



Global Knowledge®

Expert Reference Series of White Papers

Simplifying VMware vCloud

Simplifying VMware vCloud

A Clear Explanation of the Clouds Powered by VMware vCloud Suite

John A. Davis, VMware Certified Instructor (VCI), VCP5-DV, VCP5-DT, VCAP5-DCA, VCAP5-DCD

Introduction

One of today's buzzwords in the business and information technology (IT) communities is "cloud". To many, "cloud" is a positive term, much like a real cloud that relieves us from the intense heat of the summertime sun. But, to others, it is a negative term, much like real black clouds that bring us dangerous thunderstorms. Many think of a cloud as a simple environment that allows businesses to run their applications without having to get involved in underlying infrastructure. Others think of a cloud as a very complex environment that forces businesses to run their applications without the visibility and control of the underlying infrastructure. Part of the reason for varying perspectives is that today's clouds are built upon various technologies, where the cloud providers and technologies' vendors define the cloud differently. This white paper is aimed at clearly explaining clouds built upon VMware vCloud Suite.

This white paper has three main goals.

1. The first goal is to generate a better understanding of the cloud in both the business and IT communities. It describes the cloud using simple, yet effective terminology. It illustrates the cloud by providing simple examples and scenarios. It focuses on clouds built upon VMware vCloud Suite and VMware vSphere Enterprise Plus, which this paper will reference as "vClouds".
2. The second goal is to describe the major components of vCloud and the virtual data centers they provide. It provides simple explanations of how each component interacts through VMware vSphere to the underlying physical components.
3. The third goal is help businesses visualize and understand how vClouds could be beneficial in addressing their specific IT needs.

Introduction to the Cloud

Physical infrastructure is a collection of physical IT assets. It includes items like CPU, RAM, SCSI devices, Storage adapters, Ethernet adapters, and Ethernet switches. In traditional (or legacy) environments, business applications run directly upon physical infrastructure.

Virtual infrastructure is a collection of logical objects built upon the physical infrastructure. For example, a resource pool is a virtual infrastructure object that logically maps to a set of underlying physical CPU and RAM resources residing on a cluster of hardware servers. Resource pools can be used to "carve out" a specific amount of CPU and RAM for a particular use, allowing business applications to run within resource pools.

Cloud provides virtual infrastructure as a service, where sets of virtual infrastructure resources can be allocated for various uses, automatically. For example, a customer who needs to deploy new business applications may be able to obtain the necessary virtual and physical infrastructure automatically, without engaging the IT Team, by making the request through a web portal to the cloud. The cloud may meet the request by automatically manipulating the virtual resources assigned to the customer.

Clouds offer many other features, such as adaptability, availability, security, and manageability. This paper focuses mostly on the cloud's ability to provide virtual infrastructure as a service. VMware defines the cloud provided by VMware vCloud Suite as a "Software Defined Datacenter".

Details on the Cloud

In a traditional environment, business application software runs directly on physical Windows, Linux, and other servers. Each need for a new application typically requires the procurement, installation, configuration, and support of a new hardware server. Each new hardware server typically requires the configuration of supporting network and storage infrastructure, such as the configuration of a port on an Ethernet switch or the creation of a logical disk within a disk array. Once the consumer makes the request for the new server, a great deal of time may be required as the IT Team makes the appropriate procurements and configurations of the hardware, network, and storage resources.

In a virtual environment, business application software runs in virtual machines (VMs) on virtual infrastructure. Consumer requests for new VMs are typically met very quickly by the IT Team, but could require additional time if changes are necessary in the virtual and physical infrastructure. VMs can be rapidly deployed if available CPU and RAM exist in the resource pools, unused space exists in the datastores, available bandwidth exists in the network connections, and available virtual ports exist in the virtual switches. Whenever the appropriate amount of virtual infrastructure is not already available or the new VMs require additional isolation, such as a new network, then additional virtual resources must be configured. Configuring additional virtual infrastructure often requires manipulating the underlying hardware infrastructure, which requires additional time.

