VMware vCloud® Architecture Toolkit™
for Service Providers

Architecting Multisite VMware vCloud Director®

Version 2.9

January 2018

Steve Dockar

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

Contents

List of Tables

List of Figures

Introduction

1.1 Overview

Service Providers typically offer one or more availability levels to their customers based upon the nature of the solution, customer requirements, and price point of the service. Secondary or “disaster recovery” solutions might be designed with single points of failure to minimize the cost of a service the customer does not plan to use, or to rely on, often. The designs for primary or “active” sites often include reducing the risk of a single device failure affecting service by using infrastructure elements in pairs. These pairs operate in high availability (HA) mode where one device will take over seamlessly in the event of a failure of the other device. However, HA cannot mitigate against site wide outages, such as those caused by major power outages or other environmental issues. To deal with the potential loss of a service location, often as a requirement imposed by business insurers, services must be delivered from multiple, geographically separate locations. Introducing multiple geographic locations to a customer’s service presents the challenge of providing the customer with the opportunity to manage services in multiple locations while not overcomplicating the customer experience.

Previously, Cloud Providers have used customized portals to offer a single “frontend” to multiple, disparate services. This typically requires investing significant resources into realizing the benefits of the underlying platforms through their portals. Often these portals use Application Programming Interfaces (APIs) on the underlying products, but must then represent the capabilities of the products’ user interface (UI) within that of the portal. VMware vCloud Director® is a product with a comprehensive API that has been used by Cloud Providers to differentiate, add value, and present multiple, disparate vCloud Director instances through a single portal’s UI.

With the introduction of vCloud Director v9.0, Cloud Providers can now offer access to multiple, independent vCloud Director instances through a single point of access. Because of the distributed nature of this access, Cloud Providers can choose to offer a single, global access point, or, should they prefer, several, regional access points which each in turn offer access to, and management of, multiple vCloud Director instances.

1.2 Document Purpose and Scope

This white paper examines the choices available to a VMware Cloud Provider offering service from multiple vCloud Director v9.0 platforms. It looks at the design philosophy behind multisite operation in vCloud Director, examining the principles which have led to the capabilities in the latest release of the product and which will shape multisite capabilities in future releases. The document discusses the options for grouping workload platforms behind vCloud Director instances, and grouping those instances behind a single customer access point. It examines the configuration required to group instances together at the Cloud Provider and Tenant levels, considers authentication and authorization enhancements and requirements, and looks at new infrastructure requirements to support multisite capabilities.

Multisite configurations in vCloud Director v9.0 rely extensively on the vCloud Director API. An understanding of both the API structure, REST, and the XML responses the API typically returns is assumed for the sections of this document which cover those topics. A deep knowledge of those areas is not assumed or required for the rest of this document. The multisite elements incorporated in the new vCloud Director HTML5 UI are examined, but the migration of existing elements of the previous UI are out of scope of this document

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.
Organization	The vCloud Director "parent" object which provides an administrative connection between users and their allocated resources (in the form of Organization VDCs).
Tenant	The portion of the infrastructure that is used by, and provides services to, the customer.
User	vCloud Director object which represents an individual customer (or service provider) representative accessing the service through a preconfigured account.
Site	A geographic location within which a vCloud Director instance runs. Usually a single physical location, but can also be spread across multiple, connected locations. (See Metro-Cluster.)
Metro-Cluster	A VMware vSphere® deployment built on a metropolitan area network which allows the compute, network, and storage components to be deployed in different locations (where transmission latencies between locations are low enough), while remaining part of the same logical "site".
Load balancing	The distribution of load in which the proportion of the load sent to each receiver is balanced according to predetermined rules.
Load sharing	The distribution of load without consideration to the proportion of load sent to each receiver.

1.3.2 Acronyms and Abbreviations

API	Application Programming Interface is a mechanism for controlling an application programmatically, typically from an external source without the need to use the user interface of the application.
BCDR	Business continuity and disaster recovery is a term used to describe the combined planning and actions which will be enacted shortly after an incident (disaster recovery) and then, over a longer period (business continuity) to maintain the operational effectiveness of an organization.
DNS	Domain Name System is a distributed, hierarchical system for resolving domain names to IP addresses. It can also be used to resolve domain names to other "alias" domain names or return information about a domain.
DR	Disaster recovery – see BCRD.
OVDC / OrgVDC	Organization virtual data center is a collection of compute resources made available to an organization user for the placement of workloads.

PVDC	Provider virtual data center is a collection of similar compute capacity managed by the Cloud Provider from which Organization VDCs are allocated.
URI	Uniform Resource Identifier is a compact sequence of characters that identifies an abstract or physical resource, typically in the form of a location or name.
URL	Uniform Resource Locator is a member of a subset of URIs which contains the resource's primary access mechanism (typically "http", "https", or similar).
VDC	Virtual data center is a representation of resources within a Cloud Provider platform.

Multisite vCloud Director v8.20 and Earlier

Before exploring the new multisite capabilities in vCloud Director v9.0, it is worth looking at the way in which earlier versions of vCloud Director could be deployed in multiple locations to understand both the advantages and disadvantages of the most common deployment topologies.

2.1 vCloud Director Management Cluster Topologies

High availability of vCloud Director management components is achieved through the presence of multiple redundant "cells", each of which requires connectivity to a backend database, message bus, and other components. These components are typically installed on a "management" vSphere cluster, under the control of a separate VMware vCenter Server® dedicated solely to management workloads. Cloud Providers can deploy these management workloads on infrastructure in one or more locations. Examples are described briefly in the following sections.

2.1.1 Single-Site Management Cluster

In the simplest deployment topology, the management cluster which supports the vCloud Director components runs in a single site. The following figure shows a single data center within which resides a Management vCenter Server and its Management resource cluster which together host the components of a vCloud Director instance.

Figure 1. Single-Site vCloud Director Management Cluster

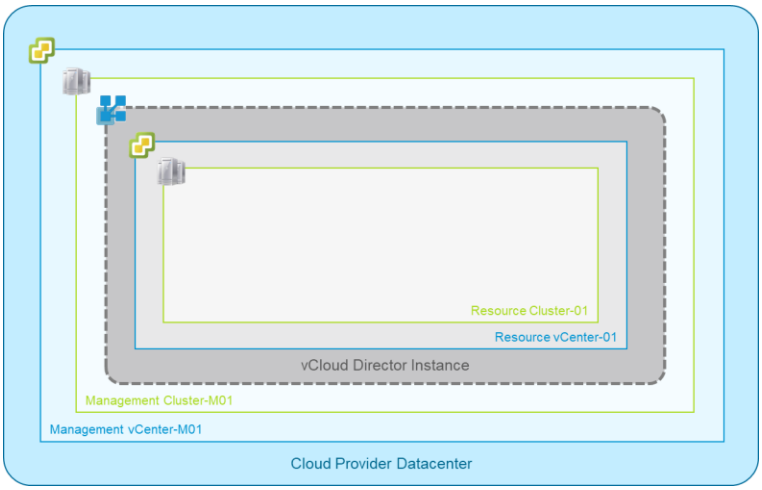

Although the tenant workload vCenter Server and associated resource cluster run on separate hardware to that of the management cluster, the two workloads are contained *inside* the vCloud Director instance in the figure to illustrate that they are managed by the surrounding vCloud Director.

2.1.2 Stretched Management Metro-Cluster

When a Cloud Provider has two data centers within close proximity to each other, it is possible to distribute the elements of the management environment across a single vCenter Server controlled resource platform which is itself distributed across the two data centers creating a single, logical data center. The following figure shows the vCloud Director instance deployed within such a management resource environment. The possible workload vCenter Server and resource cluster topologies are the subject of the following sections of this document, so are intentionally omitted.

Figure 2. Stretched Management Metro-Cluster vCloud Director

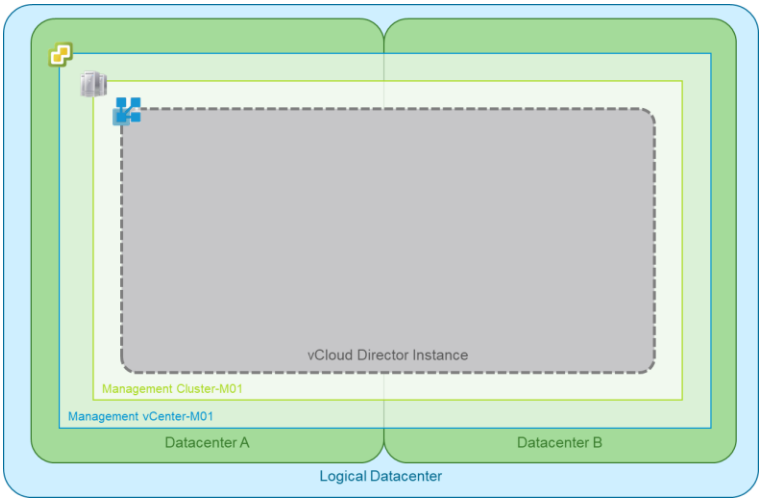

This figure shows the Management vCenter which controls resources in data center “A” and “B”. The two data centers are interconnected with sufficient bandwidth and low enough latency to enable a single compute cluster and its associated storage to operate across both. Typically, the elements of the vCloud Director instance run in one of the two data centers, and fails over to the other in the event of a failure within the active site. Some of the failover can be achieved using VMware vSphere High Availability. Other more complex elements (such as integration into external systems) might require a recovery automation tool such as VMware Site Recovery Manager™.

While it is, therefore, possible to stretch the vCloud Director management components between sites, the details of implementing a stretched management environment are out of scope for this document. Within the remainder of the document, “stretched” refers to the workload resources controlled by a single vCloud Director instance being deployed across multiple locations rather than the specific topology of the management components of vCloud Director itself. In subsequent figures, because the management resource topology is omitted, the vCloud Director instance is shown as the outermost container within, or across, data centers.

2.2 Single-Site vCloud Director

Until vCloud Director v9.0, there was no real concept of a site within the vCloud Director data model. Typically, Cloud Providers deployed an instance of vCloud Director in a single location, dedicating much of the underlying infrastructure to a single Provider VDC (PVDC) which was the closest representation to a site because it was usually constrained to a single location. The following figure shows a representation of the relationship between the elements of a single vCloud Director instance.

Figure 3. Single-Site vCloud Director Object Hierarchy

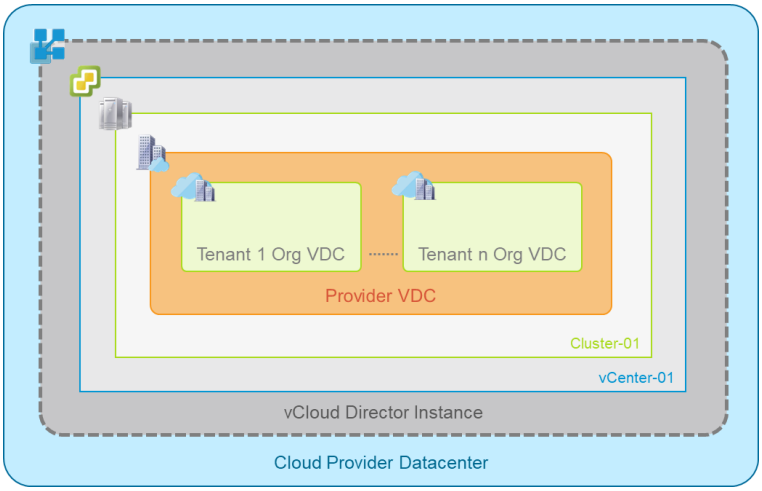

The Cloud Provider data center is the outer "container", within which the vCloud Director instance is deployed and has control of the (in this example) single VMware vCenter Server and its associated resources. Cluster-01 is allocated to a single PVDC within which individual Tenant Org VDCs are deployed. Additional PVDCs can be provisioned to present different types of compute nodes (perhaps with a different CPU or memory) within the same geographic location.

2.3 Dual-Site Stretched vCloud Director

As mentioned previously, some Cloud Providers with data centers located close together stretch the workload resource of a single vCloud Director instance across separate sites. Previously, a number of network "round-trip delay" configuration maximums constrained Cloud Providers' ability to stretch services across separate locations unless those locations were very close together. While stretching service across separate locations offers a degree of independence from a single location, it leaves a dependence upon either inter-site connectivity to bridge a "remote" site to resources in its "parent" location, or, complex networking and support designed to lessen the connectivity dependence. However, the advantage that stretching offers is the simplification of customer or management operation. Even though the underlying technology is more complex, the provider and customer gain the ability to manage services in multiple data centers as if they were in one location. VMware Cloud Providers have used two different topology models to achieve this, sometimes in combination. The two models are examined in the following sections.

