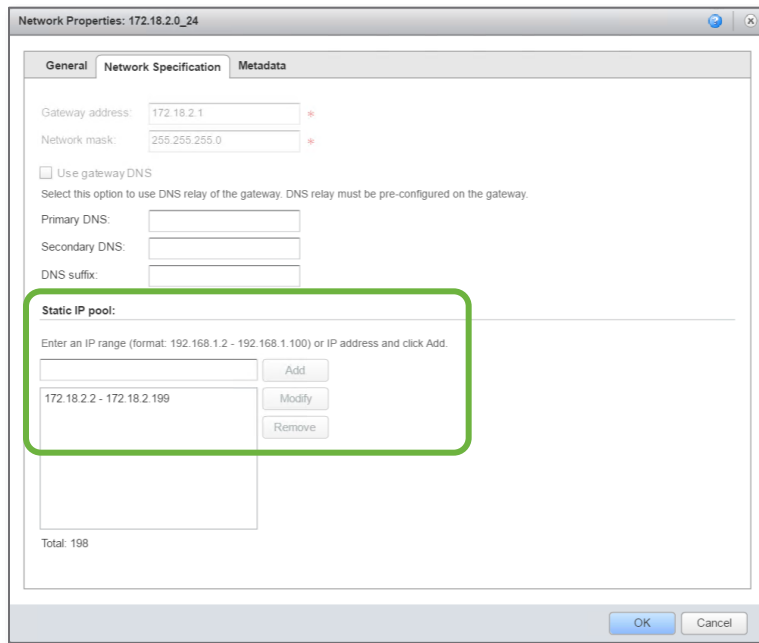
Figure 20. Static IP Pool Address Assignment

Virtual Machine	Computer Name	Networks
VM6	VM6 *	NIC 0 172.18.2.0_24 Static - IP Pool

This figure shows a new VM “VM6” being provisioned onto a network called “172.18.2.0_24”. Because the default “Static – IP Pool” method of address selection and assignment is in force, an address will be chosen from a pool assigned to the network (see the following figure).



Figure 21. Org VDC Network Static IP Pool Configuration



Unless the enhanced networking workflow is enabled on the Configure Networking page of the vApp provisioning dialogue (see the following figure), **Static IP pool** is the only option available on the Network Specification tab of the VM configuration shown in Figure 21.

Figure 22. Enabling the Enhanced Networking Workflow



5.2.2 Static – Manual Assignment

In many data center environments, the preferred option for servers is to permanently assign them a fixed IP address. Knowing that a server will retain its address after a reboot makes administration more straightforward. However, in some solutions, the specific address assigned to a VM carries additional significance. When this is the case, enabling the enhanced networking workflow and choosing **Static – Manual** (see the following figure) allows the user to manually select a specific address to be assigned to the VM being provisioned.

Figure 23. Static – Manual Address Assignment

Virtual Machin	Computer Nam	Primar	Network Adapter Type	Network	IP Assignment
VM6	VM6	NIC	E1000	172.18.3.0_24 N	Static - Manual 172.18.3.33



There are a number of points to be aware of concerning manual assignment of addresses, including the following:

- The manually entered address is validated against the network subnet and an error is raised if the address is outside of the network assigned to the network.
- A manually entered address is not validated against other address assignments already present on the network, but it is tracked. The **IP Allocations** dialog will show which VMs have been configured with the same addresses as long as they have been configured through vCloud Director, either through the UI, or API. It should also be noted however, that vCloud Director will not power on a VM with a duplicate IP address as long as it is tracking (or aware of) both allocations of the address.
- If the address assigned comes from within a range already assigned to an IP pool on the network, the allocation is recorded, and the utilization of network address space is updated to reflect the newly assigned address.
- If the address is allocated from within the subnet configured on a network but not from an IP pool on that network, the allocation is recorded but no utilization percentage is calculated for that network.
- After there are tracked addresses in use on a network, it is not possible to assign an IP pool that contains those addresses. It is, however, possible to create multiple pools “around” and “between” them, but utilization percentages are only calculated for the IP pools, not the entire subnet.

Manual assignment of IP addresses to VMs within vCloud Director allows complete flexibility, but makes the assumption that a user (or API call) that chooses manual assignment has a good reason to do so and understands the consequences. With guest customization enabled, the VM will be configured with the entered IP address which, if incorrect, could cause service issues. Manual assignment of IP addresses is useful when for example, the VM must have its IP address set without the aid of guest customization, but the administrator wants to track the address used within vCloud Director. A static manual assignment provides that the vCloud Director database is updated even if the VM address must then be set directly through the guest OS.

5.2.3 DHCP Assignment

While vCloud Director can manage and track the assignment of addresses to VMs, and can track manual allocations which it requested, it is also possible to manage IP addressing outside of vCloud Director. Selecting DHCP assignment during the VM creation (see the following figure) causes vCloud Director to configure the supported guest OS to use DHCP to acquire an IP address during its boot sequence rather than during its initial configuration.

Figure 24. DHCP Address Assignment

Virtual Machin	Computer Nam	Primar	Network Adapter Type	Network	IP Assignment
VM6	VM6	<input checked="" type="radio"/> NIC	E1000	172.18.3.0_24 N	DHCP

When a VM is configured to use DHCP to acquire an IP address, it requires either a DHCP server on the network to which it is connected, or a DHCP “helper” or “proxy” on that network that can forward the VM’s request for an address assignment to an authority on a remote network. vCloud Director cannot track address utilization when DHCP is used, but it is possible to assign part of an Org VDC network’s address space to an IP pool (which would be tracked) and part to a DHCP scope (which would not be tracked). This can be useful in situations where for example, a portion of an application stack is relatively static and the rest is dynamically scaled as needed. The static portion would utilize an IP pool and because the specific addresses assigned to temporary workloads are not significant, they would use DHCP.

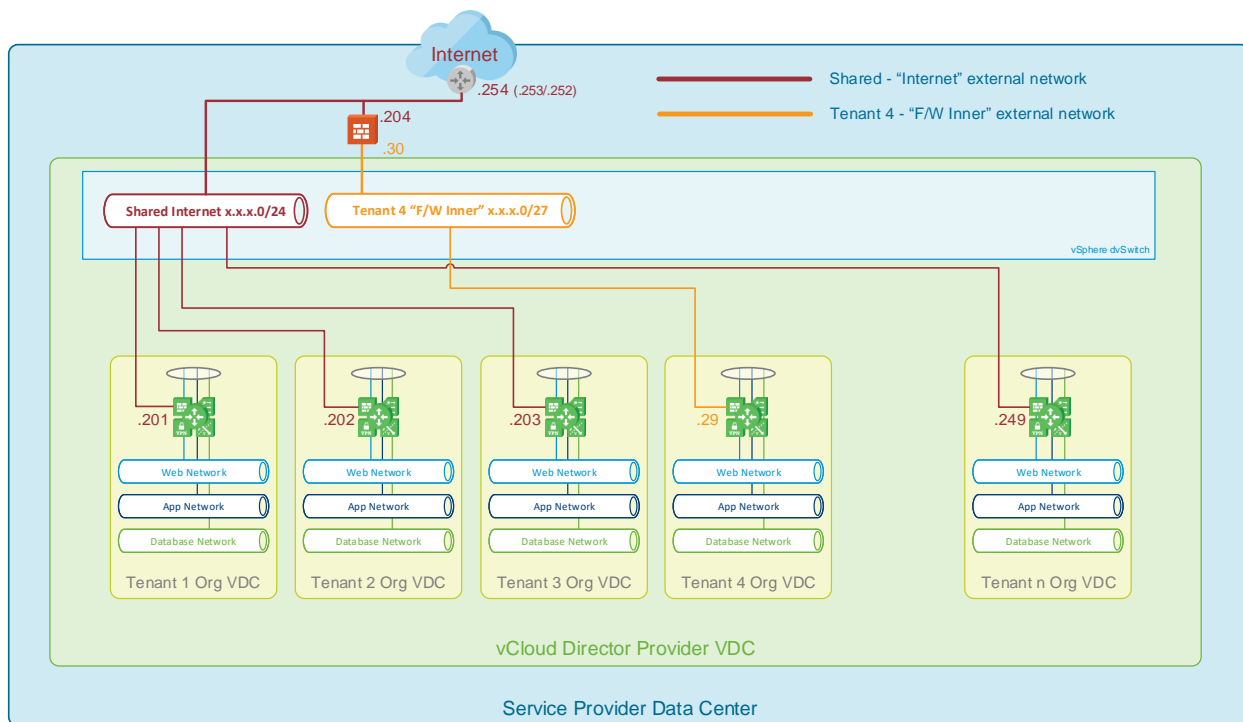


5.3 Internet Address Management

Internal, “private” addresses managed by the service provider or the customer which, if managed sensibly, have sufficient capacity so that all networks can be assigned subnets which allow room for expansion. This is not the case for registered “public” internet addresses which are in short supply. Internet access is presented to an Org VDC as an external network. Depending upon the provider data center topology, there might be one or more network hops between the vCloud Director environment and the local internet access or aggregation device. Because the precise topology is outside the scope of this document, only the final routing device is considered here.