In a cloud environment, business applications run in VMs on virtual infrastructure that is automatically provided on demand by the consumer without involving the IT Team. In this case, the provider has already produced a set of virtual infrastructure to meet the consumer's existing need and short-term growth. The provider charges the consumer for these resources and allows the consumer to easily deploy new VMs within these resources. The provider may also allow the consumer to request additional virtual resources online. The cloud automatically provides the necessary virtual infrastructure. Naturally, this requires the provider to maintain spare virtual infrastructure resources that can be automatically assigned to consumers upon request.

Multiple customers (or tenants) share a cloud, although they are unaware of the other tenants. Each tenant is granted a virtual datacenter that is isolated from the other tenants.

VMware vCloud Suite Overview

VMware is considered to be the leading supplier of server virtualization technologies. VMware vSphere is utilized in some capacity by 100% of the Fortune 100 companies and 98% of the Fortune 500 companies. Most of

these companies have virtualized a large percentage of their production servers. The largest of these companies continuously express the need to be able to deliver their vSphere resources automatically, on demand. This need led to the creation of VMware vCloud Director, which is aimed at providing Infrastructure as a Service (IaaS). More specifically, vCloud Director provides virtual infrastructure built upon VMware vSphere automatically as a service. It provides a web-based portal where consumers can request, access, and manage their assigned virtual infrastructure.

VMware vCloud Suite is a set of products that includes VMware vCloud Director, VMware Network and Scaling, and VMware vSphere, which provides the cloud functionality. VMware vCloud Suite comes in Standard, Advanced, and Enterprise editions, where Advanced and Enterprise editions add enhanced features and additional products, such as management software. Cloud providers can utilize VMware vCloud Suite to produce and manage clouds built upon VMware vSphere. In this paper, such clouds will be referenced as "vClouds".

VMware vCloud Director provides several core constructs, including organizations, virtual datacenters, organization networks, vApps, and vApp networks. It includes other features, such as Fast Provisioning, vShield Edge Gateways, and vCenter Chargeback. Some of these features and constructs seem rather complex, especially the network components. The following sections of this paper explain the major constructs and features.

vCloud: Overview of Organizations and Virtual Data Centers

Within a vCloud, the provider can configure multiple organizations, where each organization is associated to a specific customer. In a public vCloud, each customer may be a unique company. In a privately owned vCloud, each customer may be a different department or business unit. Typically, each organization can be thought of as a customer of the vCloud, which has a specific set of users. Each organization owns its own virtual assets that are provided by vCloud, including its own virtual datacenter (vDC) and its own organization networks.

An organization vDC is a construct provided by a vCloud that is assigned to one specific organization (customer). Each organization can have more than one vDC. The provider creates the vDCs for each organization and defines parameters such as the maximum amount and the guaranteed amount of RAM to assign to the datacenter. Each customer perceives that they have their own datacenters, each consisting of processor, memory, storage, and network resources.

Technically, the vCloud provider first creates provider vDCs, which can then be subdivided into organization vDCs. Each provider vDC maps to a parent resource pool in vSphere, and the subordinate organization vDCs map to child resource pools attached beneath the parent resource pool, as illustrated in Figure 1.

Each provider vDC is typically intended to provide a different level of service. For example, a vCloud may contain two provider vDCs. One provider vDC is named "Tier-1", and it contains fully guaranteed processor and memory resources and Fiber Channel-backed datastores. The other provider vDC is named "Tier-2", and it contains 50% guaranteed processor and memory resources and NAS datastores. In this example, the provider intends to charge more for Tier-1 resources. Each customer may elect to purchase both Tier-1 and Tier-2 resources. One

organization vDC may be provided by the Tier-1 Provider Datacenter, while another organization vDC may be provided by the Tier-2 Provider Datacenter, as illustrated in Figure 1.