2.3.1 Dual vCenter Server Instances – Stretched vCloud Director

In the first model, vCloud Director is installed in a single location, but manages resources under the control of VMware vCenter Server instances in each site. The maximum round-trip network latency from vCloud Director to vCenter Server, which had previously been 20ms, was increased to 40ms with v8.20 allowing one vCloud Director instance to manage workload vCenter Server instances in multiple locations. The following figure shows a representation of the relationships between the elements in this model. Tenant 1 has separate resources within Provider VDC "A" and "B". Those resources, while separate at the compute/storage layers, are connected at the network layer either by a common or shared vCloud Director network or through data center interconnects outside of vCloud Director.

Figure 4. Dual vCenter Server Stretched vCloud Director Object Hierarchy

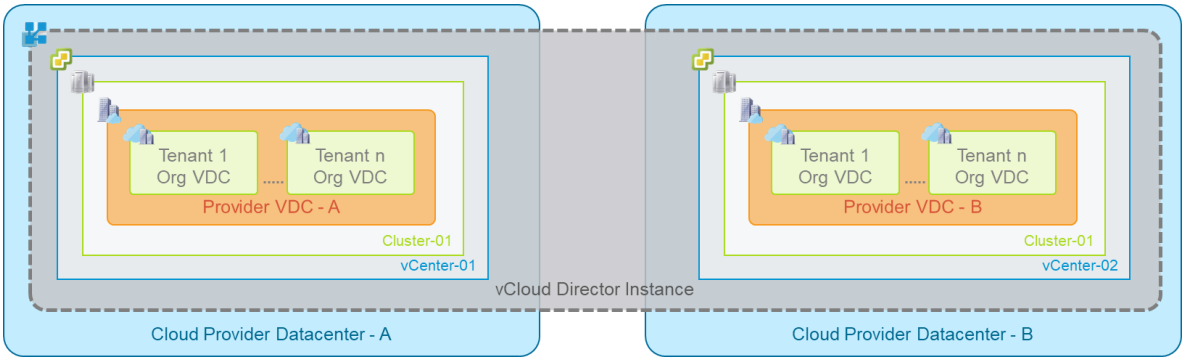

When using this model, Cloud Providers must rely on vCloud Director deployed on one of the sites, with a disaster recovery (DR) plan to recover the management platform to another site in the event of a failure. This DR plan mitigates against a complete site failure at the "primary" vCloud Director site leaving the other operational, but without a functioning customer or provider management portal.

In the event of a failure at the "secondary" or "remote" site, the remaining site operates as a single-site vCloud Director deployment with a single, local vCenter Server as shown in Figure 3. However, should the failure affect the primary site, while the remaining vCenter Server and its associated workload hosts remain unaffected at the secondary site, initially at least, there is no vCloud Director instance from which to control them. Recovering the vCloud Director to the remaining site is represented in the following figure.

Figure 5. Dual vCenter Server Stretched vCloud Director in Disaster Recovery

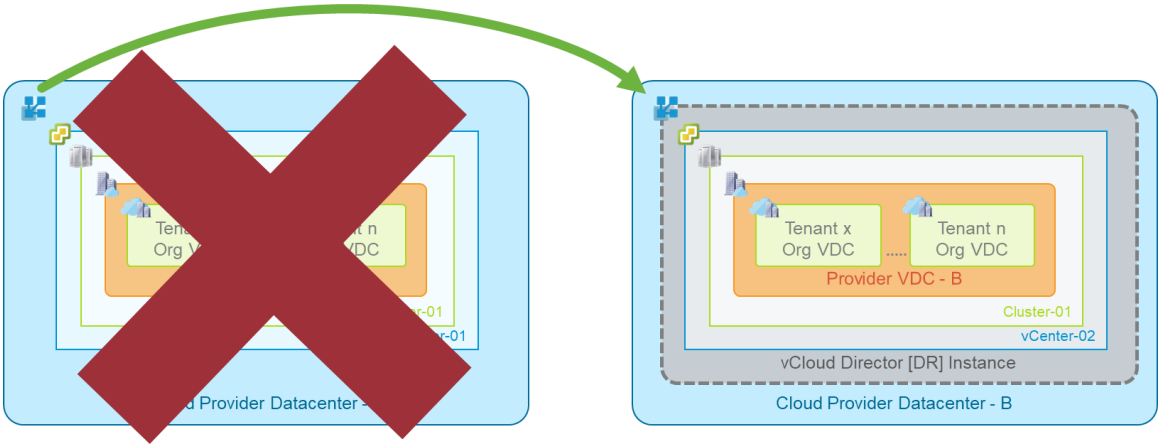

Recovery of the vCloud Director instance complete with vCloud Director cells, databases, message bus, and other supporting infrastructure, is not a trivial task. While VMware Site Recovery Manager can be used to automate much of this process, additional steps might be necessary to affect the network topology changes for the remaining vCenter Server and other Business or Operational Support Systems (BSS/OSS) to restore connectivity to the vCloud Director instance in its new location.

2.3.2 Single vCenter Server – Stretched vCloud Director

In the second model, vCloud Director is installed in a single location, but with managed resources under the control of a single vCenter Server, typically co-located within the same management environment as vCloud Director. In this model, the single vCenter server controls resources in both the local and remote sites. Where in the dual vCenter Server model, the factor limiting the distance between sites is the

network latency tolerance of vCloud Director, in this model, the limit is the network latency tolerance between vCenter Server and the resource hosts under its management. The following figure shows a representation of the relationships between the elements in this model.

Figure 6. Single vCenter Server Stretched vCloud Director Object Hierarchy

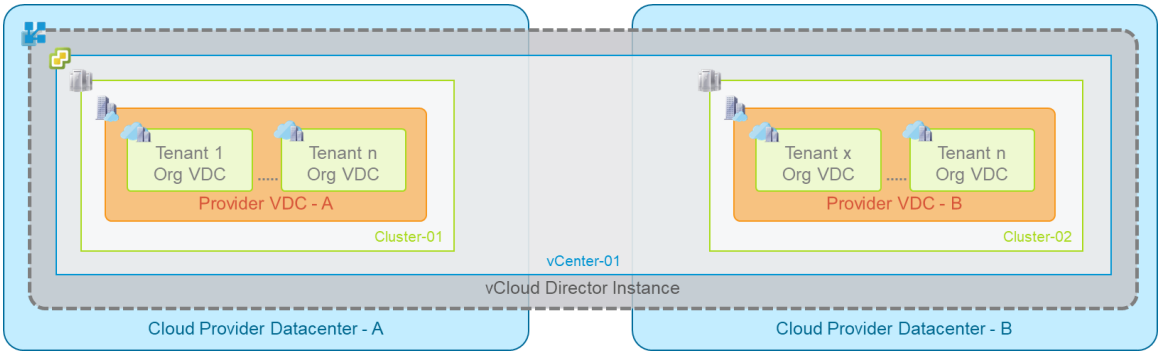

Although this model appears less complex than the dual vCenter Server model, the decrease in complexity due to the elimination of the second vCenter Server instance is more than made up for in the operational overheads in managing a stretched vCenter Server. Although beyond the scope of this document, there are considerations with regard to VMware vSphere design, storage design, and both virtual and physical network design. As with the previous model, vCloud Director is installed in a single site and Cloud Providers must once again rely on a DR plan to recover the platform to the secondary site in the event of a failure at the primary site. Unlike the previous model, if the primary site becomes unavailable, so too will the vCenter Server which manages the resources at both sites. Instead of needing to recover only the vCloud Director instance, in this model, the Cloud Provider must first recover the vCenter Server to the remote location. Recovering both vCenter Server and vCloud Director is shown in the following figure.

Figure 7. Single vCenter Server Stretched vCloud Director in Disaster Recovery

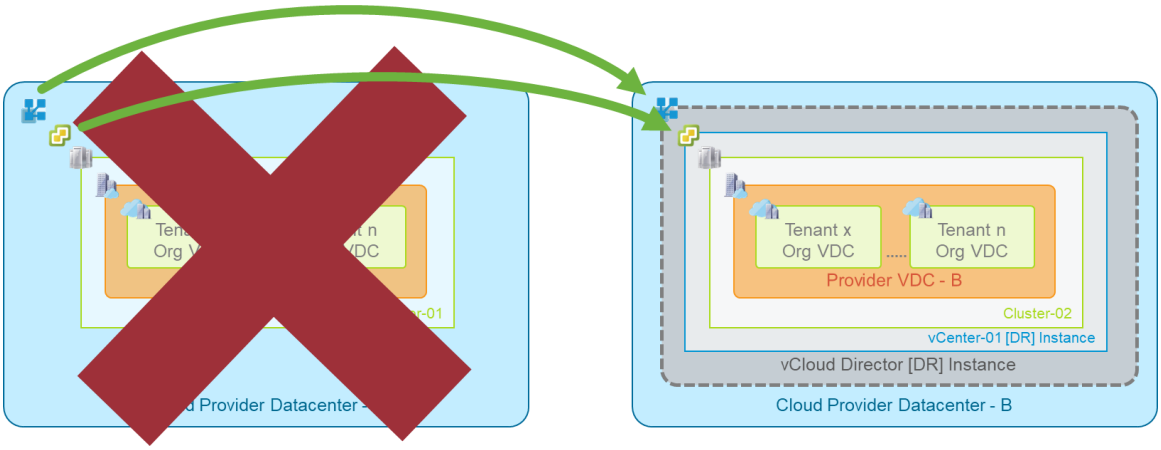

This model requires an even more complex recovery plan to restore vCloud Director and vCenter Server as well as the supporting infrastructure and network configuration. In addition, the topology which the vCenter Server manages (and the associated vCloud Director Provider VDC) has also changed, so accommodation must also be made while the platform is operating in “DR mode” so that customers cannot provision workloads to the failed or inaccessible resources.

2.4 The Benefits of a Stretched vCloud Director Solution

The two multisite stretched vCloud Director models both introduce complexity in day-to-day operations and in a business continuity disaster recovery (BCDR) planning. If Cloud Providers have gone to the trouble of implementing one or other of these models (and sometimes both), there must be a benefit to doing so either for the provider or their customer.

The dual vCenter Server model allows the provider to offer services to their customers across multiple physical locations which, subject to the network latency restriction, allows the customer to benefit from a degree of protection from a loss of service or access at a single provider data center. Depending upon the design of the network interconnectivity between them, customers can treat the two sites as two halves of a single solution, distributing elements of their service to each of the provider locations. With a DNS-based traffic management or load balancing solution in front, a customer solution can be designed to tolerate the loss of a single site without noticeable downtime, while taking advantage of synchronous replication of data between the two sites.

The single vCenter Server model offers similar benefits to those outlined above, with the advantage that, with resources managed from a single vCenter Server, Cloud Providers can offer enhanced replication, migration, or recovery services between sites through, for example, the deployment of a workload stretched metro-cluster. Because both the virtual machines and the virtual infrastructure which underpins them are under the control of a single vCenter Server, workloads moved between sites do so without the need to be “exported” from one vCenter Server and “imported” into another. Even so, careful consideration is required to make sure that the workloads can still be managed through vCloud Director after any migration or recovery.

The advantage that both models offer, beyond what has already been discussed, center on the ability of the customer to manage their vCloud Director Organization and services from a single portal. Separate vCloud Director instances at each provider location are independent of each other, requiring the provider to establish organizations and users and to carry out other administrative tasks at each site. The customer then logs in separately to each vCloud Director instance to manage services at that site. Similarly, automation tasks directed at the vCloud Director API must target the API endpoint at each location.

Despite the fact that Cloud Providers could offer these enhanced services with vCloud Director v8.20 and earlier, it forced a compromise in which the Cloud Provider had to trade a more convenient, simplified, and improved user experience offering for a more complex, difficult to deploy, manage, and support provider platform.

Multisite vCloud Director v9.0

vCloud Director v9.0 introduces more native multisite capabilities that allow a Cloud Provider to avoid the tradeoff between customer satisfaction and platform complexity. vCloud Director v8.20 introduced a new, more flexible HTML5-based interface to support the enhanced networking features of VMware NSX®. This has continued into vCloud Director v9.0 which has a new HTML5-based Tenant user interface. The previous Tenant user interface still exists and the Service Provider interface has not moved to HTML5 in this version, but to take advantage of the multisite capability, customers must use the new Tenant interface.

3.1 Multisite Concepts

vCloud Director v9.0 provides a number of new features, capabilities and concepts. Multisite capabilities allow an organization user to log in to the vCloud Director UI hosted at any of the sites where they have an Organization and Organization VDC. Upon login, the UI displays a sites icon that allows them to switch to other sites in which they have resources so that they can manage them from the same session. There are a number of foundational concepts which together create the new multisite model.

Table 1. Foundational Concepts

Term	Definition
Site	A single vCloud Director deployment that acts as a logical unit of a multisite installation.
Organization	A Tenant's presence on a site.
Member	A site or organization which is part of (that is, a member of) an association.
Association	A collection of members, either sites or organizations.