The following figure shows the internet access elements of the data center. The red network shows the shared internet access network connected to a next-hop upstream device. As before, Tenant 4 has an external physical firewall, so while the other Tenants’ Edge Services Gateways have a connection to the internet external network, Tenant 4’s Edge Services Gateway is connected to a separate vCloud Director external network with the “Inside” interface of the physical firewall as the next hop.

Figure 25. Internet vCloud Director External Network



5.3.1 Shared Multitenant External Network

In the previous example, the Internet external network is assigned an entire class C, /24 subnet of which only the last octet is shown. The next-hop gateway uses the highest address in the range (.254) and, because a physical device can run an availability protocol across redundant hardware, additional addresses (.253 and .252 in the example). Tenant Edge Services Gateways are directly connected to the network and have their interface addresses allocated from .201 upwards. This allows approximately 50 tenants to be allocated an Edge Services Gateway address from the remainder of higher end addresses on the subnet. With 50 tenants and the next-hop gateway allocated from the addresses in the /24 subnet, the “red” network still has all the addresses from .1 to .200 available for customers to use for services they want to be accessible from the Internet. These addresses can be assigned to customers in ranges of any given size using a method known as sub-allocation. See Section 5.4, External Network Address Sub-Allocation for further information on the sub-allocation process. Each tenant must also have the interface address of their Edge Services Gateway sub-allocated to itself, so that it too can be used for tenant



services. See Section 5.3.3, Outbound Internet Access for the use of the interface address for outbound Internet access.

5.3.2 Dedicated Single-Tenant External Network

Tenant 4 is consuming the .204 address on the Internet network, but is doing so on a physical device outside of vCloud Director. In this example, a single address has been assumed but a high availability firewall “pair” can in practice require two or three addresses. Tenant 4’s Edge Services Gateway requires a connection to an external network, but cannot use the “red” Internet network, so a second network and associated address subnet is required for that. Each network loses the highest and lowest addresses (network and broadcast addresses) as well as those required for the two connected device interfaces.

On a /24 network, 2 addresses out of 256 is a relatively small proportion, but if many tenants require “yellow” networks and each is much smaller in address subnet size, the overhead of losing four addresses from each can become a consideration. Despite that, the yellow network will still require an address subnet assigned to it. Certain network devices and protocols such as those used in VPNs can either be intolerant of, or complicated by, the presence of NAT, so the yellow network should, like the red network, be assigned a subnet of public internet addresses, one which is large enough for the customer’s requirements but not so large as to be wasteful. This is a particular challenge for service providers with limited public IPv4 addresses. Connecting tenants to a shared network employing address sub-allocation is a more efficient way to provide addresses to connected tenant Edge Services Gateways.

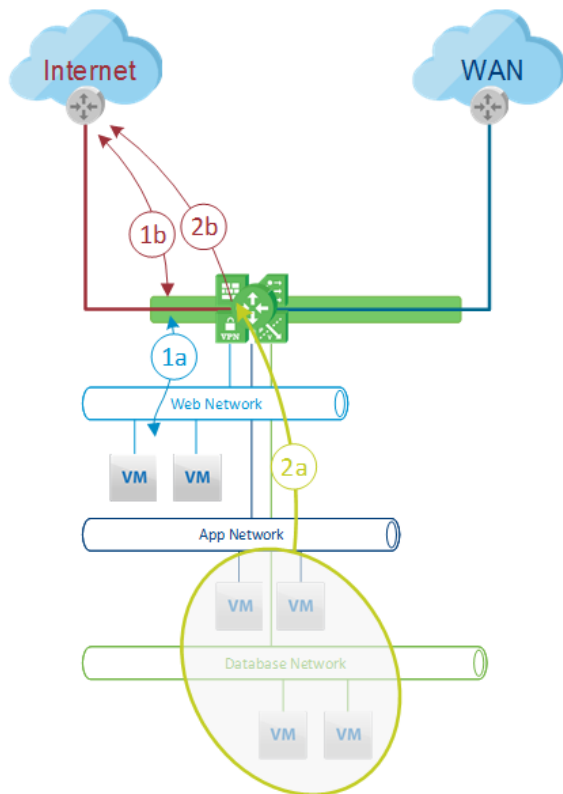
In the case of both Tenant 4 and the directly connected tenants, the addresses on their external Internet networks, after those allocated to the Org VDC Edge Services Gateway, can then be allocated to source or destination NATs or load balancer virtual server VIPs.

5.3.3 Outbound Internet Access

For tenant workloads to access the public internet, their private, internal addresses must be translated into public addresses that can be used on the Internet. Unlike the Destination NAT in Section 5.1.1, Service Provider Managed Addressing, when the outbound connection to the internet needs its source IP address changing, this NAT process is more accurately known as Source NAT (SNAT). A Source NAT is a 1:1 mapping between the IP address used by the device connecting to the internet and a suitable internet address assigned to that device. This 1:1 mapping is typically necessary when the device in question requires unsolicited inbound connections to be able to reach it, such as a web server, or email transfer host.

However, often one or more devices in a solution might require access to the internet but do not require unsolicited inbound access from the internet. These can be application servers which connect to a source of data or updates on the Internet, but should not receive inbound connections. For cases such as this, it is possible to “hide” many devices behind a single internet IP address. Each device’s outbound connection has its source IP address changed, often to the interface address of the device carrying out the translation, and the connection’s source TCP or UDP “port” number changed to a randomly selected number whose value is tracked in a connection table to allow the device to deliver the inbound half of the connection back to the correct, originating host. As both the source IP address and next layer protocol port address are both changed, this version of NAT is sometimes known as Port NAT (PNAT), or more commonly, Port Address Translation (PAT).

The following figure revisits the NAT diagram from Figure 18, but this time looking at outbound connections.

**Figure 26. Outbound Address Translation (SNA/PAT)****1. Outbound (Source) NAT**

- a. Originating connections arrive on the edge with their original (internal) source address targeting a destination on the Internet.
- b. Outbound connections from devices which also have dedicated inbound (destination) NAT addresses have their source address translated to the one used in those inbound connections.

2. Outbound Port Address Translation (PAT)

- a. Connections from devices which do not require inbound access arrive on the edge with their original (internal) source address targeting a destination on the Internet.
- b. Outbound connections from devices without a dedicated inbound (destination) NAT all have their source IP address translated to a shared internet address, often the interface address of the device carrying out the translation, and their source TCP/UDP port translated to a random, high numbered port.

As noted in Section 5.3.1, Shared Multitenant External Network, the Edge Services Gateway must have its interface address sub-allocated to itself before the address can be used for outbound PAT.

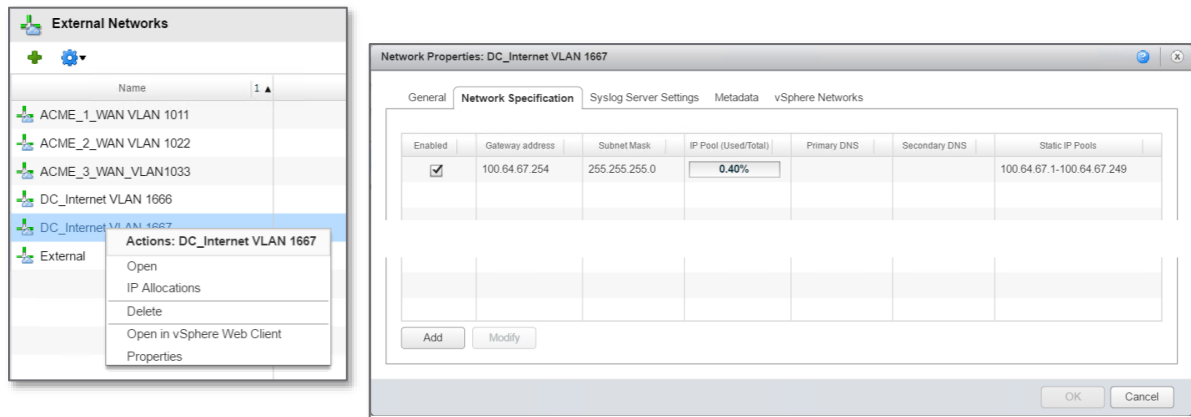
5.4 External Network Address Sub-Allocation

In a Managed Service Provider environment, the IP address space on a shared network such as the “Internet” network in Figure 25, is under the control of the service provider. If the tenant connections to that network were for example per-tenant physical firewalls under the management of the provider, the consumption of addresses on the Internet network would also, therefore, remain under the control of the provider. Allocating an internet address to a particular tenant would be accomplished by configuring NAT on the appropriate firewall, which would then respond to IP packets on the internet network destined for the address in the NAT statement.