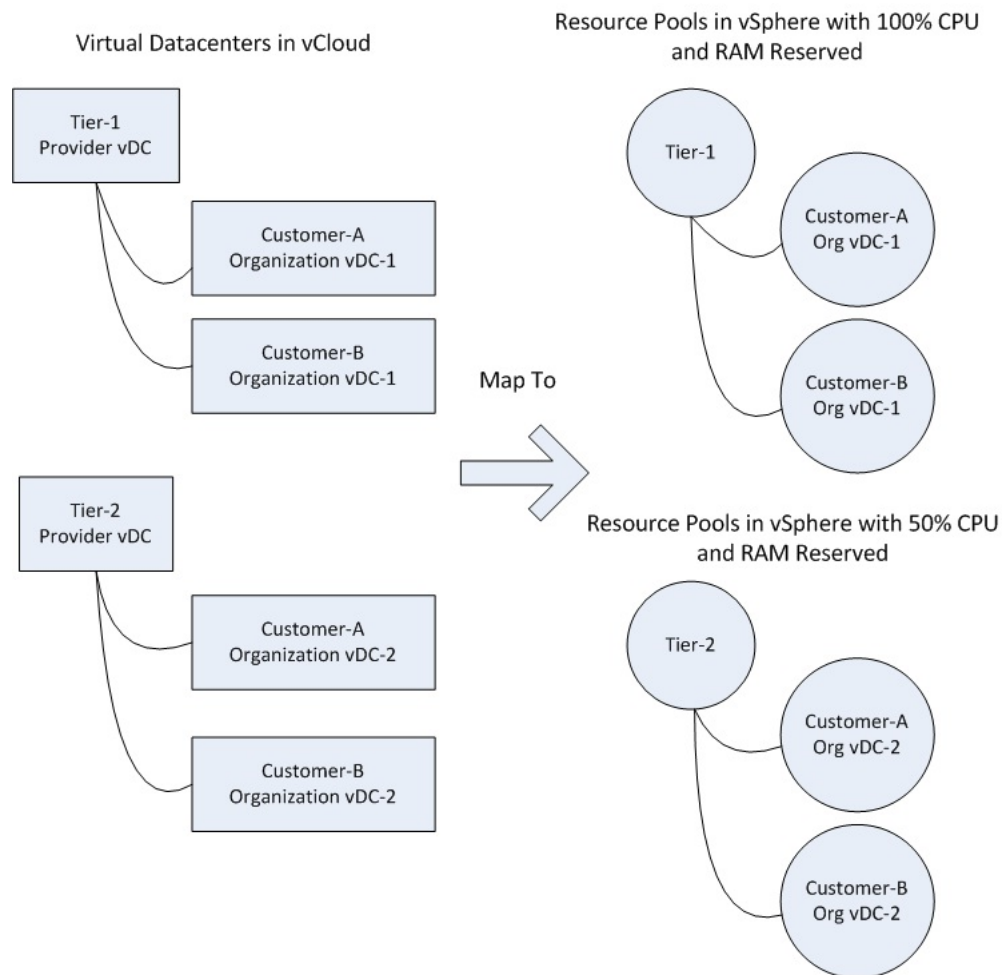


Figure-1

vCloud: Additional Features of Organizations and Virtual Data Centers

Organizations in vCloud can be used for more than just assigning vDCs. Organizations include additional features such as users, role-based permissions, catalogs, vApps, and a unique administrator web portal. Likewise, organization vDCs can be used for more than just allocating processor, memory, storage, and network resources. Organization vDCs include additional features such as allocation model, virtual disk provisioning type, and quotas. This paper only provides a brief explanation on some of these features.

- **Organization web portal** – Each organization has a dedicated web portal for configuring and managing their resources, such as VMs and users.

- **Organization Administrators** – Specific users within an organization that have the ability to configure all aspects of the organization vDC that has been delegated by the provider.
- **Organization users** – Typically these are provided by an LDAP, such as Active Directory, that is owned by the organization, but may be created manually. The organization administrator may delegate access and control to the organization's objects to various users.
- **Catalogs** – A set of vApp templates and ISO files used by the organization.
- **vApps** – One or more VMs that are treated as a single entity by the organization. The VMs in a specific vApp tend to provide a specific business application and communicate frequently with each other. A vApp may be placed on a protected network, called a vApp network.
- **Allocation Models** – The choices are:
 - **Pay as You Go** – The customer pays for just the resource used by vApps they actually run. The provider does not guarantee any resources upfront, but could proactively place a limit on the customer's use. Whenever a vApp is permitted to start, then a percentage of its resources may then be automatically reserved. This tends to be the least expensive Allocation Model choice.
 - **Allocation Pool** - The provider guarantees an agreed upon amount of resources upfront, but the customer is allowed to potentially access more resources, when available, up to a configured maximum. The provider tends to charge for just the guaranteed resources.
 - **Reservation pool** – The provider allocates (guarantees) and charges for a specific amount of resources to the customer upfront. In other words, all the customer resources are proactively reserved. The customer cannot access additional resources. This tends to be the most expensive Allocation Model choice.
- **Virtual Disk Provisioning Types** – Choices include Full Provisioning, where each VM contains its own independent virtual disks, and Fast Provisioning, where a set of VMs share a base disk and contain their own delta files.