In previous versions of vCloud Director, the concept of a "site" has always been implicit. In vCloud Director v9.0, the site becomes an explicit part of the data model and a fully-fledged REST API object. In addition to sites, the other key component of multisite operation is association. vCloud Director contains a new API which can "associate" sites and organizations. The purpose of associating objects between sites is to establish a trust relationship between them. This trust relationship enables one object to make API calls on behalf of the user, against other members of the association. This forms the foundation of the multisite operation. The following sections describe the process of creating the associations so that customers can take advantage of the new capabilities.

3.2 Site Association

The first step in the process of creating a multisite vCloud Director service is for the Service Provider to create associations between sites. Association establishes a trust relationship so that one site can execute API calls at another site with pre-approved authorization. After an association is established between a pair of sites, a heartbeat process monitors the connection between the two member sites. Until each site is associated with the other, Site "A" to Site "B" and Site "B" to Site "A", the heartbeat process will show the association as incomplete. Establishing an association between two sites requires System Administrator permissions at both sites, because authorized API calls must be made to each site during the process.

3.2.1 Site Association Process

The high-level sequence of API calls to establish a unidirectional association is a two-part process. "Association data" which cryptographically identifies the source site Site "A" is collected first, then passed to an API endpoint which creates a Task within vCloud Director at Site "B", adding Site "A" as an associated site and storing both the cryptographic data provided during the association and the URLs of the included REST endpoints for Site "B". The following figure shows this process.

Figure 8. Unidirectional Site Association Sequence Site "A" to Site "B"

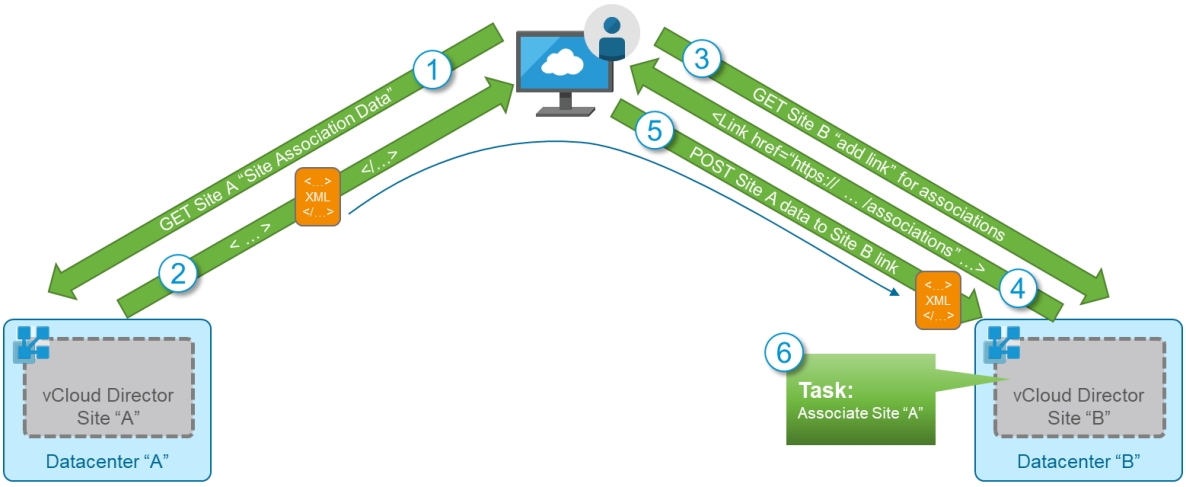

With the first part of the associated completed, vCloud Director at Site "B" has the credentials it requires to authenticate and decrypt API calls and payload data from Site "A", but no way to identify itself to Site "A" or, to encrypt API payload data in a way in which Site "A" will be able to reciprocate. The association sequence is then repeated in the opposite direction to establish bidirectional authentication and encryption. The following figure shows this process.

Figure 9. Unidirectional Site Association Sequence Site "B" to Site "A"

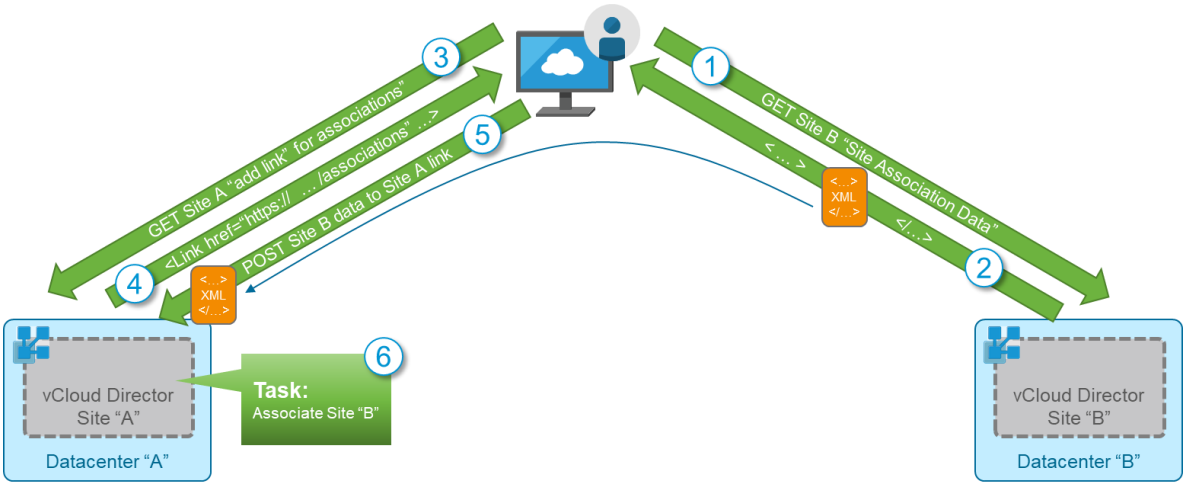

With the second part of the association complete, both members of the association are now aware of each other, have the URLs of the remote site's REST API endpoints, and have the credentials to be able to use them securely. The API endpoints in the current version (v29.0 at the time of writing) allow the administrator to collect the association data from one vCloud Director site and submit it to another (as shown above) using the following API workflow.

The requests and responses show only the relevant elements. See the vCloud Director API guide in the References section for full details. The workflow shows the sequence of steps, the requests (➔) and their relevant headers, and the responses (🡐) and their relevant content.

3.2.1.1 Site Association Workflow

Site A – https://Site-A.example.com		
1.	Establish the login Uniform Resource Locator (URL) for required API version.	
	➔	`GET https://Site-A.example.com/api/versions`
	➔	`<Version>29.0</Version>` `<LoginUrl>https://Site-A.example.com/api/sessions</LoginUrl>`
2.	Create an authenticated login session.	

	➔	`POST https://Site-A.example.com/api/sessions`
		`Accept: application/*;version=29.0`
		`Authorization: Basic <administrator@system> <password>`

	➔	`x-vcloud-authorization →460f7bbe33b2453fa93e96bc11b2ee5d`
3.	Retrieve the site association URL from the Site object.	

	➔	`GET https://Site-A.example.com/api/site`
		`Accept: application/*;version=29.0`
		`x-vcloud-authorization: 460f7bbe33b2453fa93e96bc11b2ee5d`
		(these headers are present in future requests to Site-A but will be omitted for clarity)

	➔	`<Site`
		`<Link`
		`rel="down"`
		`href="https://Site-A.example.com/api/site/associations" … />`
		`</Site>`

4.	Retrieve the site association data URL from the Site Associations object.	
	➔	`GET https://Site-A.example.com/api/site/associations`

	➔	`<SiteAssociations`
		`<Link`
		`rel="down"`
		`href="https://Site-A.example.com/api/site/associations/localAssociationData" … />`
		`…`
		`</SiteAssociations>`

5.	Retrieve the site association data from the SiteAssociationMember element.	
	➔	`GET https://Site-A.example.com/api/site/associations/localAssociationData`

	➔	`<SiteAssociationMember … >`
		`…`
		`<RestEndpoint>https://Site-A.example.com/api</RestEndpoint>`
		`<RestEndpointCertificate>-----BEGIN CERTIFICATE-----`
		`MIIDDjCCAfagAwIBAgIJAPVFVZ64w… -----END CERTIFICATE-----`
		`</RestEndpointCertificate>`
		`<SiteId>urn:vcloud:site:5e4b381d-dabb-49d9-9352-05f7ac0be7f5</SiteId>`

		<SiteName>5e4b381d-dabb-49d9-9352-05f7ac0be7f5</SiteName> <PublicKey>-----BEGIN PUBLIC KEY----- MIIBIjANBgkqhki … QIDAQAB -----END PUBLIC KEY----- </PublicKey> </SiteAssociationMember>
Site B – https://Site-B.example.com		
6.	Repeat steps 1-3 using the Site-B URL to establish an authenticated session and retrieve the site association URL from the Site object.	
7.	Retrieve the “add associations” URL from the Site Associations object.	
	➔	GET https://Site-B.example.com/api/site/associations Accept: application/*;version=29.0 x-vcloud-authorization: 9af7e9762efc435daceba456d28261c2 (*these headers are present in future requests to Site-B but will be omitted for clarity*)
	➔	<SiteAssociations …> <Link rel="add" href="https://Site-B.example.com/api/site/associations" type="application/vnd.vmware.admin.siteAssociation+xml" … /> </SiteAssociations>
8.	Post the SiteAssociationMember response body received in Step 5 to the “add associations” URL from Step 7.	
	➔	POST https://Site-B.example.com/api/site/associations Content-type: application/vnd.vmware.admin.siteAssociation+xml <SiteAssociationMember … > … <RestEndpoint>https://Site-A.example.com/api</RestEndpoint> <RestEndpointCertificate>-----BEGIN CERTIFICATE----- MIIDDjCCAfagAwIBAgIJAPVFVZ64w… -----END CERTIFICATE----- </RestEndpointCertificate> <SiteId>urn:vcloud:site:5e4b381d-dabb-49d9-9352-05f7ac0be7f5</SiteId> <SiteName>5e4b381d-dabb-49d9-9352-05f7ac0be7f5</SiteName> <PublicKey>-----BEGIN PUBLIC KEY----- MIIBIjANBgkqhki … QIDAQAB -----END PUBLIC KEY----- </PublicKey> </SiteAssociationMember>
	➔	<Task … operationName="siteUpdate" > <User href="https://Site-B.example.com/api/admin/user/..." name="administrator" type="application/vnd.vmware.admin.user+xml"/> <Organization href="https://Site-B.example.com/api/org/... " name="System" type="application/vnd.vmware.vcloud.org+xml"/> </Task>

As noted in the previous figures, this results in a Task within vCloud Director at Site B to add the association with Site A. When this process is complete, it can then be reversed, retrieving the SiteAssociationMember data from the localSiteAssociation URL of Site B and POSTing it to the “add

association" URL of Site A. After both unidirectional associations are complete, the two sites are associated and the site-to-site heartbeat is established.

After sites are associated, the login request must be modified to request authorization tokens for all associated sites. This is accomplished by adding an extra parameter "multisite=global" to the Accept header of the standard login request. The modified version of the standard login request in step 2 above then becomes the following:

	Create an authenticated multisite login session.	

	➔	`POST https://Site-A.example.com/api/sessions`
		`Accept: application/*;version=29.0;multisite=global`
		`Authorization: Basic <administrator@system> <password>`

	➔	`x-vcloud-authorization` `→9cd477276e5d444896fca73ffb7ed5c0,2d95e0b896a9419ca7e63e19bbd0ed91`

3.2.2 Site Association Mesh

The process described in the previous section established a full, bidirectional association between two sites. If the Cloud Provider wants to include additional member sites in the association, the new sites must be associated with the existing member sites. For predictable operation, VMware recommends that all members are associated with each other within the association. This enables the organization associations discussed in the following sections to be established between any combinations of sites within the group. Associating a third site with the two sites illustrated earlier results in six unidirectional associations as shown in the following figure.