If, as may be the case in a Cloud Service Provider environment, the customer could configure their own firewall NATs or load balancing VIPs, they could, without some form of control, consume as many of the internet addresses as they wanted, and conceivably, configure an address which is already in use within another tenant on the network, causing service disruption. To manage this scenario, vCloud Director introduces the concept of address sub-allocation. Before an Edge Services Gateway can consume addresses on an external network (other than its assigned interface address), a range of addresses must be sub-allocated to it. This is a two-stage process. First, the assignment of the external network’s subnet is made in the **External Networks** properties dialog as shown in the following figure.

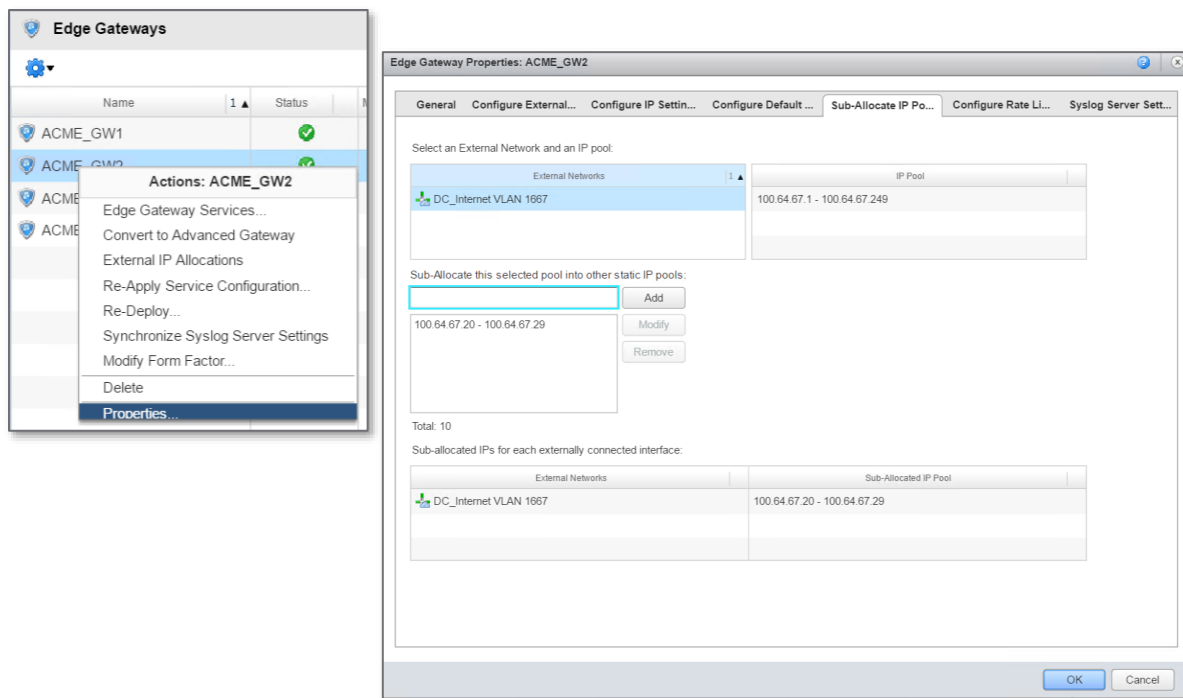


Figure 27. External Network IP Address Assignment



Then, after the external network has at least one IP pool assigned, the pool can be sub-allocated to connected Edge Services Gateways. The sub-allocation is carried out in the Edge Services Gateway configuration as shown in the following figure.

Figure 28. Sub-Allocation of External Network Addresses



In this example, the Edge Services Gateway “ACME_GW2” has been assigned the range of addresses from 100.64.67.20 to 100.64.67.29. These addresses are on a connected network, so do not need to be specifically routed to the Edge Services Gateway and therefore do not need to fall on subnet boundaries. Should the tenant require more addresses in the future, they could be added with a subsequent allocation from the IP pool. Should the parent IP Pool become exhausted, it is possible to add another IP pool to the external network from which further sub-allocations can be made, but care should be taken when doing so, because the upstream, gateway device will need to be configured with a secondary interface address, and default gateway and routing configurations become more complex to configure and troubleshoot.



5.5 Routing in a Multitenant Service Provider Environment

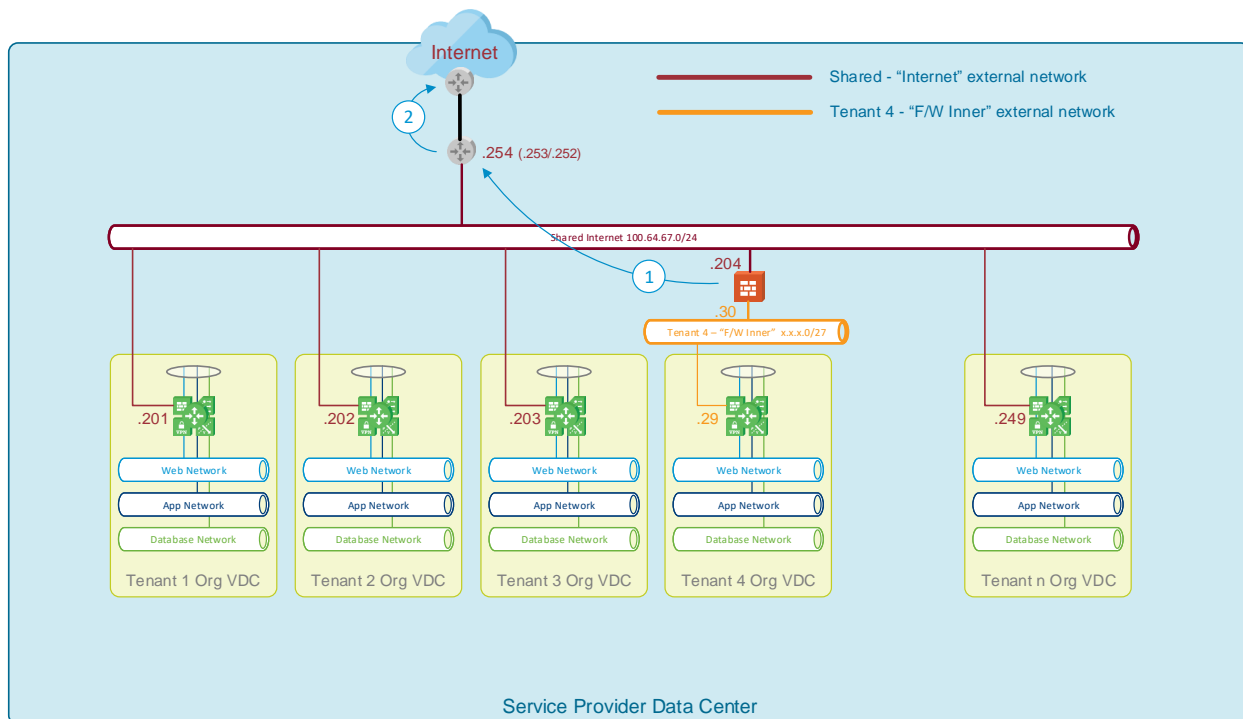
As vCloud Director allows (service provider permitting) customers to create new networks within their tenant Org VDCs, or set up WAN and internet external facing services, there is a requirement to make the addresses in question reachable from external networks. Although it is possible to manage this using static routing, configured on the routing devices within the solution, dynamic routing protocols provide a mechanism to automate the distribution of routing information without the need for manual intervention or complex, bespoke orchestration of routing configuration.

Tenant traffic routing within a service provider environment generally falls into two categories, public internet, and private, per-tenant. Each of these is examined in more detail in the following sections.

5.5.1 Provider Data Center Internet Routing

Using the model from Section 5.3, Internet Address Management, there are two types of addresses which need to be advertised from the vCloud Director environment to the upstream data center Internet access, and both are subnets allocated to the “Internet” external networks in the earlier section. The “red” shared internet network is directly connected to the upstream router and so is automatically present in that router’s routing table.

