Global Knowledge®

Expert Reference Series of White Papers

Shortcuts to Speed Your vCloud Deployment

Shortcuts to Speed Your vCloud Deployment

Rebecca Fitzhugh, VMware Certified Instructor and Consultant

Introduction

VMware vCloud Director is VMware's cloud solution built on VMware's technologies and solutions as an abstraction layer above vSphere in order to deliver cloud computing. This approach provides logical pooling of memory, compute, storage, and networking resources in order to allow for on-demand user self-service provisioning of VMware server-based virtual machines. Cloud computing services can be achieved with three different layers of service delivery:

- Software as a Service (SaaS)
- Platform as a Service (PaaS)
- Infrastructure as a Service (IaaS)

vCloud Director focuses solely on the IaaS layer of service delivery.

vCloud Director is not for every VMware customer. Most enterprise customers that decide to implement vCloud Director have multiple "tenants" or organizations; these tenants are potentially hosting customers or different divisions of a company but, ultimately, the tenants represent a set of logical users that need secure isolation from one another and access to self-provision their own infrastructure resources.

vCloud Director can be used to deploy different types of clouds. A private cloud is an infrastructure whose resources are only used internally. A public cloud is an infrastructure made available to external customers for a price. A hybrid cloud combines two or more clouds with some kind of standardized technology, like VMware vCloud Connector, while each cloud maintains its own unique identity.

There are many different layers and components that make up vCloud Director and the first deployment of this product can be challenging for many administrators. This white paper is designed to help demystify, highlight and clarify several topics in order to provide shortcuts to a new vCloud Director deployment.

Decide Which Cloud Components Are Needed

The most important part of any vCloud Director deployment is a solid design. There are many components that make up a vCloud infrastructure, and it is important to understand which each provides and whether or not the component is needed to determine the design requirements. Though not all of the discussed components are a requirement for a vCloud deployment, many should be used for ease of cloud administration.

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.

Table 1 lists many components of a vCloud infrastructure and whether or not the component is required for deployment.

vCloud Component	Required	Description
vCloud Director cell	Y	Manages connectivity to the cloud and provides both API and UI end-points, cloud coordinator, abstracts vSphere resources Requires vCloud Director Database
vCloud API	Y	API used to programmatically interact with vCloud
vSphere	Y	Provides underlying virtualized resources Consists of vCenter, vCenter Database, clustered ESXi hosts and vSphere Management Assistant
vCloud Network and Security	Y	Provides network security services Edge devices are automatically deployed by vCloud Director
vCenter Chargeback	N	Provides resource metering and reporting to aid showback
vCenter Orchestrator	N	Facilitates orchestration at the vCloud API and vSphere level
vCloud Request Manager	N	Provides provisioning workflows for requests and approvals, assists with software license tracking
vCloud Connector	N	Facilitates transfer of powerful entities from a local vCloud/ vSphere to a remote vCloud

Table 1. Components of a vCloud Infrastructure

Multiple vCloud Director cells should be used to address availability and scalability concerns. This is typically achieved by load-balancing or content switching this front-end layer. Load balancers assist in presenting a consistent address for services regardless of the underlying node responding. The load balancer can spread session load across cells, monitor cell health and add/remove cells from the active service pool.

VMware vCloud Director 5.1 supports Oracle version 10g and 11g as well as SQL databases. All vCloud Director servers share a single database instance; this is an important design consideration to remember.

The vCloud infrastructure design will most likely have an NFS server behind the vCloud Director cells, or to be more precise, between the servers. This will provide scratch space for when users are uploading vApp templates into an organization's catalog. These vApp templates are actually staged on the NFS share during the upload. NFS is used in this case because it is easy to deploy, for its file locking and file sharing properties, and because NFS is a storage platform that is accessible from many operating systems. The size for this volume will vary, depending on how many concurrent uploads are in progress. Once an upload completes, the vApp is moved to permanent storage on the datastores backing the catalogs for each organization, and the data will no longer reside on the NFS volume.

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.

There is a 1:1 ration between the vCenter Servers that back vCloud and vCloud Director Network and Security servers. The vShield Manager will be paired with a single vCenter and the vCloud Director cells will provision Edge devices as needed during the network creation process.

While a vCloud Director administrator can create local users within vCloud, it is best practice to use an external LDAP source. Table 2 shows the supported LDAP platforms and which authentication methods each provides. If vCenter Single Sign On is configured and deployed, vCloud Director can be configured to use SSO to import users and give permissions within the cloud.