Figure 10. Association Mesh Between Three Member Sites

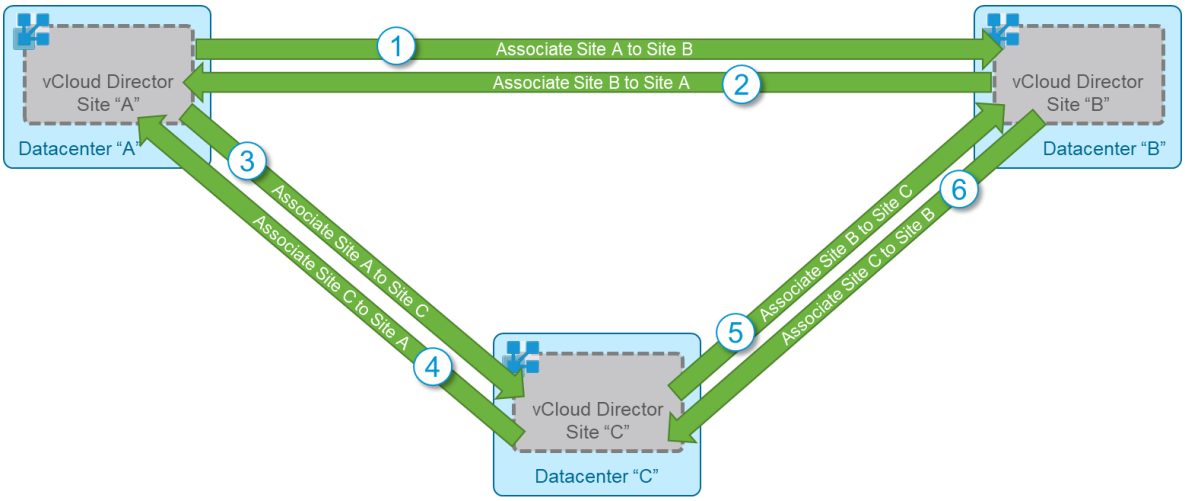

3.3 Organization Association

After the site-level associations are shown as completed, individual organizations can be associated with their peers in the other member sites. In the Cloud Service Provider environment, it is likely that the member organizations at each site will belong to the same customer, but this does not have to be the case, and vCloud Director does not require any explicit connection between the administrative bodies behind the member organizations. However, to associate organizations, either System Administrator or

Organization Administrator credentials are required for each site or organization to be associated, which implies common administration on the part of either the customer or Service Provider.

3.3.1 Organization Association Process

The organization association process is similar to that used to associate sites. Once again, the process requires credentials to be collected from one vCloud Director and passed to the other after which the process is reversed to complete the association. However in this case, the association data cryptographically identifies the Organization rather than the site. The high-level sequence of API calls to establish a unidirectional association in each direction is illustrated in the following figures.

Figure 11. Unidirectional Organization Association Sequence Org "A1" to Org "B1"

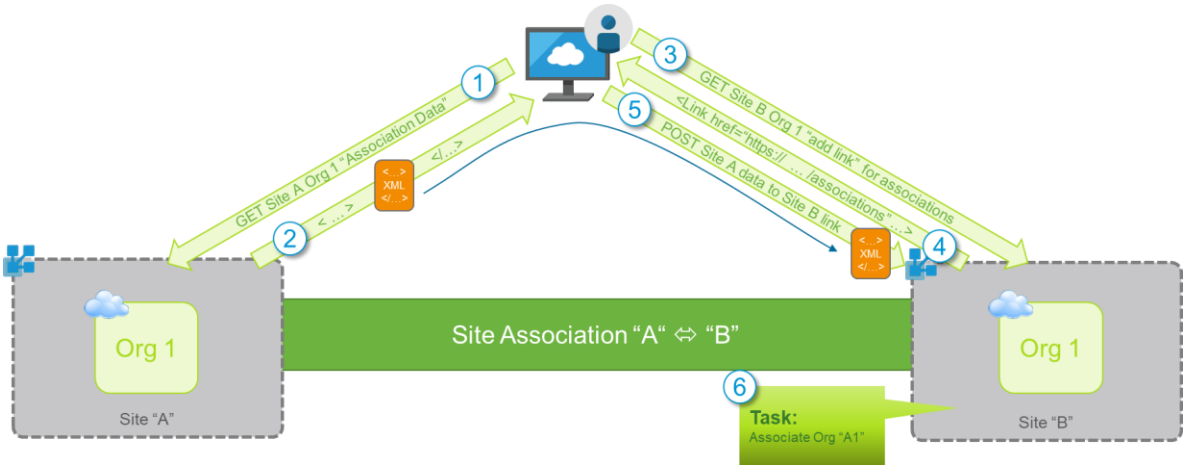

With this complete, “Org 1” at Site “B” has the credentials it requires to authenticate and decrypt API calls and payload data from “Org 1” at Site “A”, but no way to identify itself to “Org 1” at Site “A” or to encrypt API payload data in a way in which “Org 1” at Site “A” will be able to reciprocate. The association sequence is then repeated in the opposite direction to establish bidirectional authentication and encryption. The following figure shows this process.

Figure 12. Unidirectional Organization Association Sequence Org "B1" to Org "A1"

With the second part of the association complete, both member organizations are now aware of each other, have the URLs of the remote organization's REST API endpoints, and have the credentials to be able to use them securely. The API endpoints in the current version (v29.0 at the time of writing) allow the system or organization administrator to collect the association data from one vCloud Director organization and submit it to another using the following API workflow. The requests and responses show only the relevant elements. See the vCloud Director API guide in the References section for full details. The workflow shows the sequence of steps, the requests (➔) and their relevant headers, and the responses (🡐) and their relevant content. While the workflow can, as noted, be carried out by a user with Organization Administrator credentials, the following example shows the process carried out by a System Administration user and omits the session login steps illustrated in the site association workflow.

3.3.1.1 Organization Association Workflow

Site A – https://Site-A.example.com		
1.	Retrieve the ID of the organization to be associated.	
	➔	`GET https://Site-A.example.com/api/org`

	🡐	`<OrgList … >`
		`  <Org`
		`  href="https://Site-A.example.com/api/org/{ID}"`
		`  name="TestOrg" … "/>`
		`  …`
		`</OrgList>`

2.	Retrieve the organization association URL from the Org object.	
	➔	`GET https://Site-A.example.com/api/admin/org/{ID}`

	🡐	`<AdminOrg name="TestOrg" ... >`
		`  ...`
		`  <Link rel="down"`
		`  href="https://Site-A.example.com/api/admin/org/{ID}/associations"`
		`... />`
		`  ...`

		`</AdminOrg>`
3.	Retrieve the organization association data URL from the Org Associations object.	
	➔	`GET https://Site-A.example.com/api/admin/org/{ID}/associations`

	➔	`<OrgAssociations`
		`  …`
		`  <Link rel="down"`
		`  href="Site-`
		`A.example.com/api/admin/org/{ID}/associations/localAssociationData"…`
		`  />`
		`</OrgAssociations>`

4.	Retrieve the organization association data from the OrgAssociationMember element.	
	➔	`GET https://Site-` `A.example.com/api/admin/org/{ID}/associations/localAssociationData`

	➔	`<OrgAssociationMember …>`
		`  <Link … />`
		`  <SiteId>urn:vcloud:site:79423eb2-0983-454e-88b2-71c9ffcd4c5e</SiteId>`
		`  <OrgId>urn:vcloud:org:93a6def0-85ba-447c-b4be-453cf60854e1</OrgId>`
		`  <OrgName>TestOrg</OrgName>`
		`  <OrgPublicKey>-----BEGIN PUBLIC KEY-----`
		`  MIIBIjANBgkqhki…VybCa7wIDAQAB -----END PUBLIC KEY-----</OrgPublicKey>`
		`</OrgAssociationMember>`

Site B – https://Site-B.example.com		
5.	Repeat steps 1 and 2 using the Site-B URL to establish an authenticated session and retrieve the org ID and its association URL from the Org object	
6.	Retrieve the “add associations” URL from the Org Associations object.	
	➔	`GET https://Site-B.example.com/api/admin/org/{ID}/associations`

	➔	`<OrgAssociations …>`
		`  <Link`
		`  rel="add"`
		`  href=https://Site-B.example.com/api/admin/org/{ID}/associations`
		`  type="application/vnd.vmware.admin.organizationAssociation+xml"/>`
		`</OrgAssociations>`

7.	Post the OrgAssociationMember response body received in Step 4 to the “add associations” URL from step 6.	

	➔	`POST https://Site-B.example.com/api/admin/org/{ID}/associations`
		`Content-Type: application/vnd.vmware.admin.organizationAssociation+xml`
		`<OrgAssociationMember …>`
		`  <Link … />`
		`  <SiteId>urn:vcloud:site:79423eb2-0983-454e-88b2-71c9ffcd4c5e</SiteId>`
		`  <OrgId>urn:vcloud:org:93a6def0-85ba-447c-b4be-453cf60854e1</OrgId>`

		<OrgName>TestOrg</OrgName> <OrgPublicKey>-----BEGIN PUBLIC KEY----- MIIBIjANBgkqhki…VybCa7wIDAQAB -----END PUBLIC KEY-----</OrgPublicKey> </OrgAssociationMember>
	←	<Task … operation="Creating association Organization TestOrg ... " operationName="orgAddAssociation" … > <Owner href="https://Site-B.example.com/api/admin/org/{ID}" name="STFTestOrg" … /> <User href="https://Site-B.example.com/api/admin/user/{ID}" name="administrator" … /> <Organization href="https://Site-B.example.com/api/org/{ID} " name="System" … /> … </Task>

As noted in the previous figures, this results in a Task within vCloud Director at Site B to add the org association with the organization from Site A. When this process is complete, it can then be reversed, retrieving the OrgAssociationMember data from the localAssociationData URL of the Org in Site B and POSTing it to "add associations" URL of the Org in Site A.

3.3.2 Organization Association Using the GUI

It is quite likely that Cloud Service Providers will want to automate much of the provisioning activity associated with onboarding a new tenant, or a tenant to a new site. The organization association process must therefore be scripted using the API flow outlined above. However, it is also possible that the Provider will choose to leave the association process to the customer or might, if they offer multisite association as part of a premium product offering, need to carry out the process as a day-2 activity at some point later in the Tenant lifecycle. Even though the API process can be used to create the association at any time, it might be more convenient for the Provider or customer to be able to complete the association using a GUI in place of the API.

The Administration page of the vCloud Director HTML5 Tenant UI includes a multisite tab which shows current associations and allows the user to set up new associations. In much the same way as the API process, this requires the creation of two unidirectional associations, each formed by retrieving the association data from one organization and sending it to the other. These two actions are represented in the GUI as shown in the following figure.

Figure 13. Organization Association Through the HTML5 UI

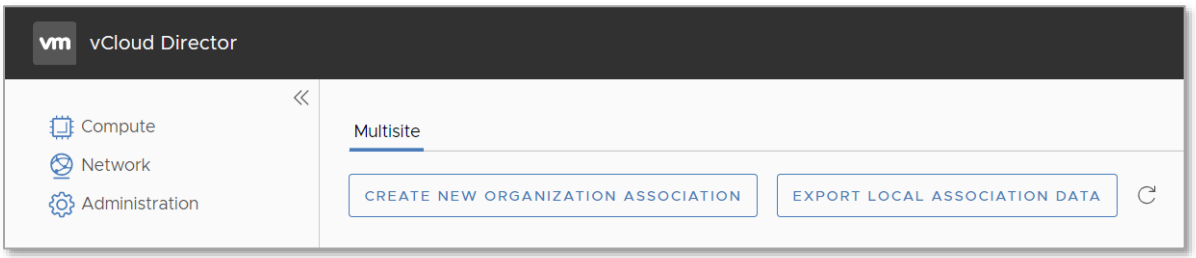

The **Export Local Association Data** option sends the organization's localAssociationData as an XML file for the user to save locally. The **Create New Organization Association** option opens a dialogue allowing the user to submit the localAssociationData from the remote organization. The dialog is shown in the following figure.

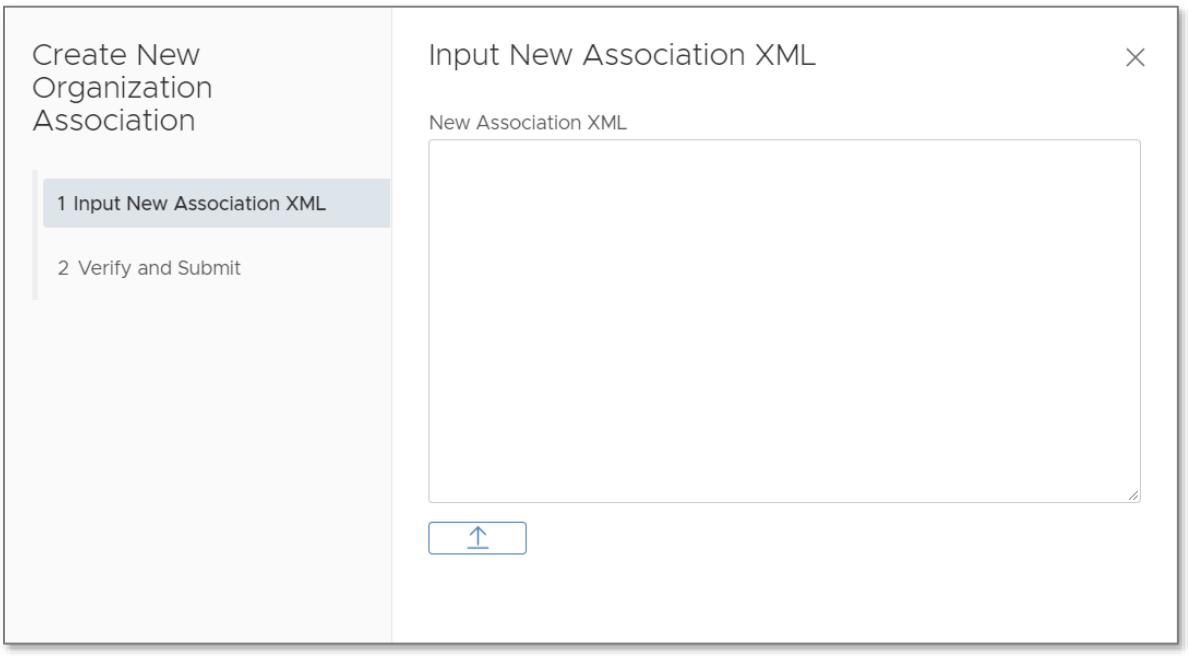

The dialog allows the user to either upload the XML file (previously downloaded from the other site to be associated) or simply paste its contents into the XML field directly. Once verified, the dialog allows the user to submit the association data into the local vCloud Director. This process is then repeated within the GUI of the other site to be associated, completing the bidirectional association. After completion, the association is then shown in the Multisite tab in both vCloud Director sites. The following figure shows the association between TestOrg locally in Site-A and its remote peer at Site-B.