Networking – The networking features in a vCloud are provided in the next section.

Networks in a vCloud

The most difficult concept to comprehend in a vCloud is the network. vCloud uses constructs called External Networks, Network Pools, vShield Edge Gateways, Organization vDC Networks, and vApp Networks. When first learning about vCloud Directory, many administrators struggle to fully understand the purpose and use of each of these objects. The many layers of networking, as well as the many associated settings, tend to be a major source of confusion. In this paper, a simple explanation is provided to help the reader understand the major concepts of networking within vCloud. This paper does not attempt to address all the low-level networking features and settings.

External Networks are networks created by the provider that connect to the outside world (networks external to the vSphere infrastructure). It connects to a specific virtual switch port group in vSphere that connects to the outside world. It is configured to provide parameters to any objects or networks that may connect to it. These parameters include the Gateway, Mask, and DNS settings. It is also configured with a set of allowed IP addresses that may be assigned to any objects that may attach to it. Conceptually, the provider intends to connect multiple Organization vDC Networks to each External Network and allow organizations to create additional

networks. Figure 2 shows an example of the provider's conceptual view of how an organization's networks connect to an External Network that is connected to a port group named External-Access.

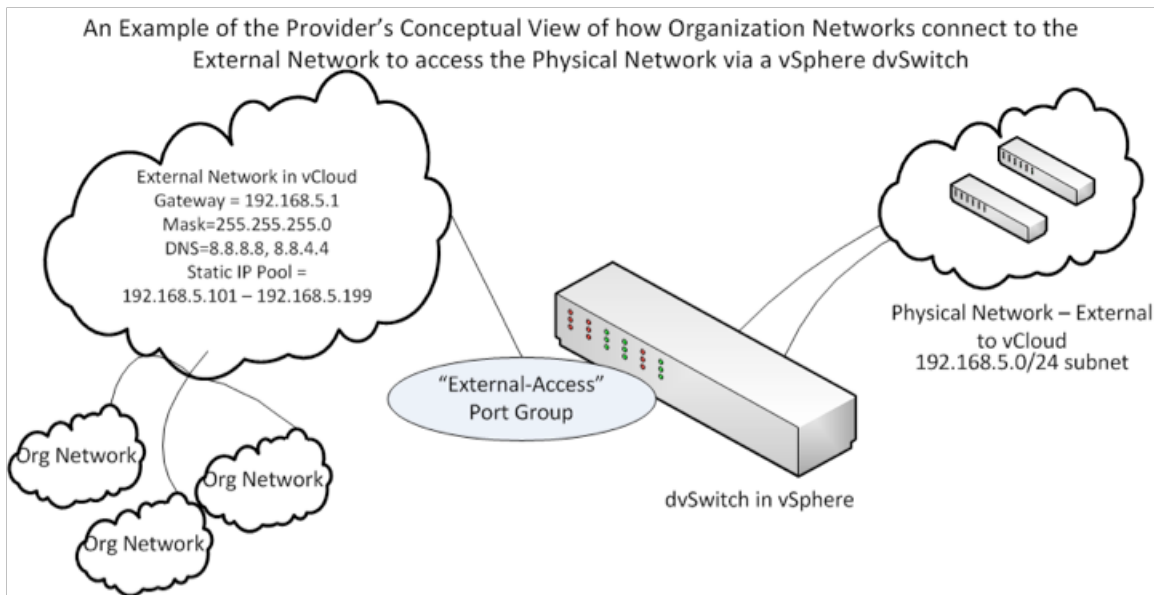


Figure 2

Network Pools are templates created by the provider that can be used to rapidly create Organization vDC Networks and vApp Networks. Network Pools provide configuration settings that will be used for each network created from pool. Each organization vDC is assigned exactly one Network Pool. Effectively, organization administrators may create new networks from the Network Pool resources.