Platform	LDAP Server	Authentication Method(s)
Windows Server 2003	Active Directory	Simple, Simple SSL, Kerberos, Kerberos SSL
Windows Server 2008	Active Directory	Simple
Windows 7 (2008 R2)	Active Directory	Simple, Simple SSL, Kerberos, Kerberos SSL
Linux	OpenLDAP	Simple, Simple SSL

Table 2. Supported LDAP Platforms and the Authentication Methods Each Provides

Understanding what function each of the vCloud infrastructure components provide aids in determining which should be included in the design.

Understand the Relationships Between vCloud Constructs

vCloud Director introduces logical constructs with virtual datacenters and security boundaries with organizations to facilitate and preserve multi-tenancy utilization of resources. Table 3 lays out the major constructs of vCloud Director 5.1 and a brief description.

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.

vCloud Director Object	Description
Provider Virtual Datacenter	Collection and abstraction of vSphere CPU, memory, and storage resources. Uses underlying vSphere cluster or resource pool and one or more datastores
Organization	Logical group of users to which IT services will be presented. Provides a security boundary, so appropriate resources and controls can be set up for that given group of users.
Organization Virtual Datacenter	Partitioned from a provider vDC to provide resources to an organization. Allocates resources using one of three models: Pay as you go, Reservation, and Allocation.
External Network	Network that connects to outside the cloud using a predefined vSphere network port group.
Organization Network	Network defined and visible within an organization. Can be a direct or routed connection to an external network or internal only to the vApps within the organization.
vApp Network	Network defined and visible within an organization. Can be a direct or routed connection to an external network or internal only to the vApps within the organization.
Network Pool	Predefined collection of vSphere network resources that can be used by vCloud Director to create a number of private and NAT-routed organizations and vApp networks as needed.

Table 3. The Major Constructs of vCloud Director 5.1

Provider virtual datacenters are built directly on top of the vSphere resources and can be backed by either an ESXi cluster or a resource pool. The current best practice is to have a 1:1 mapping between a provider vDC and a cluster. Service levels for vCloud tenants are differentiated at the provider vDC level. It is best to divide storage into different tiers of cost based on storage device expense and storage device speed and then associate each tier with a different provider vDC.

An organization virtual datacenter allows the provider vDC to share resources with multiple tenants or organizations. The tenants don't have the ability to see the actual resources available in the provider vDC, only what is available in the organization vDC. Resource allocations models can be applied at this level, ensuring performance and control of capacity requirements while still maintaining security.

There are a few rules of thumb to remember for the relationship between provider and organization vDCs:

- An organization can have multiple organization vDCs
- Each organization vDC uses resources from a single provider vDC
- A provider vDC's resources can be used by multiple organization vDCs
- An organization vDC cannot be created until a provider vDC exists

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.

vCloud Director delivers our IT services in packages that are called vApps, which exist within organizations. These vApps are actually composed of one or more virtual machines. The virtual machines communicate over networks that are built into the vApp package and use the different resources and services in its deployed environment. The package also includes an Open Virtualization Format (OVF) descriptor, which provides general application information, hardware requirements, deployment instructions, and policies that are enforced during the vApp's runtime.

Differentiate Between Allocation Models

When the vCloud administrator is creating an organization vDC, it is a requirement to choose an allocation model. The allocation model defines how the provider vDC commits resources to the organization vDC and how the customer will be billed for these resources. There are three different allocation models available:

- Pay as you go
- Allocation
- Reservation

The **pay as you go model** is usually the easiest model to understand and administer. The easiest way to think of it is that customers pay for what they use. When a vApp powers on, the resources are committed. If a vApp is not powered on, then the customer is not being billed for resources. The customer can also specify a guaranteed percentage of resources, which allows you to overcommit resources. Even though the customer is billed as soon as a vApp is powered on, only a percentage of the resources are guaranteed to be available for utilization.

The **allocation pool model** configures a virtual container of resources, allocating a subset of resources. But it will guarantee to a tenant only a percentage of what has been allocated, which means that the provider has the ability to overcommit resources when using this model.

The **reservation pool model**, generally the most expensive allocation model offered to customers, essentially configures a physical container of resources. The benefit is that customers are in complete control of the resources they use, and all resources are guaranteed. The customers have the same controls that a vSphere administrator would have over resource pool settings, such as shares and reservations, to manage over-commitment of resources between their workloads. Effectively, this means that over-commitment is possible, but the customer controls it.