Figure 14. Active Organization Association Displayed in the HTML5 Tenant GUI

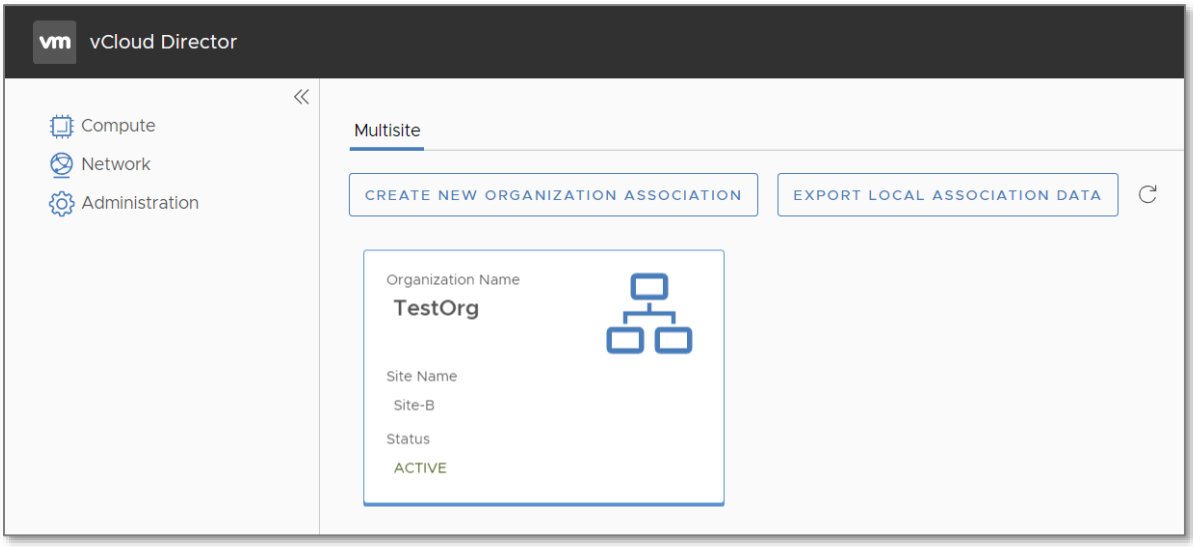

By clicking on the association status panel, the user can see the details of the association data and, if they want to, delete the local side of the association. The following figure shows this.

Figure 15. Organization Association Details Dialogue

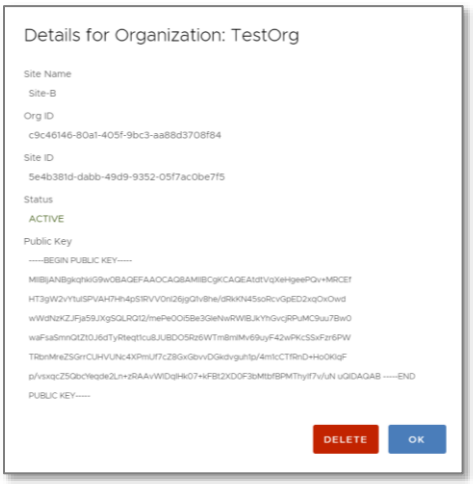

3.3.3 Organization Association Mesh

As with the site association process, the organization process establishes a full, bidirectional association between member organizations at two sites. If the Cloud Provider wants to include additional member organizations (at other member sites) in the association, the new organizations must be associated with the other sites/organizations in the existing group. Associating a third organization with the two illustrated earlier results in six unidirectional associations as shown in the following figure.

Figure 16. Association Mesh Between Organizations at Three Sites

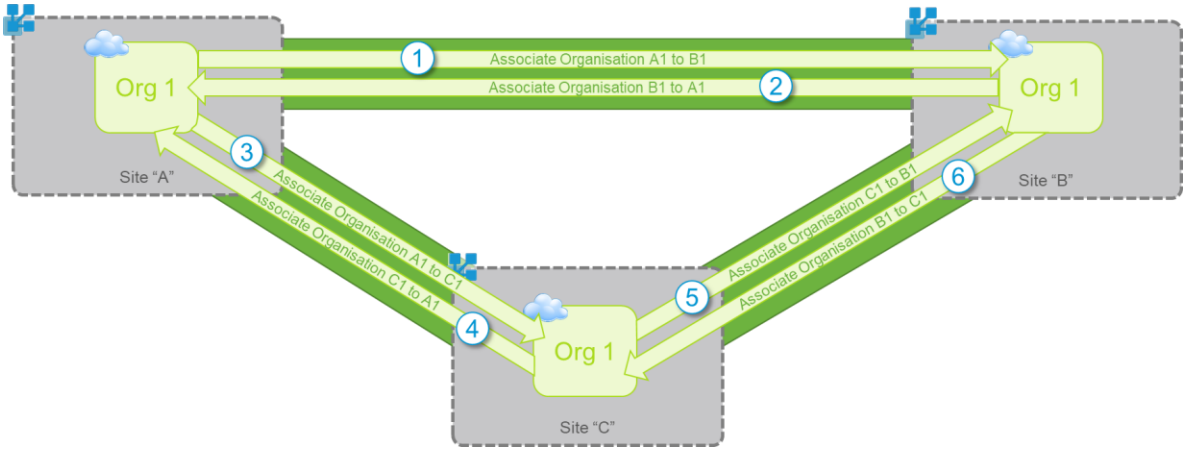

3.4 Multisite Tenant User Interface

With vCloud Director v9.0, end users can log into the vCloud Director environment at any site within which there is an organization where they have an user account. When a user logs into the new, HTML5 user interface of an organization that does not have associations at any other sites, the right-hand side of the toolbar at the top of the screen shows the Tasks and Help/About icons as well as the User/Role area containing the session logout menu option as shown in the following figure.

Figure 17. HTML5 UI Toolbar of a Single-Site User

However, when a user logs in to an organization that is a member of an association, the user additionally see the Sites icon.

Figure 18. HTML5 UI Toolbar of a Multisite User

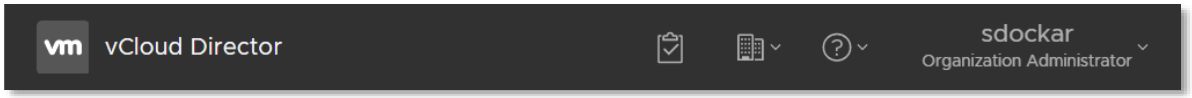

When clicked, the Sites icon activates a drop-down list of associated Orgs (excluding the current Site/Org) which the user can switch to as shown in the following figure.

Figure 19. Switching Sites with the HTML5 UI Toolbar

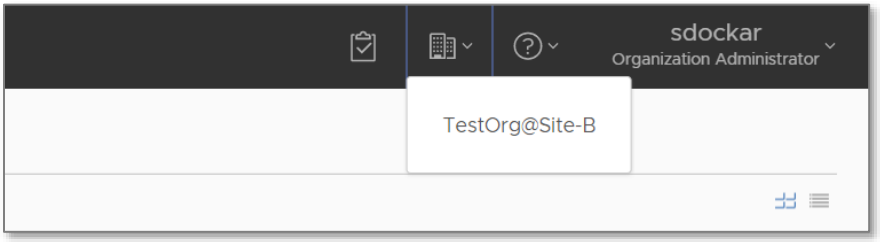

User Access to a Multisite vCloud Director UI

With the introduction of vCloud Director v9.0, users can access the portal interface in different ways. This section examines a number of them both from the user perspective and from that of the underlying technology which facilitates them. The URL format to access the new HTML5 UI is the site's fully qualified domain name (FQDN) followed by the organization identifier which (while configurable) defaults to the format "https://<FQDN>/tenant/org-id". For the examples used throughout this section of the document, that is "https://<Site>.cloud.example.com/tenant/testorg". For brevity, although "/tenant/testorg" is part of the complete URL, it is omitted from the following examples.

4.1 Direct Site Access

Users can log into any vCloud Director site in which there is an organization where they have a user account. This does not change even if the organization is an association member. If the user wants to make changes to the service at a single site they can log in through their tenant URL at that site. If the user needs to make changes in other member sites, they can still log into them separately, authenticating each time.

Because the vCloud Director user interface utilizes certificates for integrity and confidentiality, those certificates must contain the site's FQDN. While this is not an uncommon practice, it can become complicated when there are other access routes into the service.

4.2 Switching Between Associated Sites

When a user logs into an organization at a site to which they have access, their credentials are checked against the identity management source configured for their account. In vCloud Director v9.0, if the organization they log into is associated with others, the UI will render the Sites icon to allow them to switch to other sites with which their organization has existing associations. When the user switches sites, their browser session is redirected to the chosen site without the need to re-authenticate the new session. The sequence of a successful login to Site "A" followed by a switch to Site "B" is shown in the following figures.

Figure 20. Logging in to an Associated Organization

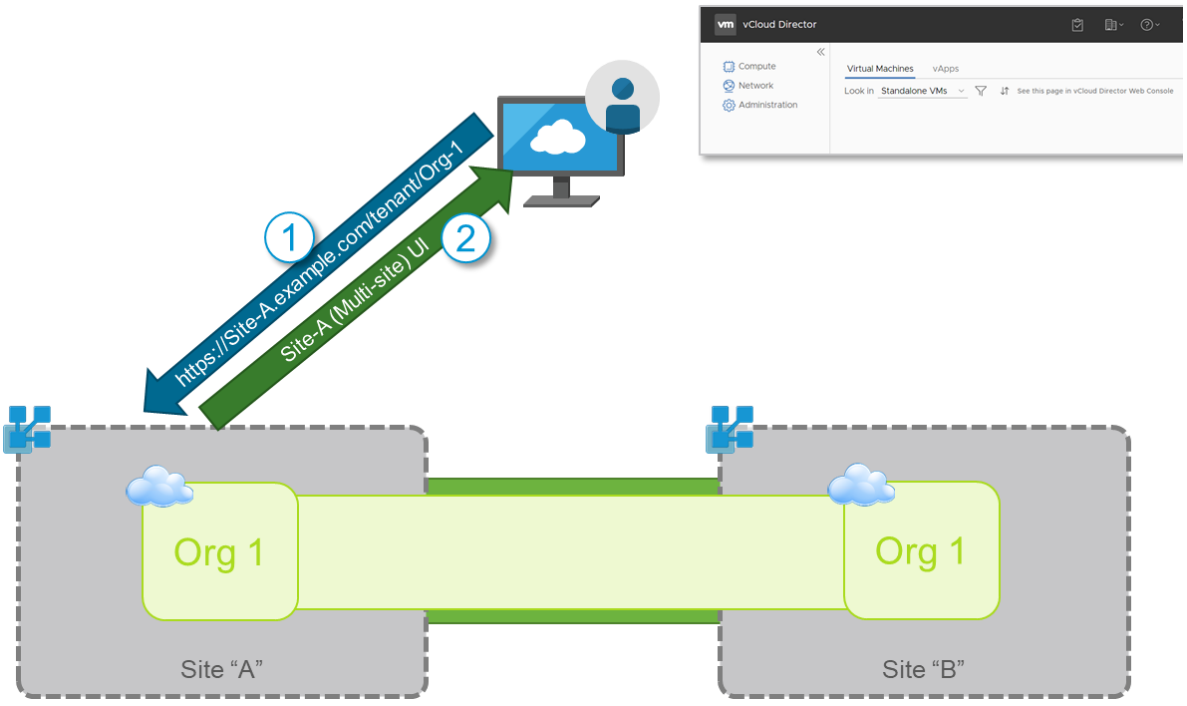

After the user has successfully logged in to Site "A" they can use the Sites menu to switch to another site.