Figure 29. Provider Internet Routing



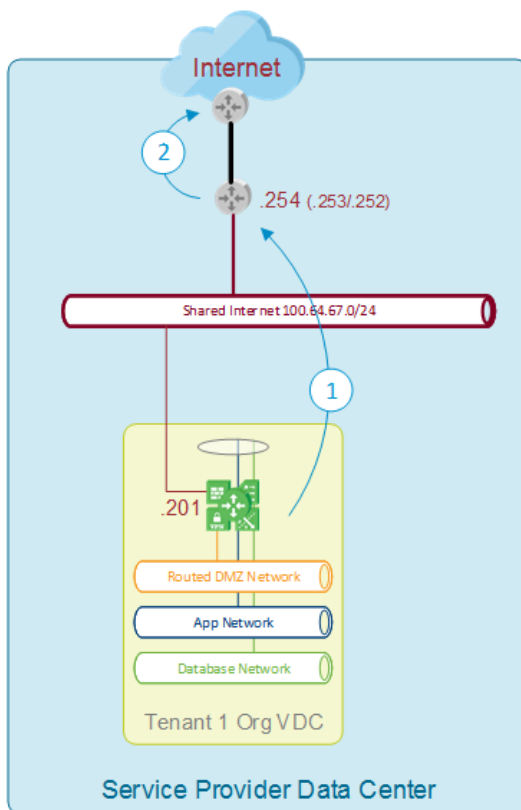
The first routing exchange in this figure is the advertisement of the Tenant 4 F/W Inner subnet to the next-hop gateway of the Shared Internet external network. This is necessary because the F/W Inner network is not directly connected to the Shared Internet network and would otherwise be unreachable. The choice of routing protocol used between the dedicated firewall and the upstream router depends upon a number of factors, but typically follow the service provider’s established standard. Some service providers do not run dynamic routing protocols on physical firewalls as a security precaution and, were that to be the case here, a static route would need to be configured on the Shared Internet router directing traffic to the F/W Inner subnet through the “.204” address on the connected, shared network. Similarly, the dedicated firewall would need its default gateway setting to the “.254” address of the upstream router.



The second routing exchange would be the advertisement of the Shared Internet subnet (in this example 100.64.67.0/24) to the next-hop Internet router within, or upstream of the provider data center. This would usually be carried out using Border Gateway Protocol (BGP), and the relative Autonomous Systems (AS) within which the two routers sit would determine whether this would be Internal BGP (iBGP) or External BGP (eBGP). After the Shared Internet router establishes a peering relationship with its upstream neighbor (2) it will advertise reachability of the connected Shared Internet subnet and the F/W Inner subnet, which it learned through the dynamic or static routing process in (1).

While not explicitly discussed earlier, it is possible for a customer to use registered internet addressing on an Org VDC network. This might be a range allocated by their service provider, or a Provider Independent Address (PIA) range which the customer already owns. To illustrate this, the usual web network in the Tenant 1 Org VDC in the following figure has been replaced with a routed “DMZ” network. Although not shown, the network would be addresses from a range of public internet addresses, with the Edge Services Gateway interface consuming one address, and the remaining useable addresses allocated to an IP pool for assignment to vApps connected to that network.

Figure 30. Routed DMZ Org VDC Network



The Routed DMZ network, in a similar way to the Tenant 4 F/W Inner network in Figure 28, is not reachable from the internet without the subnet assigned to it being advertised into the upstream internet connection. This is accomplished in two parts:

First, the Shared Internet router needs to learn that the routed DMZ network is reachable via the Tenant 1 Edge Services Gateway's “.201” address. This can be accomplished through the configuration of a static route on the internet router or, through a dynamic routing protocol such as BGP used in the earlier examples.

After the route is present in the Shared Internet router routing table, it can be redistributed to the upstream router over the BGP neighbor relationship already established.

To mitigate potential security issues, take care whenever the service provider learns routes from customer-managed devices that only appropriate routes are accepted and redistributed.

5.5.2 Provider Data Center Per-Tenant Routing

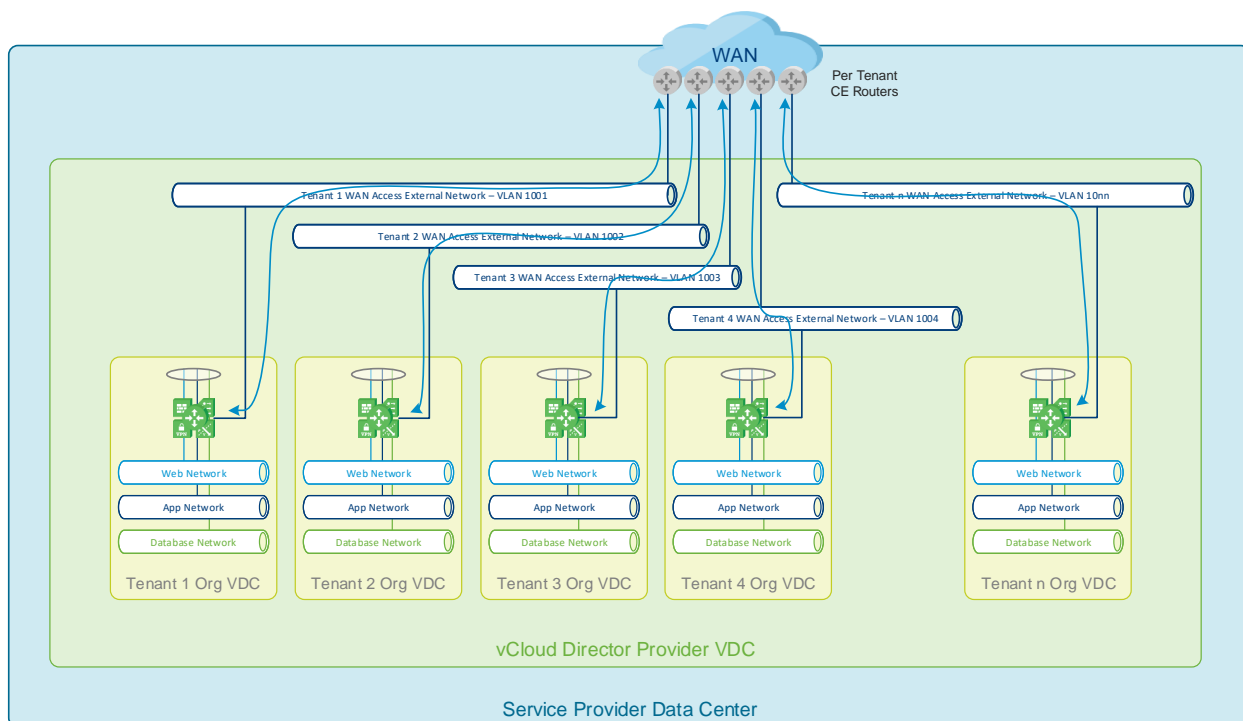
As noted in Section 5, IP Address Management and Routing, addresses used within tenant Org VDCs can be managed by either the provider or the customer. In the majority of deployments, whoever manages the addresses will do so from the “private” address ranges provided by RFC1918. Even if their usage can be controlled within the provider environment, it is likely that within the customers' WAN environments, addresses will overlap between customers. For this reason, even when the provider makes sure that addresses within the data center are unique, the need to uniquely route traffic to and from non-unique WAN addresses means that separate routing tables must be maintained for each tenant.



One case where this is not required is when the provider uses registered, public internet addresses in the external network NAT layer, which hides tenant Org VDC network addresses completely and only provides external access through the public Internet. In this case, both the tenant (NAT) addresses and the remote connection addresses are globally unique, and a single routing and forwarding table can be shared across all customers using the same mechanisms described in the previous section.

Overlapping customer addresses are separated by the Org VDC networks whose VXLAN backing creates the same Layer 2 separation as traditional VLAN-backed networks. The only places at which the overlapping addresses could therefore clash is if the separated networks were then connected to shared routing devices. In the examples used throughout this document, each tenant's Org VDC has a dedicated Edge Services Gateway, and is connected to the respective customer WAN over a discrete vCloud Director external network. This VLAN-backed network typically terminates on a dedicated, per-tenant WAN CE router or a shared multi-tenant PE router in which the each VLAN is internally mapped to per-tenant VRF (as described in Section 4.2, vCloud Director Multitenant Data Center Networking in vSphere). The Edge Services Gateway, CE router, or PE VRF maintain independent routing tables, allowing each customer to use identical addresses within their tenant Organizations without affecting other tenants.

Figure 31. Per-Tenant WAN Router Peering



This figure shows the peering between each tenant Org VDC Edge Gateway and that customer's WAN access router. The Edge Services Gateway will advertise the Org VDC networks to which it is directly connected, to the WAN access router, and from there to the rest of the customer's WAN. Any new Org VDC networks that are created by the service provider or the customer and which are connected to their respective Edge Services Gateway will also be advertised to their WAN. Because the Edge Services Gateway will typically default route towards the internet, it must learn the address ranges in use on the WAN to reach those destinations. This is carried out over the same routing protocol as that used to advertise the Org VDC networks to the WAN, except in the opposite direction.

The Edge Services Gateway supports both OSPF and BGP routing protocols, either of which can be used to peer with a WAN access router.