When choosing an allocation model, factor virtual machine admission control into the decision. Admission control is whether or not a DRS-enabled cluster will allow a virtual machine to be powered on. It is based on available resources. The different allocation models directly impact how admission control is used in the DRS-enabled cluster. Best practice currently includes using separate organization virtual datacenters for each type of allocation model. Mixing the different types may lead to unpredictable resource usage.

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.

Identify Network Configurations Needed

External networks are the first network objects created within vCloud Director. This provides a connection from the cloud to the outside world; this is not necessarily an Internet connection but can be a connection to the underlying vSphere layer. An external network is created from a predefined vSphere port group, and multiple external networks can exist on the same physical LAN as long as they are separated by VLANs. An external network can be dedicated to a sole organization or shared across multiple organizations.

An organization network provides network services and connectivity to a specific organization. vApp networks provide network services and connectivity to a specific VM or set of VMs within a vApp. Both of these network types can be created as isolated, routed, or directly connected to a higher network layer. A direct-connect network connects an organization to an external network or, in the case of vApp, the vApp directly to an organization network. There is no Edge device deployed for a direct-connect network. Routed networks connect via an Edge device, which provides services such as DNS forwarding, port forwarding, NAT, fencing, etc. Lastly isolated networks are internal only to that organization and vApp, and no traffic is routed externally.

Network design and creation within vCloud Director can quickly become complex due to the number of layers involved; therefore, here are a few rules to remember:

- Organization networks cannot be directly connected to each other. The organization networks should be connected to a shared external network for communication.
- Multiple vApps cannot connect to a single vApp network. Instead, connect multiple vApp to a common organization network.
- A vApp network can only be connected to a single organization network.
- Virtual machines within a vApp can be multi-homed; however, each vNIC can only connect to a single network.
- Networks can only be deleted if all dependencies have been removed first.
- Limit network names to 33 characters or shorter because vCloud Director adds a unique identifier, as long as 47 characters, to the end of the network name when an Edge device is provisioned.

Network pools are used by vCloud Director to create private networks to facilitate VM to VM communication and NAT-routed networks. There are four types of network pools that can be used; these are displayed in Table 4. Multiple network pools may be created, each a different type, but at a minimum at least one pool must be defined in order to create routed networks.

Network Pool Type	Description	Considerations
vSphere Port Group-backed	Port groups are manually pre-created in vSphere and specified for network pool use.	Only network pool that can use Standard vSwitches and the Nexus 1000V; least flexible, difficult to manage.
VLAN-backed	VLAN range is specified and port groups are created dynamically as the VLANs are used by vCD.	Physical switches need to be trunked according for VLAN range; flexible, best network performance; could exhaust VLAN IDs.
vCloud Network Isolated-backed	Creates an overlay network that used MAC-in-MAC encapsulation for each isolated network. vCD creates port groups as needed.	MTU must be modified to account for the VCD-NI header due to the MAC-in-MAC encapsulation. Not-routable.
VXLAN	Uses MAC-in-UDP encapsulation to provide Layer 2 abstraction regardless of physical location. Automatically created in vCD 5.x when a provider vDC is created.	ESXi9 hosts must be prepared to use VXLAN through the vCloud Network and Security appliance. The Segment ID Pool and Multicast address must also be specified. MTU must be modified to account for VXLAN header.

Table 4. Four Types of Network Pools

Conclusion

VMware vCloud Director fashions the provisioning of the software-defined datacenter layer to allow for a full virtual datacenter deployment within a short period of time. A vCloud consists of many layers and can quickly become a complex architecture. Before any deployment, requirements should be defined so that the vCloud can be designed to offer those services needed. It is imperative to understand the many components of vCloud, how each vCloud construct fits, which allocation models are available, and what network options can be chosen.

References

For more information on vCloud Director and the features mentioned in this paper, see the following documents on VMware's website:

VMware vCloud: Architecting a vCloud – Technical White Paper

vCloud Director Administrator's Guide [v5.1]

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

VMware vCloud Director: Install, Configure, Manage [v5.1]

VMware vCloud: Design Best Practices [v1.5]

VMware vCloud: Networking and Security for vSphere Professionals [v5.1]

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

About the Author

Rebecca Fitzhugh is a VMware Certified Instructor and consultant whose primary focus is on VMware virtual infrastructure products and vCloud Director. Prior to becoming an instructor and consultant, she served five years in the United States Marine Corps where she assisted in the build-out and administration of multiple enterprise networks residing on virtual infrastructure. Rebecca currently holds the VCAP-DCA, VCP-DCV, VCP-DT, and VCP-Cloud certifications, as well as various other industry certifications.

Copyright ©2013 Global Knowledge Training LLC. All rights reserved.