For example, a network pool can be configured to use a specific dvSwitch, provide VLAN backing, and use a specific VLAN range. With these settings, any network created from the pool will be automatically assigned a VLAN number and attached to an automatically created port group on the specified dvSwitch. In this case, consider the Network Pool as the set of VLANs and the virtual switch on which an organization may automatically create new networks. Figure 3 provides an example of two Network Pools. Networks created from Network-Pool-1 will automatically create a new port group on dvSwitch1 with a VLAN from the 101-150 range. Networks created from Network-Pool-2 will also automatically create a new port group on dvSwitch1, but with a VLAN from the 201-250 range.



Figure 3

Network Pools do not specify which External Network will be used when creating a new network. However, organization administrators should plan to select an External Network when using a Network Pool to create a new Organization vDC Network.

Organization vDC Networks are networks that are built for specific organizations. An organization may have multiple Organization vDC Networks. Before building the first Organization vDC Network, the provider assigns a Network Pool to the organization and configures a vShield Gateway for the organization. The vShield Edge Gateway is configured to connect to a specific External Network using an available IP address from the External Network's pool of IP addresses.

After configuring the vShield Gateway and assigning a Network Pool to the organization, then Organization vDC Networks may be created. Each Organization vDC Network may automatically utilize the assigned Network Pool to determine which dvSwitch should be used to create the underlying port group and which VLANs to assign. It may automatically connect to the organization's vShield Edge Gateway to gain access to the External Network. Each Organization vDC Network may be assigned its own static IP Address Range, gateway, and mask. Additionally, other services, such as DHCP and a specific range of IP addresses for DHCP may be configured in each Organization vDC Network.

Organization vDC Networks may be built by the organization, but must be within the constraints established by the vShield Edge Gateway and the Network Pools that are assigned by the provider. In other words, the vShield Edge Gateway and Network Pools are tools used by the provider to allow organizations to be delegated to create their own networks within the boundaries established by the provider.

Figure 4 shows an example where two Organization vDC Networks that were created using Network-Pool-1 (from Figure- 3) connect to the same External Network via a vShield Edge Gateway. One Organization vDC Network is assigned to VLAN 101 and subnet 192.168.101.0/24. The other Organization vDC Network is assigned to VLAN 102 and subnet 192.168.102.0/24. Any traffic between the outside world (external to the vCloud) and either of the Organization vDC Networks must travel through the vShield Edge Gateway. The vShield GateVM_way supplies DHCP services and utilizes a separate pool of IP addresses for each of the Organization vDC Networks.