5.6 IPv6 Considerations

The current version of vCloud Director cannot manage IPv6 addresses for Tenant devices or networking. Some elements of a customer solution such as the virtual machine guest OS, and the underlying NSX components might offer IPv6 support, but will need to be managed outside of the control of vCloud Director itself. Should customers require IPv6 within their tenant organization, their specific requirements should be assessed to establish if it makes sense for the service provider to accommodate them with a managed service until such time as versions of vCloud Director are released with native support.

Commercial Considerations

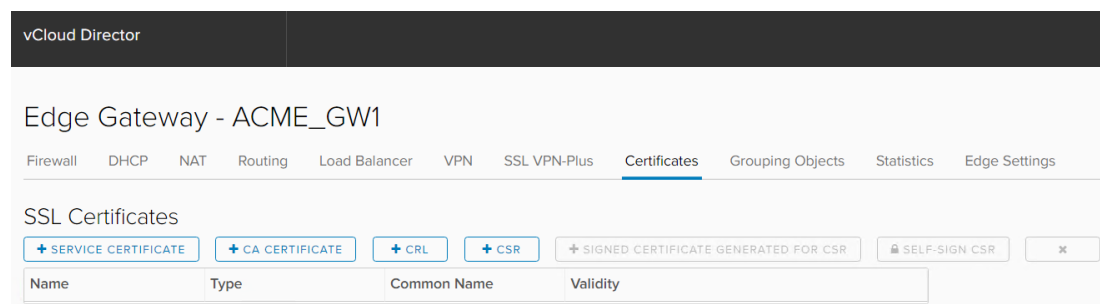
6.1 Managed Service or Self Service

VMware Managed Service Providers have been responsible for much of the configuration and ongoing management of their customers' solutions. The opportunities for value-add services throughout the lifecycle of a solution have allowed providers to differentiate their offerings based on their areas of expertise. Cloud services with higher levels of automation and standardization can, in some cases, restrict opportunities to differentiate.

Using vCloud Director to offer customers direct access to manage their tenant configuration does not preclude a VMware Cloud Provider from offering additional services to provide differentiation. While it is out of scope for this document, the Role-Based Access Control (RBAC) within vCloud Director allows the granular control of access to many of the capabilities within the vCloud Director User Interface². However, as noted earlier, restricting customers' ability to manage changes themselves can be seen as a shortcoming with the platform. To maximize customer benefits while maintaining the provider's ability to differentiate, vCloud Director allows both the provider and the customer to access the solution through the same interface. This provides that changes made by either party are consistently applied, and reduces the risk of the customer adversely affecting a solution for which the provider is accountable within a Service Level Agreement.

The ability for both the provider and customer to manage the same solution is particularly useful within the networking services part of a provider service, where customers might have in-house skills to manage their business applications which reside on the provider platform, but might not have the networking or security skills to be able to set up a complex micro-segmentation distributed firewall policy. Similarly, if the customer wants to take advantage of NSX Edge load balancing, while they might understand their applications load balancing requirements, they might prefer that the service provider be responsible for procuring, installing, and managing SSL certificates for those same services.

Figure 32. Managing Certificates in the Edge Services Gateway



² Most user permissions are managed through the vCloud Director user interface, but some, notably those to do with advanced networking and the distributed firewall must be managed through the vCloud Director API. See the References section for more information.



The certificate management dialog in vCloud Director might be quite straight forward to Content Delivery Engineers who understand signing requests and revocation lists, but customers might be willing to pay for that expertise from their VMware Cloud Provider.

While the management of SSL certificates is just one example, it illustrates that simply by providing customer access to manage part of their service, the provider does not preclude offering to manage those same services for their customer. Even if a customer tries and fails to configure an element of their service themselves, the provider knows that they will be able to access the failed configuration through the same interface that the customer used and either, help the customer to remediate the error themselves, or take over and complete the deployment on the customer's behalf. In either case, customer satisfaction is assured through offering them self-service, support service, or a managed service, as required.

6.2 Additional Product Licensing

Deploying vSphere into a managed hosting environment requires the hosts and their vCenter Server to be appropriately licensed to enable the required feature set. The introduction of vCloud Director and NSX similarly require appropriate licenses. For service providers using the VMware vCloud SP Bundles to license their current estate, vCloud Director is included³ in the 5-point bundle, but this does not include NSX. NSX for Service Providers is available in three different versions depending upon the features required. Those versions are known as NSX-SP Base, NSX-SP Advanced, and NSX-SP Enterprise.

If a tenant pod does not use the distributed firewall, dynamic routing or L2VPN connectivity, the Edge Services Gateway (with static routing only) and VXLAN-based networking features are available in the NSX-SP Base edition which is included in the “Advanced” 7-point bundle. To enable the distributed firewall and deliver the benefits of micro-segmentation to the tenant, or add dynamic routing or L2 VPN, the customer's workloads must be licensed using NSX-SP Advanced which is included in the “Advanced with Networking” 9-point bundle. Should the customer require the integration of externally hosted, physical devices into their VXLAN networks, this would require NSX-SP Enterprise which is included in the “Advanced with Networking and Management” 12-point bundle.

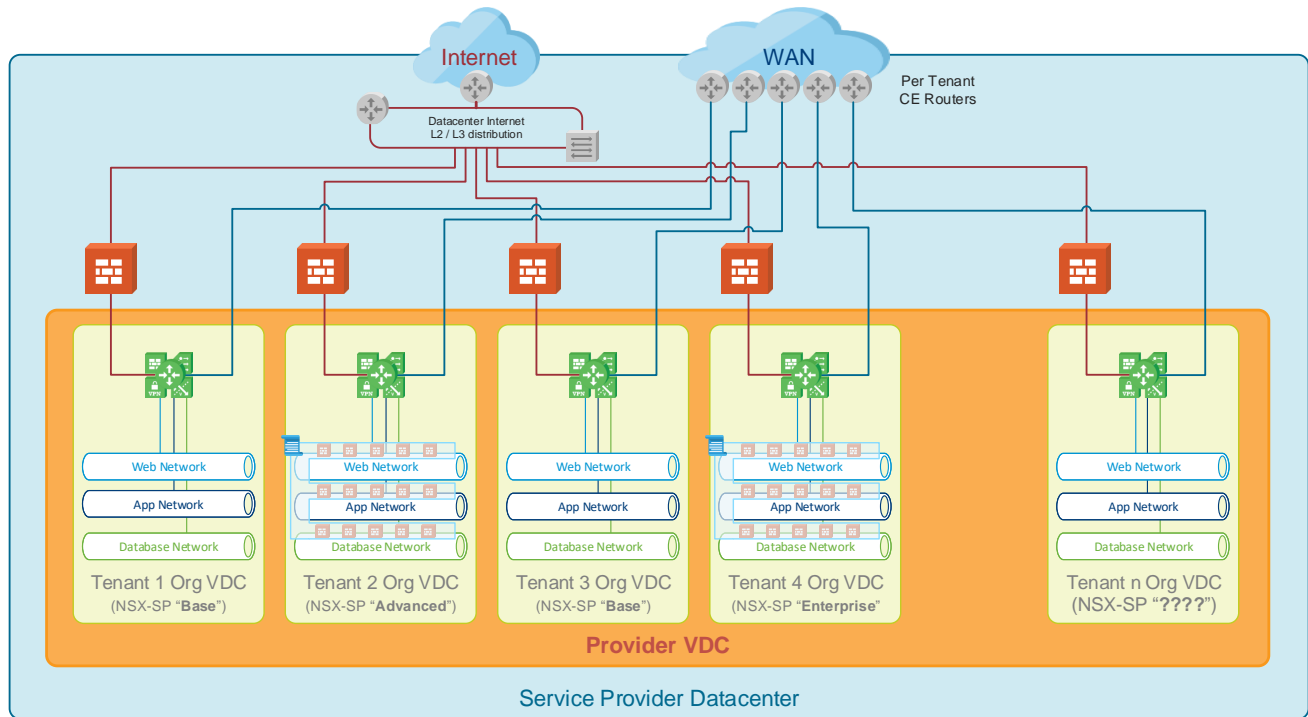
It is quite likely that an existing managed hosting provider introducing cloud services will want to offer the service to their existing customer base. The availability of the benefits of NSX such as micro-segmentation does not mean however, that all customers will immediately be increased to a 9-point or 12-point bundle. VMware vCloud Usage Meter can be used to report on multiple tenants who each use different features, allowing the service provider to upsell new capabilities to their existing customer base to increase revenue in line with the increase in license points, rather than having to increase their costs in advance of any incremental revenue being generated to recoup an upfront investment.

The following figure illustrates the service provider data center model used earlier with the addition of NSX-SP Advanced or Enterprise in some tenants, while others require only NSX-SP Base with the resulting smaller increase in points consumption there.