Figure 21. Switching to a Remote Organization at an Associated Site

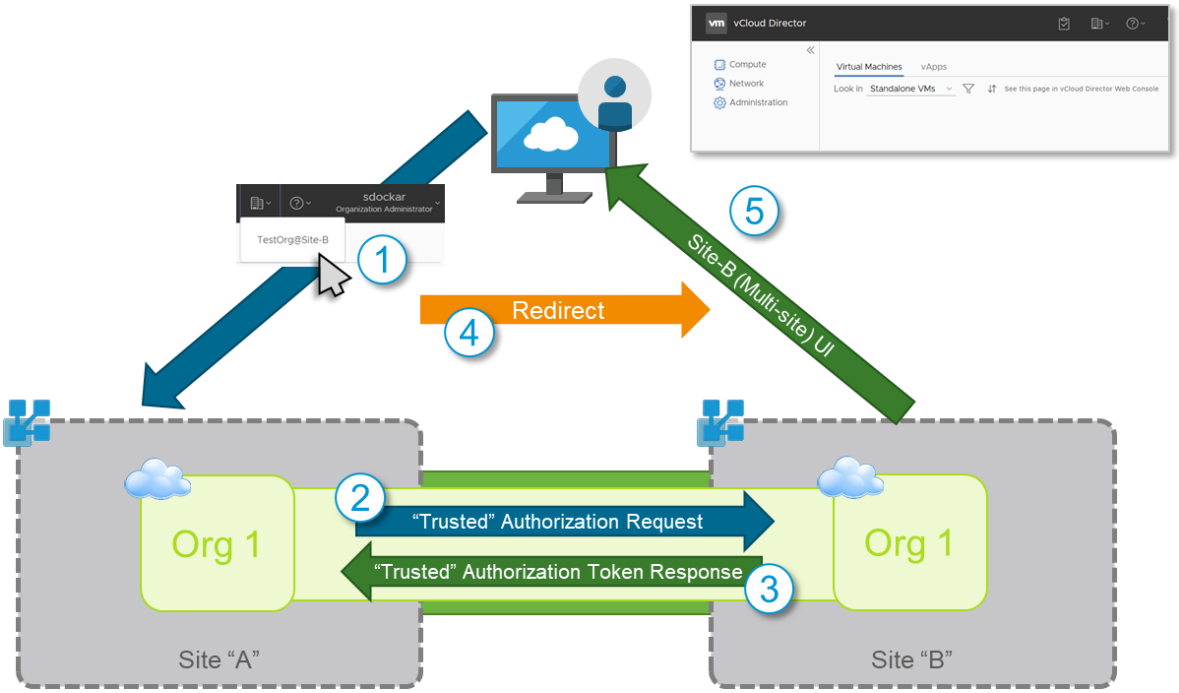

When the user selects a remote organization/site from the Sites menu, their browser session is redirected to the remote vCloud Director instance. The new session is preauthenticated using credentials retrieved from the target site over the organization association.

Note In the current version, login validation to the other member organizations does not remove sites from the drop-down list in the UI for locations where the user does not have authorized access. This means that the user can "switch" to the inaccessible sites, but doing so will leave the browser session at the Logout page *for the newly selected* site. Attempting to log back in will fail until the user establishes a new connection to a site to which they do have access.

4.3 Global Site Access

Now that a user can switch between different sites without the need to reauthenticate, making changes across a number of sites becomes more straightforward. A user can login to the first site in which they want to make changes, then switch to the second and so on until they have completed the changes in each site. However, because it is possible for users to log into any vCloud Director site as long as they have a user account within an organization there, vCloud Director v9.0 offers Service Providers the chance to provide a more resilient access mechanism across their associated sites. While customers can still log in at any site, in the event of a failure at the site they are attempting to log in to, their session will fail. With the introduction of a load sharing or load balancing mechanism in front of the users' login to a particular site, service providers can offer a higher availability service level against the associated sites. It should be noted that, even though using the techniques described in the following sections offers increased availability, a failure at an associated site will potentially prevent access to the site even if a user is able to log in to another site and then attempts to switch to the failed site.

For the purposes of this document, the ability to distribute login sessions across multiple sites is split into two models. The first involves those options which require an HTTP(S) connection to be made from the user's browser to a service location and the second involves those options which rely on intelligent DNS

service to steer the user's initial connection to the required service location. Both models appear similar at a high level. The following figure illustrates the conceptual traffic flow for both.

Figure 22. Global Site Access Conceptual Overview

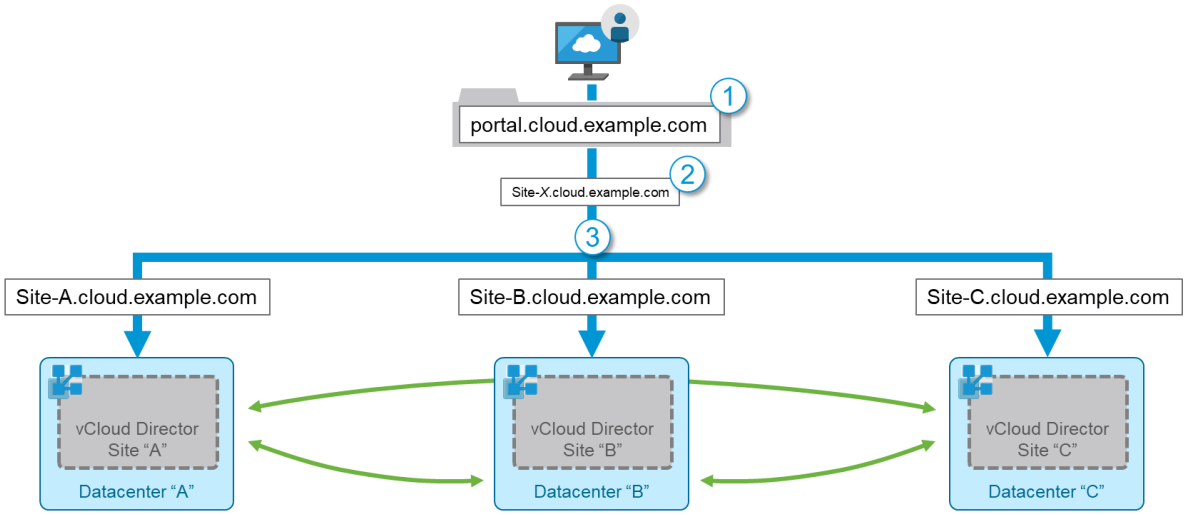

To access one of the associated organizations, the user connects to a single, global URL (1) which is "https://portal.cloud.example.com/tenant/orgname". The request is directed to one of the service location sites accessed through the site's "Site-*X*.cloud.example.com" host name and IP address (2). The site that receives the connection sends back the login page for the user to enter their credentials (3).

4.3.1 Global Site Access with Traffic Load Sharing

Connections to the global URI ("portal.cloud.example.com" in this example) are delivered to the IP address resolved by a DNS query for the FQDN portion of the full tenant interface URL. After a connection is established, the terminating equipment advises the client (the user's browser in this case) that the resource it requested has moved. The response also includes an alternate location at which the resource can now be found, and this results in the client repeating the connection process to the newly received location. The following figure illustrates this process.

Figure 23. Global Site Access Using Web Server Redirection

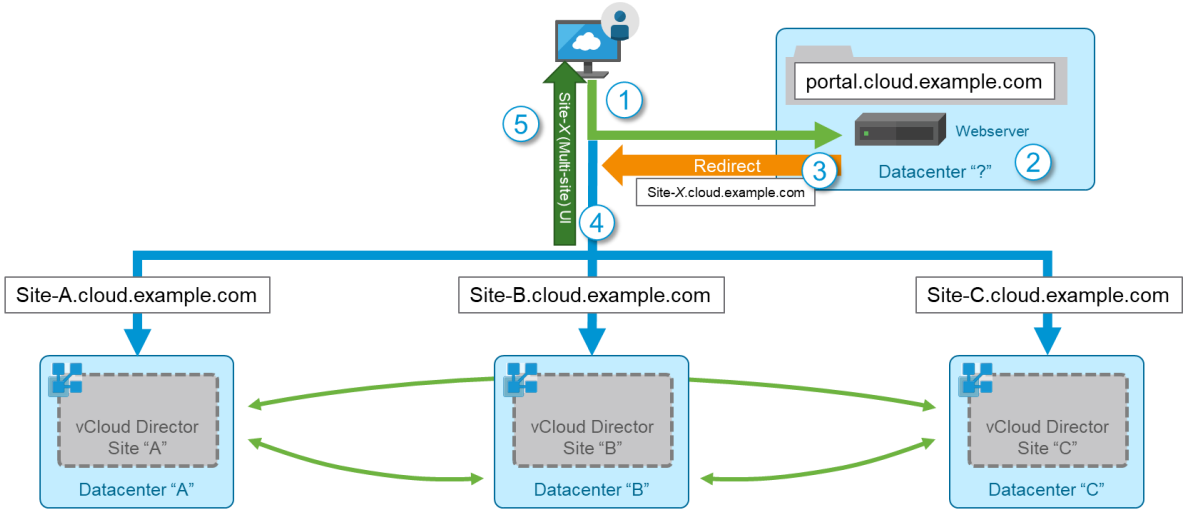

1. The user enters the global URL into their browser and, following DNS resolution, a connection is made to a server.
2. The server uses a standalone algorithm (random, round-robin or similar), or might use interaction with the vCloud Director sites, to determine the load at each site and then selects the site to which this connection should be sent.
3. The server issues an HTTP 3xx redirection response to the client and includes the site-specific URL of the chosen vCloud Director instance.
4. The client makes a new connection to the site contained within the redirect.
5. The selected site receives the new connection from the client and responds with the login page.

This approach has the benefit of using fairly common and well understood technology, but has the disadvantage that without additional complexity, the server providing the redirection runs in a single location. While the infrastructure providing the redirection can be made resilient, the location can form a single point of failure. It is therefore recommended that customers are provided with both the site-specific and global URLs so that in the event of a failure they can connect directly to any surviving sites.

In this model, the user's initial connection to "https://portal.cloud.example.com/" terminates on the service which will redirect the connection to the target site. This service must contain an SSL certificate which is valid for the "portal" FQDN. When the client is told that the resource has moved and it should instead connect to the site-specific URL, its new connection will be to "https://site-*X*.cloud.example.org" and the service terminating the connection (typically a local load balancer in front of the vCloud Director cells) must contain an SSL certificate which is valid for the "Site-*X*" FQDN. If the same local load balancer is responsible for both "global" site redirection and local site termination, it must hold certificates which are valid for both FQDNs. If it presents the redirection and local site on separate IP addresses, separate SSL certificates, each valid for one of the FQDNs can be used. If the same IP address (or Virtual IP "VIP") is used for both services, it must hold an SSL certificate that is valid for both names. See the following section for more information.

4.3.2 Global Site Access with DNS Load Balancing

An alternative to the traffic-based load sharing model uses intelligent DNS resolution to distribute connections across multiple vCloud Director instances. In the previous section, connections to the "portal.cloud.example.com" URL only reached the redirection server as a result of a DNS query/resolution. Typically, a DNS query will return a single IP address which will be used to connect to the target system. Some DNS queries will return multiple answers and the querying application or service will then decide how to select the one (or more) that it will use. A third category of DNS responses, mostly invisible to the user, are those which apply some form of intelligence to the response they select. This can take the form of identifying the approximate location of the user making the request, and returning the IP address of a more local server to minimize latency and maximize responsiveness. Alternatively, the intelligent DNS service might be able to interrogate each target site to identify the one currently returning the fastest responses, or the one with the lowest number of connections. Or, at its simplest level, the intelligent DNS service might only have a list of sites and cycle through them, returning each in turn in a round-robin method.

The specifics of the method used are beyond the scope of this document, but the basic principle is shown in the following figure.