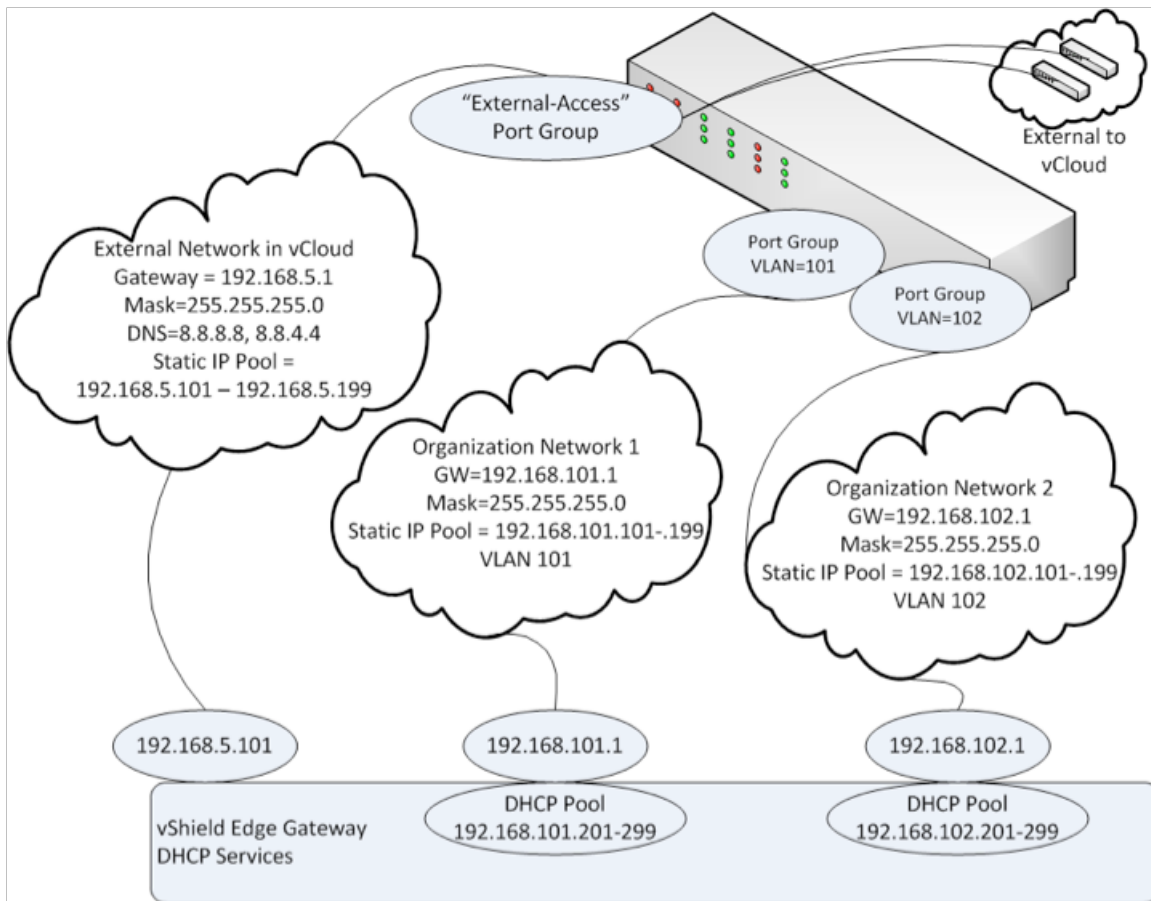


Figure 4

The lowest level network that vCloud provides is the vApp Network. The vApp Networks connect to Organization vDC Networks via vShield Edge Gateways, much like Organization vDC Networks connect to External Networks via vShield Edge Gateways. Network Pools are used to create both vApp Networks and Organization vDC Networks.

Many vApp Networks may connect to each Organization vDC Network, much like many Organization vDC Networks may connect to each External Network. Each vApp Network has configuration settings similar to that of Organization vDC Networks, such as static range of IP addresses, gateway, mask and DNS. DHCP and a range of DHCP IP addresses may be configured for the vApp Network, as well as other features such as port forwarding. Figure 5 shows an example of two vApp Networks that connect to an Organization vDC Network via a vShield Edge Gateway.

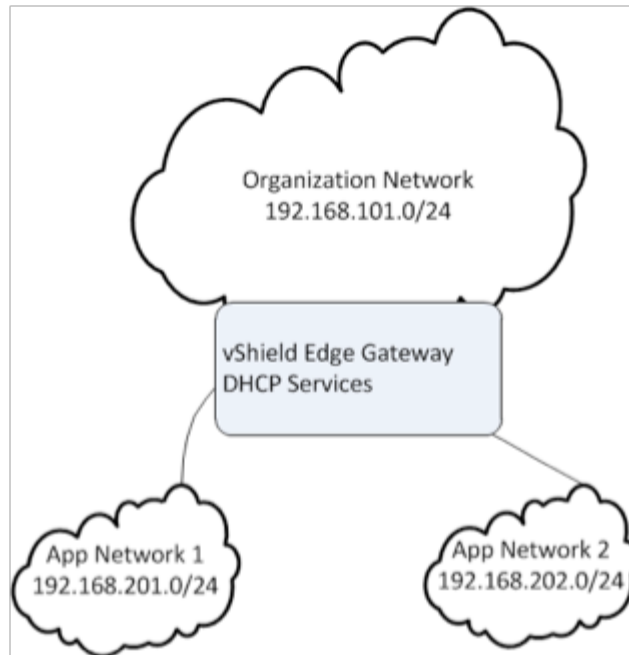


Figure 5

To summarize, organizations perceive they have their own physical network (External Network), which they may choose to sub-divide into multiple networks (Organization vDC Networks) as allowed by the provider (using vShield Gateway and Network Pool settings). Each Organization vDC Network may then be sub-divided into smaller networks (vApp Networks).