³ Product bundle comparison using the VMware vCloud Air Network Product Usage Guide Q2 CY2017



Figure 33. Cloud Service Provider Data Center with Per-Tenant Licensing Options



In this figure, Tenant 1 and Tenant 3 are using the Edge Services Gateway in VMware vCloud Networking and Security™ mode and VXLAN logical networking, features which, together with vCloud Director are available in the “Advanced” 7-point bundle. Tenant 2 is using the Distributed Firewall feature which requires NSX-SP Advanced which in turn requires the “Advanced with Networking” 9-point bundle, whereas (although not shown in the diagram) Tenant 4 is using features that require NSX-SP Enterprise and the “Advanced with Networking and Management” 12-point bundle.

It is important to note, that there is no additional points consumption when the service provider makes the features available in a particular location. The increase in points consumption only arises when a service provider’s customer takes the additional services, which also means an increase in revenue for the service provider as a result.



References

Additional information pertinent to this document and its topics are listed in the following table.

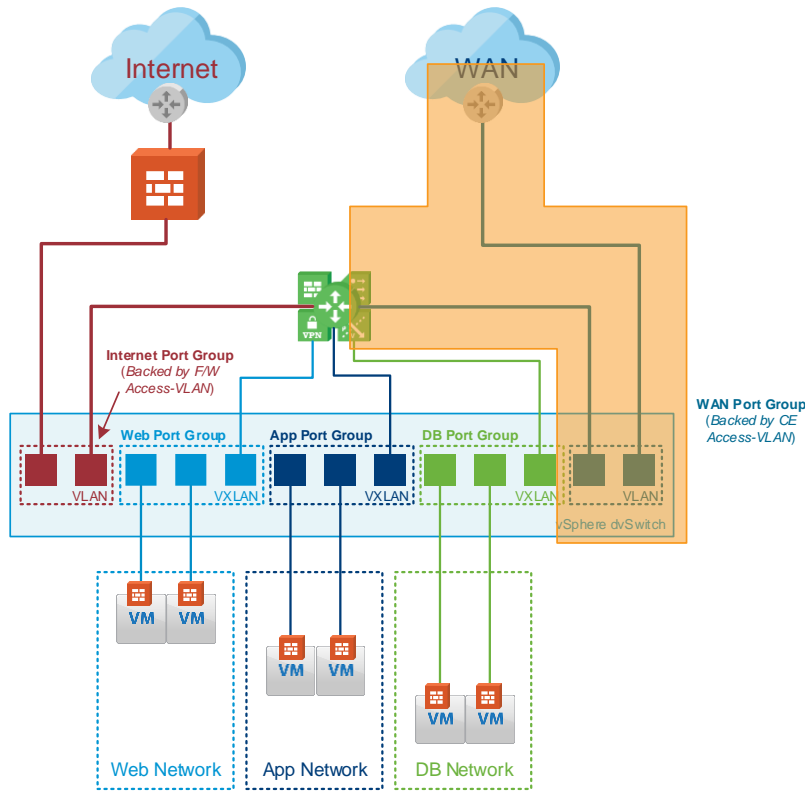
Document Title	Link or URL
<i>VMware vCloud Architecture Toolkit for Service Providers</i>	https://www.vmware.com/cloud-computing/cloud-architecture/vcat-sp.html
<i>vCloud Architecture Toolkit (vCAT) Blog</i>	https://blogs.vmware.com/vcat/
<i>Architecting a VMware vCloud Director Solution for VMware Cloud Providers</i>	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vcat/vmware-architecting-a-vcloud-director-solution.pdf
<i>Customer Onboarding with VMware NSX L2VPN Services for the VMware Cloud Provider Program</i>	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vcat/vmware-customer-onboarding-with-nsx-l2vpn-services.pdf
<i>NIST Definition of Cloud Computing</i>	http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistpecialpublication800-145.pdf
<i>Forrester Leveraging Micro-Segmentation report</i>	https://www.vmware.com/ciovantage/article/forrester-report-available-zero-trust
<i>VMware NSX for vSphere Network Virtualization Design Guide</i>	https://communities.vmware.com/docs/DOC-27683
<i>Using the vCloud API to Grant Distributed Firewall and Advanced Networking Services Rights in vCloud Director 8.20</i>	https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2149016
<i>Address Allocation for Private Internets</i>	https://tools.ietf.org/html/rfc1918
<i>vCAN Product Usage Guide</i>	https://vmware.my.salesforce.com/06980000000bpsy



Appendix A: Provisioning an External Network in vCloud Director

Section 3.2, Basic vCloud Director Tenant Topology and Section 3.3, Multitenant Networking introduced the concept of an external network that connects resources within the tenant Organization VDCs to resources in the Cloud Service Provider’s data center and beyond. It was noted in Table 2. vCloud Director Tenancy and Network Elements in NSX and vSphere, that because these networks bridge the vCloud Director management domain and the physical data center outside of the vSphere environment, they cannot be created or managed from within vCloud Director alone. The following figure shows the sequence of activities needed to connect a new WAN CE router to an interface on a newly provisioned tenant Edge Services Gateway.

Figure 34. Adding a New External Network



This figure shows the vSphere layer of the Tenant networking analysis from Section 4.1, Tenant Networking with the vCloud Director external network highlighted. In this example, a new WAN connection will be created between the vSphere environment behind the Provider VDC on which the tenant Org VDC and its Edge Services Gateway are configured. The new network details are listed in the following table.

Table 3. Appendix A Configuration Parameters

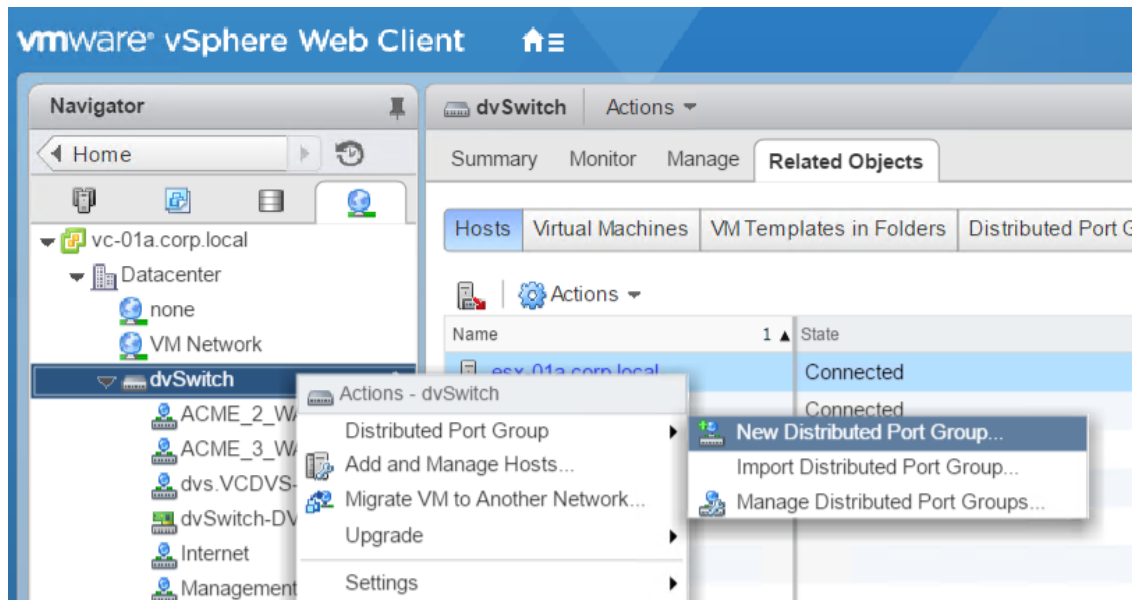
Configuration Element	Configuration Detail
Tenant WAN access VLAN ID	1011
VLAN Subnet	172.16.11.0/24
WAN CE Interface Address	172.16.11.254