Figure 24. Global Site Access Using DNS-Based Load Balancing

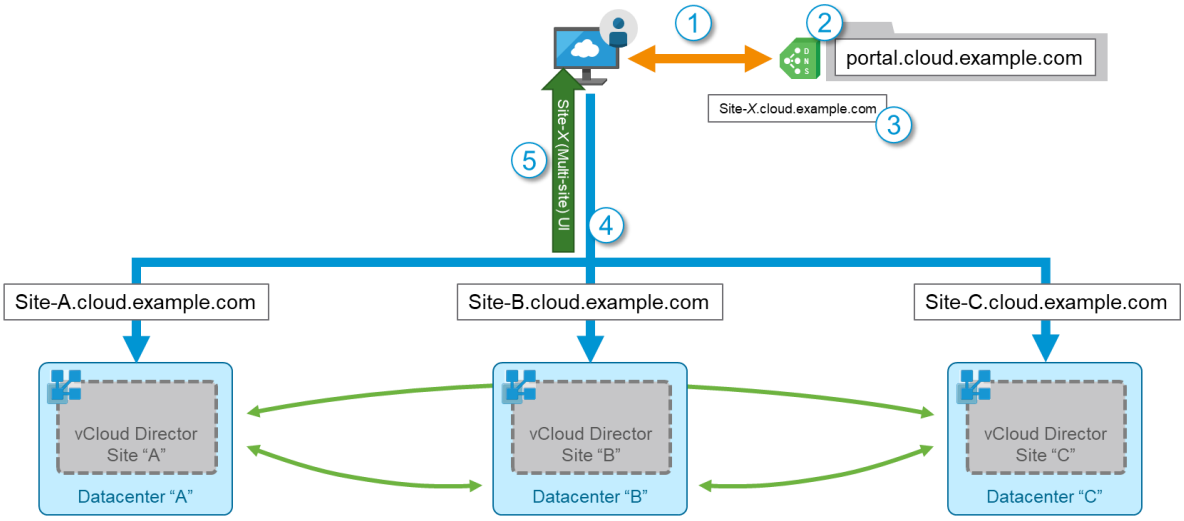

1. The user enters the global URL into their browser client which results in a DNS query.
2. The intelligent DNS application chooses a vCloud Director instance based upon its selection algorithm.
3. The chosen site is returned to the client within the DNS response.
4. The client makes a connection request to the selected site.
5. The selected site receives the new connection from the client and responds with the login page.

Unless already present within the Provider environment, this approach requires the deployment of intelligent DNS applications, ideally at multiple, resilient locations. Those applications might require connectivity to each other as well as to the target sites to better coordinate their responses. This represents a potential increase in both up-front and operational expenditure. The advantage of this approach is that by its nature, DNS is a resilient, low-bandwidth application. There is no need for the relatively intensive step of establishing an SSL/TLS session only to be told the resource has moved and to repeat the process again at the chosen vCloud Director site. During the process of DNS resolution, the identity of the intelligent DNS server responsible for the "portal" FQDN is established using DNS queries. Those queries can return the details for multiple intelligent DNS servers at different provider sites, a process inherent to the operation of DNS. This makes the deployment of resilient global access more straightforward than the traffic approach in the previous section. If coordinated intelligence can be incorporated into the solution, connections can be "balanced" between the sites rather than just "shared" among them.

Unlike the traffic-based approach in the previous section, the redirection in this model takes place within the IP address resolution process of DNS. The user enters "https://portal.cloud.example.com" into their browser and, following the multistage resolution process explained earlier, connects to the IP address of the vCloud Director user interface. No HTTP redirect takes place, so the client expects to see an SSL certificate that is valid for the "portal" FQDN. However, if each site can also receive connections to its http://Site-X.cloud.example.com/" URL directly, without raising a security warning in the user's browser, the SSL certificate associated with the site will need to be valid for the site-specific FQDN too.

There are two common methods for making sure that an SSL certificate is valid for both access routes. The first is to use a "wildcard" to match any value in the left-hand position in the FQDN. In this example, each vCloud Director instance has an SSL certificate matching "*****.cloud.example.com", where the asterisk would match both "portal" and the site's "Site-*X*" name. The same certificate could then be deployed in all three sites and would be valid irrespective of whether the user entered the global "portal" URL in their

browser or one of the site-specific ones. While this seems simple, many organizations are uncomfortable with wildcard certificates because there is no control over the portion of the URI in the wildcard position. If misused, or compromised, the certificate might be used to represent the company on a malicious, but apparently “secure” URI using, for example, “customer-admin.cloud.example.com” for which it would still be valid.

The second method is to include more than one valid name within the certificate. Using the Subject Alternative Name (SAN) field, a certificate can contain multiple FQDNs that are considered valid. The disadvantage of this approach is that the individual certificates are required for each combination of names (for example “portal” / “Site-A” or “portal” / “Site-B”), but the advantage from a security perspective is that the certificate can only represent the names encoded within it, and cannot be used to represent any other company resource either accidently or maliciously. This is often preferred over the wildcard approach, but does increase the administrative effort involved. Both methods will work with vCloud Director but company security guidelines and policies must be considered before either method is selected for a production deployment.

4.4 Association Partial-Mesh Access

The approaches described in the previous section made the implicit assumption that a user connecting to the global URL has an account in organizations present in each service location, and that the site selection process is free to choose from any site. Depending upon the Cloud Service Provider offering, this might not be the case. Customers might choose to take service in only some locations and not others. Providers might be constrained by capacity at some locations and be unable to onboard new customers to those sites. Whatever the reason, if a user logging in through the global URL is presented with the login page of a site in which their organization is not present, their login will fail. To avoid this situation, if the provider does not provide ubiquitous access, the simple global access model requires modification.

4.4.1 Per-Customer Partial Mesh-Access

If a provider offers customers the ability to choose any combination of service locations, the provider must provide the customer a bespoke “global” access capability, or have the customer manually select one of the sites in which they have service for their initial login. If a provider chooses the manual selection option, after the user is logged in to their chosen site, vCloud Director multisite associations allow them to switch between the other provider sites in which they have associated organizations. This option does not represent a poor user experience, and remains the most straightforward to implement, requiring only service design effort to make sure each customer knows which sites their users can log into.

An alternative choice is to embed intelligence into the global access model to identify the appropriate sites in which the user’s organization is present. Unfortunately, because the user’s identity is not known until they attempt to log in, this method is limited to organization-level validation. This approach only works with the traffic-based method described earlier, because the full URL is available to the site selection server. In this model, the global site selection logic is enhanced to examine the full URL, including the tenant / organization name elements when the connection is received. The following figure shows the sequence of events in this process.

Figure 25. Per-Tenant Global Site Selection

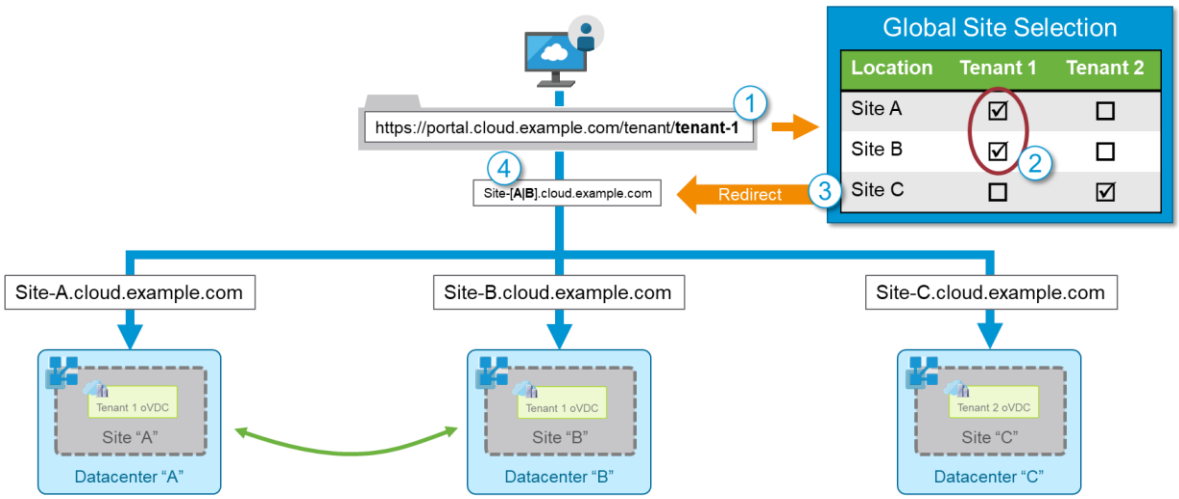

1. The full URL is passed to the global site selection system.
2. The tenant is identified from the organization name at the end of the URL, and this is matched in a site lookup table.
3. The sites which are valid for that tenant/organization are returned.
4. The site selection logic chooses a vCloud Director instance from within the returned set, using the appropriate business logic.

While it is not possible to employ this approach in the exact manner of the DNS-based site selection model, you can combine the two approaches. As the DNS query only contains the FQDN, the organization name is not available to the intelligent DNS server, and so it cannot identify the tenant or sites at which there are organizations present. However, depending on the technology deployed, it might be possible to utilize the local site load balancer, which does receive the full URL, to identify the tenant in question and issue a redirection to a valid site if it determines that the user does not have an organization configured at the local site. This represents a distributed version of the traffic load sharing model and imposes the need to provide an updated list of Tenants (organizations) and associated sites whenever there are changes.

4.4.2 Regional Partial-Mesh Access Model

Another approach to providing distributed access where users do not have organizations in all vCloud Director sites, is to have the customer provisioned to all sites in a particular geographic region or other defined group of service locations. This breaks the single, global access model into multiple, smaller groups. If a customer organization only exists in one site, they can log in directly. If their organization is provisioned in all sites within a "region" or "zone", they can log in at that level and be directed to one of the sites within that associated group. If the customer has service in all sites, they can still log in to the global level and be directed to any of the provider's sites. The following figure shows a hierarchical DNS naming model with access at the site, region, or global levels.

Figure 26. Regional Access with Hierarchical Naming Model

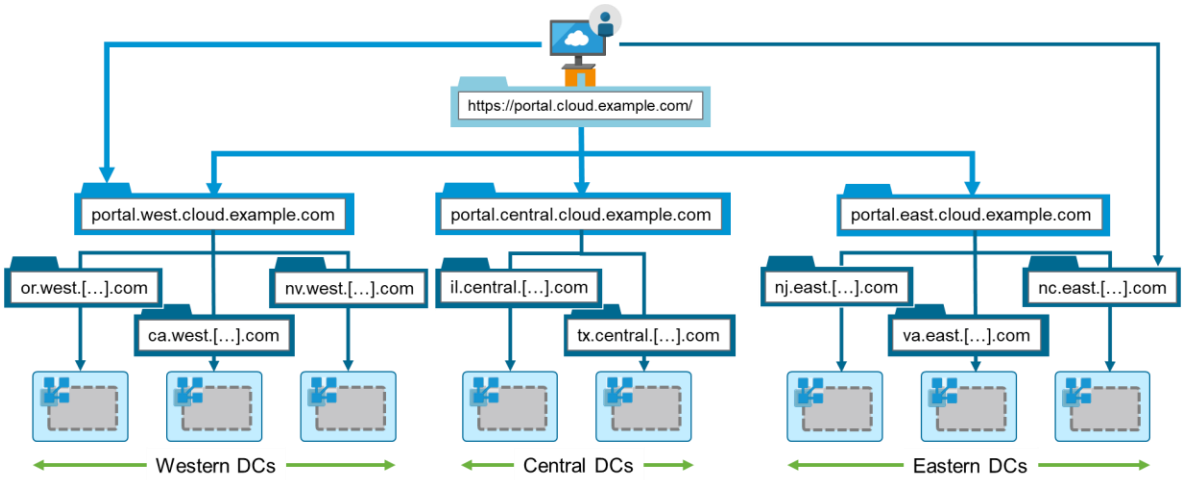

4.5 User Account Requirements for Multisite Access

To transfer user sessions between member organizations, the source member site must authenticate the new user session on the destination member site before redirecting the user's browser session. At the point at which it does so, it only has access to the user identity and organization authentication method from the current session. Currently, to transfer the session to another member site, both the user identity and organization authentication method must match. Similarly, when the user switches to another member site, their roles and rights are determined from those effective at the site to which they have switched. The user identity management types, member configuration requirements, and roles and rights configuration required for effective multisite operation are detailed in the following tables.

vCloud Director can validate user identities from three different sources.

Table 2. vCloud Director User identity Management Types

User Type	Identity Management
Local users	User name, password and permissible roles are all managed entirely within vCloud Director.
Local users with LDAP	User names are managed through LDAP (either Provider or Tenant managed) but are imported into vCloud Director. Their passwords are managed in LDAP and their roles are managed by vCloud Director.
External Identity Provider (for example, SSO/SAML/OAuth)	User names and passwords are managed within the external identity provider (IDP) but roles can either be managed by vCloud Director or by group membership in the IDP ("Defer_to_IDP" in vCloud Director user configuration). The IDP is external to vCloud Director but can be either Provider or Tenant managed.

To switch between sites, the associated organizations must acquire their user identities in the same way. Because the initial vCloud Director site only has access to the user ID and password it validated during the user's login, the follow requirements must be met.

Table 3. vCloud Director Multisite Tenant Configuration Requirements

User Type	Tenant Configuration Requirements
Local users	Associated organizations must all be set to use local users. The same user names and passwords must exist at all sites. The use of Local User account is not recommended because there is currently no mechanism for validating password matching between member organizations and, when password changes are necessary, they must be carried out manually at all sites.
Local users with LDAP	Member organizations must be set to use the same LDAP servers (or federated/synchronized equivalents). Member organizations must have imported the same set of user names from their LDAP sources.
External Identity Provider	Member organizations must use the same external IDP servers (or federated/synchronized equivalents).