Why Small Businesses Should Consider Migrating to Public vClouds

Small- and medium-sized businesses that elect to migrate their applications to a public vCloud have the opportunity to save a great deal on IT expenses. Effectively, the cost of the underlying physical and virtual infrastructure is being shared with other businesses running in the same public cloud. And, the cost for labor related to configuring and supporting infrastructure evaporates as virtual infrastructure is automatically provided by the vCloud. Additionally, they tend to begin paying strictly for IT operating expenses (OPEX) versus IT capital expenses (CAPEX).

A public vCloud allows small business to migrate their applications and data to an environment with high levels of availability, resiliency, security, and scalability that, historically, only larger businesses could afford. Most small businesses simply could not afford to purchase the enterprise-quality storage systems, network devices, security systems, and virtual infrastructure that are required to achieve the levels of protection that enterprises can afford. Such businesses may determine that migrating their applications to a public vCloud is the most cost-effective way to provide the availability, resiliency, security, and scalability they desire.

Why Medium-to-Large Businesses Should Consider Migrating to vCloud

Many vSphere administrators in medium and large companies consider their vSphere environment to be a “private cloud”. They already enjoy the many benefits of virtualization, such as guest operating systems and applications being isolated from physical infrastructure. But as a vSphere environment grows, the administrator realizes that it is lacking certain features that a true cloud should offer. For example, in environments based solely on vSphere, the IT Team must be engaged each time a new network is required. By implementing vCloud Suite, tenants would be permitted to create their own networks under the pre-configured control of the IT Team.

For enterprises, which perceive each department, business unit, or some other organization entity as a separate customer, vCloud Suite can be a great fit. It allows the IT Team to treat each customer as a separate tenant, who perceives that they have their own isolated datacenter and network. Without the cloud, the enterprise administrators would likely have to manually configure traditional firewalls and routers to isolate their multiple customers.

Additionally, vCloud Suite Enterprise offers many other benefits aimed at improving manageability and accountability. Such benefits include those from the vCenter Operations Manager Suite, such as:

- **vCenter Chargeback** – Provides metered utilization and customer chargeback reports.
- **Application Awareness** – Automatically discovers software applications and their dependencies.
- **Performance and Capacity Optimization** – Provides dashboards for performance, utilization, and capacity analysis.
- **VM Configuration Compliance** – Provides guest operating system patching and configuration

For these reasons and more, enterprises may see a real benefit of migrating to a private vCloud. But even enterprises may wish to migrate to public vClouds. They may choose to utilize both a private vCloud and a public vCloud, which is considered a Hybrid vCloud. With a Hybrid vCloud, enterprises may easily migrate VMs from the private vCloud to the public vCloud and vice-versa. For example, companies with seasonal workloads may choose to normally run most applications in the private vCloud, but can easily deploy VMs in the public vCloud to cover additional workload as demand increases. This reduces the need for the companies to procure a great deal of spare, seldom used physical resources.

Summary

Clouds built upon VMware vCloud Suite and VMware vSphere may initially seem rather complex. But many businesses, from small companies to enterprises, may benefit by migrating to private, public, and hybrid vClouds. It is very worthwhile to gain familiarity with the features, components, and constructs in VMware vCloud Suite. This paper can serve as a good starting point.

Learn More

To learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge, Global Knowledge suggests the following courses:

VMware vCloud Director: Install, Configure, Manage [V5.1]

VMware vCloud: Deploy and Manage the VMware Cloud [V1.5]

VMware vCloud: Design Best Practices [V1.5]

Visit www.globalknowledge.com or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

John A. Davis has been a VMware Certified Instructor (VCI) and VMware Certified Professional (VCP) since 2004, when only a dozen or so VCIs existed in the USA. He has traveled to many cities in the USA, Canada, Singapore, Japan, Australia, and New Zealand to teach. He splits his time between teaching and delivering professional consulting services that are 100 percent focused on VMware technology.