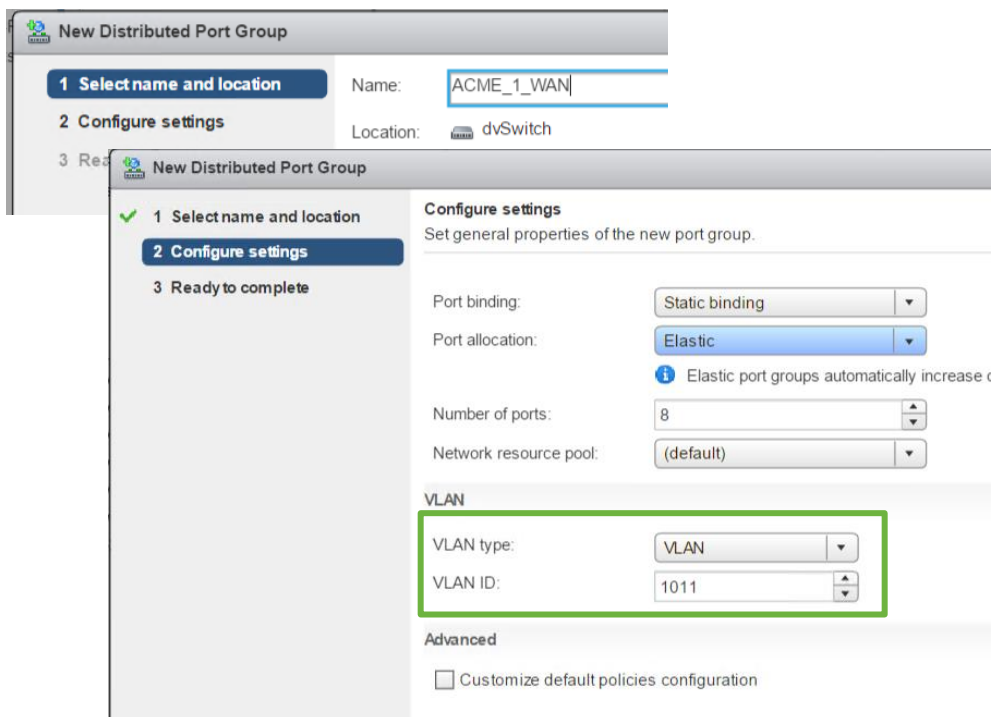
Configuration Element	Configuration Detail
Edge Services Gateway Interface Address	172.16.11.1

Configuration Steps

1. Create the new VLAN-backed Distributed Port Group in the PVDC vCenter Server to connect the (already configured within the data center network infrastructure) WAN Access VLAN to the vSphere dvSwitch.

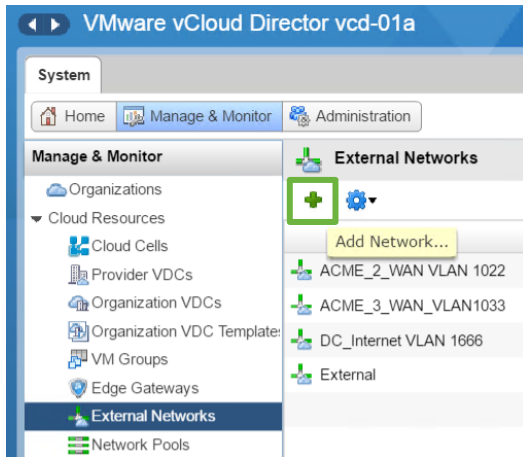


2. Name and configure the new dvPortGroup.

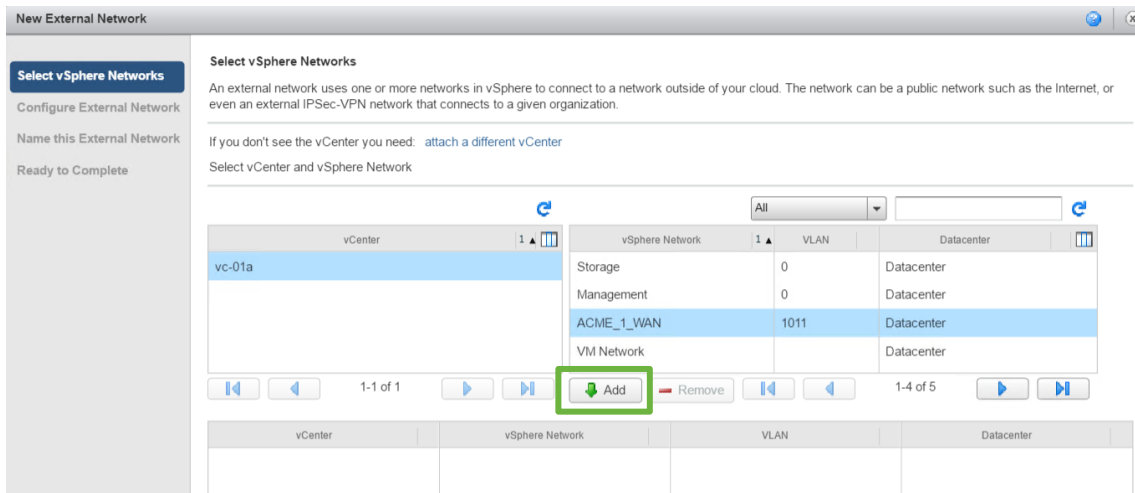




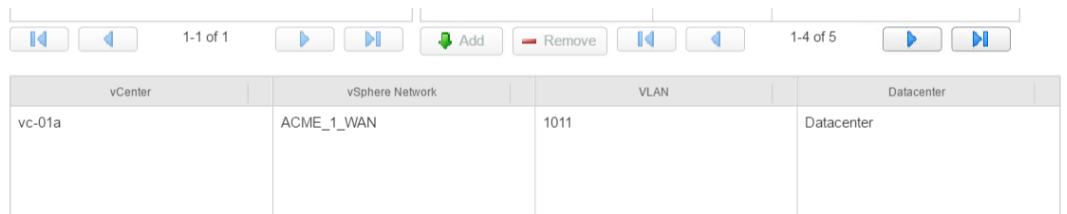
- After the new dvPortGroup is available in vCenter Server, it can be added as a new External Network within vCloud Director. In the External Networks view in the vCloud Director **Manage and Monitor** tab, click the green “plus” icon to open the **Add Network** dialog.



- In the **Add Network** dialog, select the vCenter Server to show available vSphere networks, and select the new ACME_1_WAN network with the correct VLAN next to it.

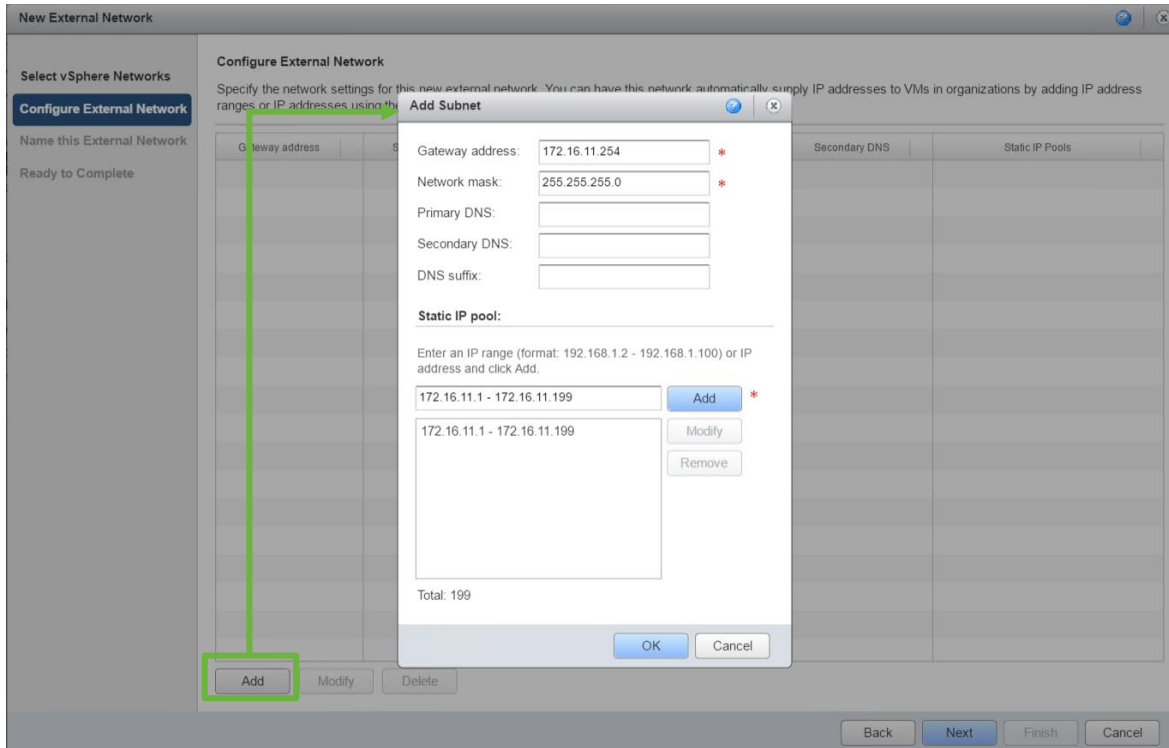


- After the network appears in the lower table, click **Next** to continue.