Roles and rights are not propagated between associated organizations and are, therefore, determined by the user account at the currently selected member site. The method in which this is controlled varies between user types and is detailed in the following table.

Table 4. vCloud Director Multisite User Roles and Rights Control

User Type	Roles and Rights Configuration	
Local users	Roles and rights constrained by the user account on the current site.	
Local users with LDAP	Roles and rights constrained by the user account on the current site.	
External Identity Provider	Roles and rights set explicitly at each site.	These settings take precedence over other roles and rights settings.
	Defer to IDP with identical group mapping at each site.	Users will have the same rights at all sites.
	Defer to IDP with different group mapping at each site.	Users will have rights dependent upon the specific group mappings at each site.

Note For predictable operation it is recommended that wherever possible user roles and rights are the same in all member organizations. It is however recognized that if different member organizations contain workloads with different security and access considerations, this might not be possible. Where this is the case, and roles and rights are different, users might be able to switch to a site at which they have either no access, or access with limited (or no) effective rights. While this is expected behavior, it might confuse users.

Multisite vCloud Director Design Decisions

Design decisions to consider before upgrading to v9.0, include whether or not to keep or deploy stretched vCloud Director instances and any requirement to deploy complex global access mechanisms. This section examines these two areas and offers some concluding notes.

5.1 The Need for Stretched vCloud Director Instances

Prior to the release of vCloud Director v9.0, Cloud Service Providers built infrastructure resilience and increased availability by distributing elements of the resource platform, managed by vCloud Director, across multiple sites. While providing availability benefits to customers, doing so brought operational challenges. So, does the introduction of multisite capabilities within the latest release make these older topology models unnecessary? The answer, as with many elements of infrastructure design, is "it depends".

While the benefits that can be realized remain, it is possible to simplify the management of a vCloud Director deployment using the multisite enhancements in v9.0. Looking at the two stretched vCloud Director models from Section 2.3, the "dual vCenter Server" model offers the ability to manage two separate resource environments from a single UI, whereas the "single vCenter Model" model offers the ability to manage two separate resource environments as if they were a single environment. Both models also brought with them additional complexity in recovering the management platform in the event of a failure at the site at which vCloud Director was operating.

The dual vCenter Server model's key benefit was the single UI across two sites, a feature that is replicated more elegantly in the multisite capability of v9.0. The dual vCenter Server model requires infrastructure for the "standby" site to recover the management platform in the event of a failure at the "active" site. If that is replaced by an active vCloud Director at both sites with association enabling access to resources at both sites, no additional infrastructure is required, and the complexity of recovering vCloud Director and the rest of the management infrastructure to a remote site is replaced with a simpler operation at both sites.

The "stretched vCenter Server" model's key benefit is a more seamless distribution, operation, and migration of resources across two discrete locations. This use case is not addressed in the current version of vCloud Director because the multisite resilience was previously delivered through the vSphere resource layer rather than the UI. If a Cloud Service Provider wants to continue to offer this level of resilience, at least in the current version of vCloud Director, the stretched vCenter Server model still has a place.

What if the Service Provider already has one, or both, of these multisite models in place within their product offerings? Does the upgrade of a production environment force the provider to adopt multisite association immediately? The simple answer is no. While it can be argued that the introduction of multisite association in place of the dual vCenter Server model offers some operational simplification, the introduction of a second vCloud Director at the "standby" site, and migration of that site's vCenter Server and associated resources, might be outside of the immediate resource budget. The simplest option might be to leave any existing stretched vCloud Director environments as they are initially, and then review the benefits, to both customers and the provider, of moving to a multisite associated model as a separate project. The following figure illustrates how existing stretched vCloud Director instances can be included in a multisite association.

Figure 27. Incorporating Stretched vCloud Director Instances into an Association Mesh

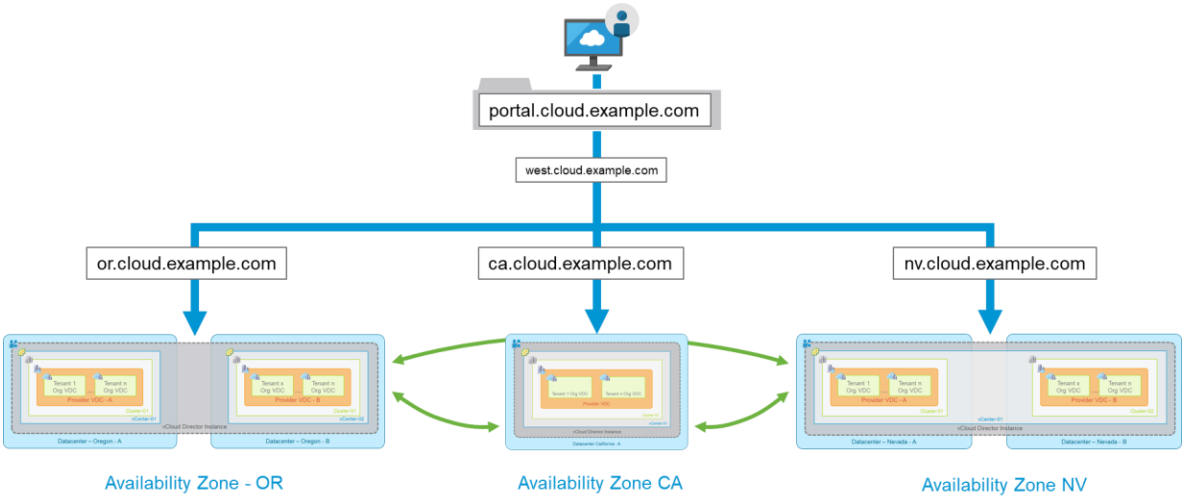

This figure shows the inclusion of existing stretched single and dual vCenter Server vCloud Director multisite instances into a site association mesh which also includes a single-site, non-stretched instance. The sites represent the availability zones within a single region. The Oregon zone is comprised of two data centers under a single, stretched vCloud Director each data center with their own vCenter Server. The Nevada zone has two data centers under a stretched vCloud Director and stretched vCenter Server. For simplicity, the California zone is shown with as a single data center. This could represent a provider deployment following the upgrade to v9.0 in an existing three vCloud Director region. The Oregon zone could then, as a follow-up project, be split into two discrete vCloud Director instances, each joined to the regional mesh as shown in the following figure.

Figure 28. Converting a Stretched vCloud Director Deployment into an Associated Pair

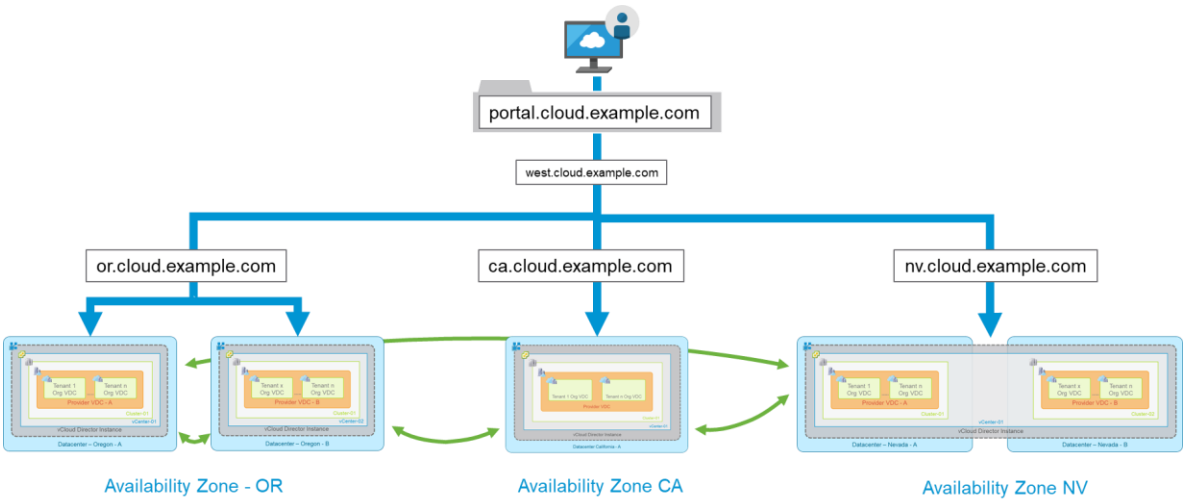

Although not all associations are shown, the provider makes sure that the required site associations are in place to enable tenants to establish organization associations as appropriate.

5.2 The Need for Global Access

The design and deployment of a single, hierarchical access model is not a trivial undertaking. Neither is the task of maintaining this platform as new sites are added or capacity exhaustion causes the available sites within a zone to change over time. Because of the advantages to offering this resilient access model, it is worth considering as part of the deployment of vCloud Director or the upgrade to v9.0. However, this level of complex access capability is not necessary to realize the benefits of multisite access. A Cloud Service Provider might choose to deploy a more local version in which each site has its own direct access URLs. In addition, sites that are typically used in availability pairs or zones can have an access URL that distributes connections across those few sites, as shown in the following figure.

Figure 29. Local Availability Zone Access

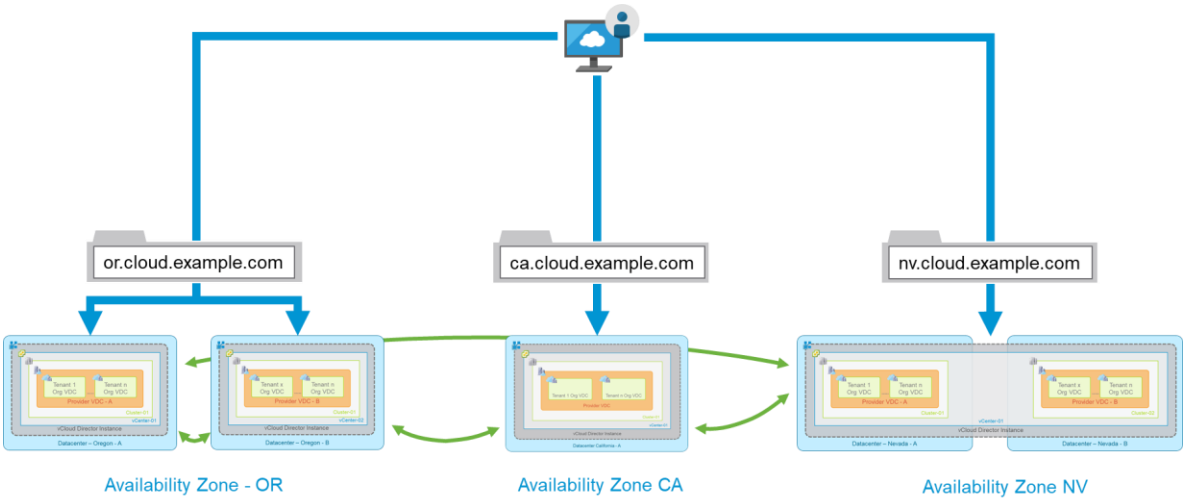

Customers might choose to take service in sites which the Service Provider does not consider part of a single zone or region, perhaps at opposite ends of a particular geography. Unless the provider chooses to offer a bespoke per-customer access model, there might not be the justification for a multilevel hierarchy of DNS names and URLs, and direct, per-site access might be the only sensible option.

5.3 Prerequisites for a vCloud Director v9.0 Upgrade or Deployment

vCloud Director v9.0 brings with it a number of new enhancements in addition to the multisite capabilities discussed in this document, such as the new HTML5 UI and support for the NSX distributed logical router. There are prerequisites required to take advantage of the multisite capabilities which must be addressed before an upgrade or greenfield deployment of v9.0. Consider them as part of the larger picture of customer and provider value that can be realized through upgrading to the latest release.

If a deployment or upgrade is technically possible (check the product release notes, upgrade documents, the VMware product interoperability matrix, and consider any external dependencies), you do not need to have a fully fleshed-out availability zone model with bespoke per-tenant intelligent DNS load balancing access before the other benefits of the latest release can be realized. Multisite access can be enabled in most existing deployments (subject to network connectivity and security considerations), offering simplified user access across their estate. Together with the other new features in this release, the benefits to Cloud Service Providers and their customers make the move to vCloud Director v9.0 a compelling proposition.

References

Additional information pertinent to this document and its topics is provided here.

Document Title	Link or URL
VMware vCloud Architecture Toolkit for Service Providers	https://www.vmware.com/solutions/cloud-computing/vcat-sp.html
vCloud Architecture Toolkit (vCAT) Blog	https://blogs.vmware.com/vcat/
vCloud API Programming Guide for Service Providers	https://code.vmware.com/doc/preview?id=5695#/doc/GUID-86CA32C2-3753-49B2-A471-1CE460109ADB.html