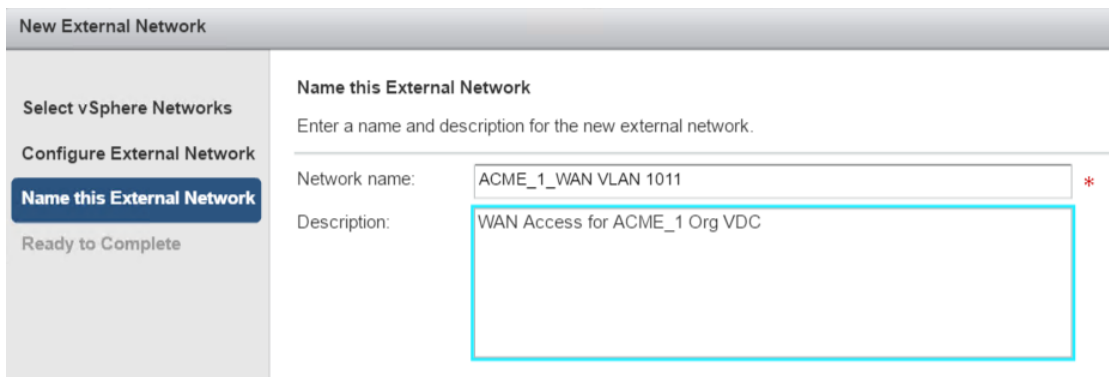


- Next configure the IP addressing details for the new network. This allows vCloud Director to understand the gateway address on the WAN CE router and the range of addresses on the new network which it can allocate.



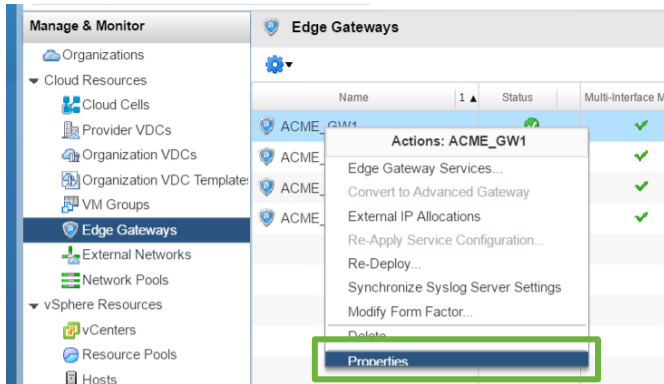
In this example, the pool of addresses which vCloud Director can allocate from has been restricted to addresses between 172.16.11.1 (which will be used for the edge interface address in a later step) and 172.16.11.199. Addresses from 172.16.11.200 – 253 could be added to the range, but are held back so that if the customer runs out of addresses, the provider can release the remainder and offer assistance with expansion.

- Give the external network a name and description and click through to complete the dialog to finish adding the new external network.

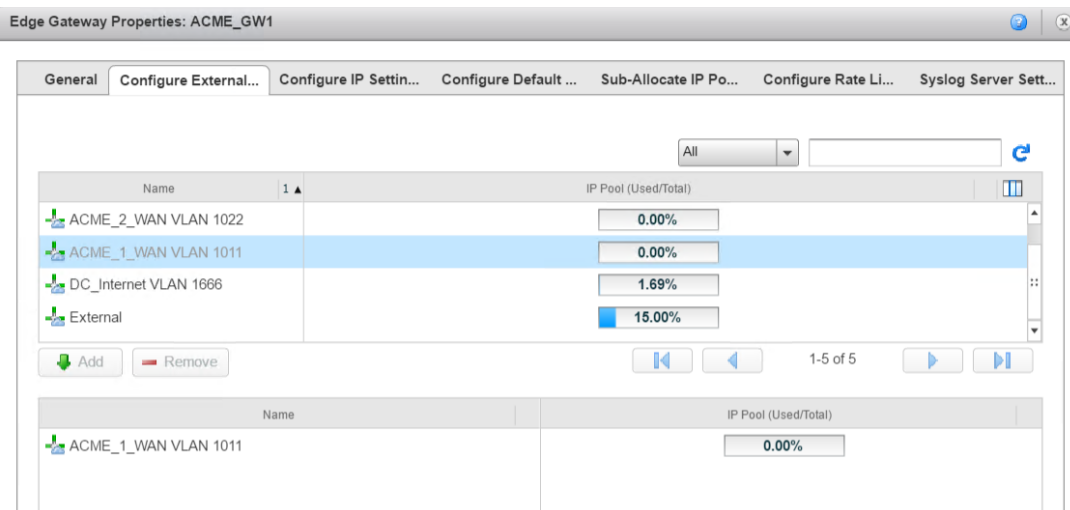




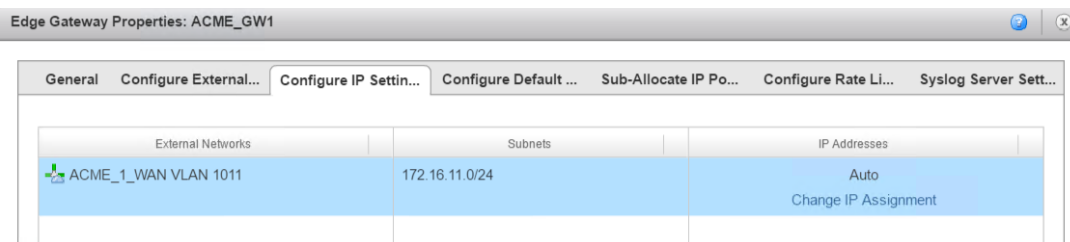
- External networks are not confined to a single Organization within the Provider VDC in whose vCenter Server they are configured, so the service provider must take care with the next step which presents the new external network to the required Edge Services Gateway. From the **Actions** menu of the Edge Services Gateway, select **Properties**.



- From the **Configure External Networks** tab of the Edge Services Gateway's Properties dialog, select the new external network from the list in the top table of candidate networks and click **Add** to copy the network into the lower table of connected networks.

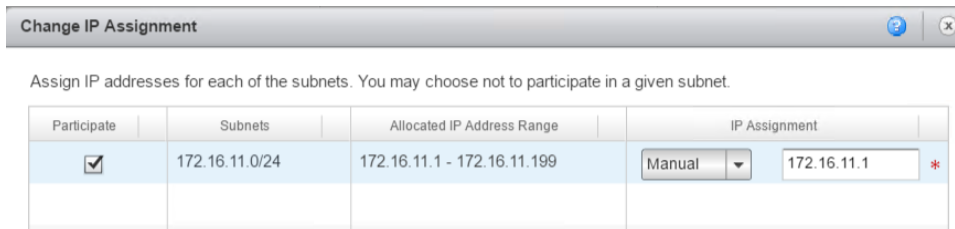


- On the **Configure IP Settings** tab, select the **Change IP Assignment** link on the new external network row to assign the address manually.

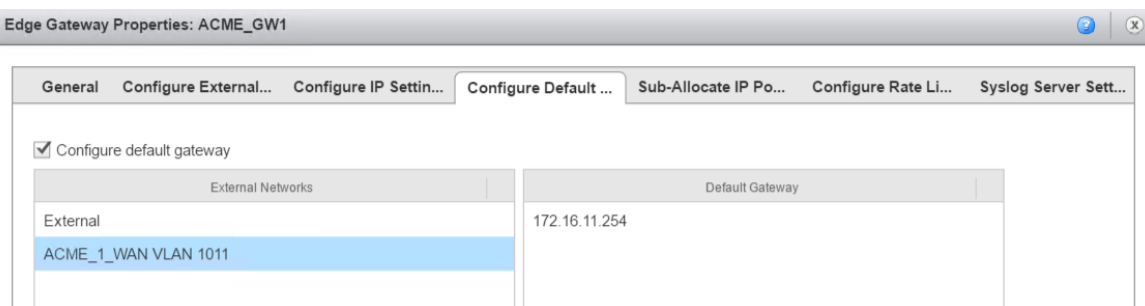




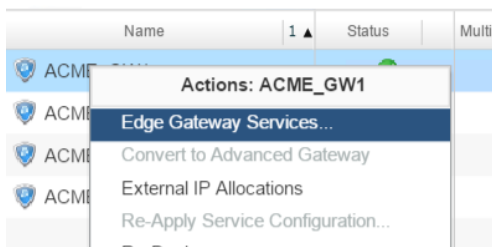
- In the **Change IP Assignment** dialog, set the **IP Assignment** mode to **Manual** and enter the required interface address.



- If the new network is to be the Edge Services Gateway's default route, the external network should be selected in **Configure Default Gateway** tab. Because the network is selected, the Default Gateway configured in Step 6 should appear in the right-hand column. After the changes are complete, click **OK** to close the dialog.



- After the infrastructure configuration is complete, the Edge Services Gateway might need additional changes to add extra static routes, configure peering over the new network, firewall rules or NAT entries. These configuration changes are carried out through the **Edge Gateway Services** option of the gateway's **Actions** menu.



- And the options are then configured in new tab using the HTML5 interface.

