VMware vCloud® Architecture Toolkit™
for Service Providers

Migration Strategies for Hybrid Cloud

Version 2.9

January 2018

Edward (Allen) Shortnacy

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

Contents

Migration Overview 9
1.1 Migration Center of Excellence 11
Migration Conceptual Architecture 12
2.1 Discovery and Assessment 12
2.2 Deeper Introspection 15
Migration Logical Architecture 18
3.1 Migration Use Cases 18
Migration Physical Architecture 25
4.1 vRealize Infrastructure Navigator 25
4.2 Capacity Planner 26
4.3 V2C Self-Service 27
Migration Reference Implementation 38
5.1 V2C Self-Service 39
Migration Operational Considerations 42
Conclusion 43

List of Tables

Table 1. Application Lifecycle Stages 13
Table 2. Application Characterizations 14
Table 3. vRealize Infrastructure Navigator Minimum Requirements 25
Table 4. vRealize Infrastructure Navigator Port Requirements 25
Table 5. vCloud Connector Size Requirements 27
Table 6. vCloud Connector Network Requirements 27

List of Figures

Figure 1. Venn Diagram of VMware vCloud Services 9
Figure 2. Migration Center of Excellence 11
Figure 3. Migration Functional Categories 12
Figure 4. Migration Workload Considerations 13
Figure 5. Dissecting Application Containers 15
Figure 6. Migration Logical Architecture Diagram 18
Figure 7. vRealize Infrastructure Navigator Architecture 19
Figure 8. vRealize Infrastructure Navigator Natively Supported Applications 20
Figure 9. vRealize Infrastructure Navigator Application 20
Figure 10. vCenter Converter Sources and Targets 21
Figure 11. vCloud Connector Logical Architecture 22
Figure 12. vCloud Connector in vSphere C# Client 22
Figure 13. vCloud Connector Content Sync 23
Figure 14. Offline Data Transfer with vCloud Connector 24
Figure 15. vCloud Connector Offline Data Transfer 24
Figure 16. vCloud Connector Ports and Data Flow 28
Figure 17. NFS Mount Path Relative to vCloud Connector Node 31
Figure 18. VMware Cloud Provider vCloud Director Tenant Org URL and Credentials 31
Figure 19. VMware Cloud Provider vCloud Director VDC and Catalog 32
Figure 20. VMware Cloud Provider vCloud Director Network Pairing and Power State 32
Figure 21. Keep Catalog Warning 33
Figure 22. Confirmation of Details 33
Figure 23. Export Progress 34
Figure 24. Export Completion 34
Figure 25. Export Verification on vCloud Connector Node NFS Mount 35
Figure 26. vCloud Connector Service Provider Import 35
Figure 27. vCenter Server Credentials Supporting vCloud Director Org VDC 36
Figure 28. vCenter Server Target Datastore 36
Figure 29. Verification of Details 37
Figure 30. vCloud Connector Tasks View 37
Figure 31. vCloud Director View of the Reference Implementation 38
Figure 32. Active Directory DNS View of the Reference Implementation 38
Figure 33. vSphere View of the Reference Implementation 39
Figure 34. vCloud Connector Server on vSphere 39
Figure 35. vCloud Connector Node on vSphere 40

Figure 36. vCloud Director vCloud Connector Node 40
Figure 37. vCloud Connector Node on vCloud Director 40
Figure 38. Linking vCloud Connector Nodes to vCloud Connector Server 41
Figure 39. vCloud Connector View in vSphere Client..... 41

Glossary and Abbreviations

Center of Excellence = COE

VMware vCloud® Architecture Toolkit™ for Service Providers = vCAT-SP

Software-Defined Data Center = SDDC

Development for Operations = DevOps

Virtual to Cloud = V2C

Physical to Virtual = P2V

VDC = Virtual Data Center

Migration Overview

Migration, in literal sense, is a relocation of a population to a place usually considered more optimal for that population's survival or quality of life. Migration may happen periodically or over time and may occur in either direction, to or from a place previously considered home. While this applies mostly to populations of living things, the phenomenon of migrating to the cloud is similar in many ways. For instance, a move from a mostly winter climate to the tropics, where you quickly learn that suntan lotion must always be in plentiful supply or you will pay a handsome sum for it at the beach. Or, when relocating to a different country for a new job, it is often beneficial to learn the native language to gain proficiencies in your post-migration lifestyle. This analogy for migration of applications to the cloud is used throughout this document to illustrate a running theme. VMware wants customers showing up prepared and able to take advantage of this new real estate, called hybrid cloud, for their information technology and business operations.

Ideally, once the initial move is completed in a migration, preparations can be made for future moves, the eventual visits, or even a permanent return to home. In the case of VMware Cloud Providers™, this capability is delivered through the hybrid cloud network coupled with migration services that seek to help customers understand which portions of their workloads will not only survive, but thrive under the new conditions. This section of the *VMware vCloud Architecture Toolkit™ for Service Providers* (vCAT-SP) explores the idea of and preparation for migration, including hybrid cloud readiness. The following figure is a Venn diagram of some of the tools required to establish each of these capabilities. While not all of the tools shown in the figure are required, this document covers the VMware solutions available to VMware Cloud Providers and their collective purpose.

Figure 1. Venn Diagram of VMware vCloud Services

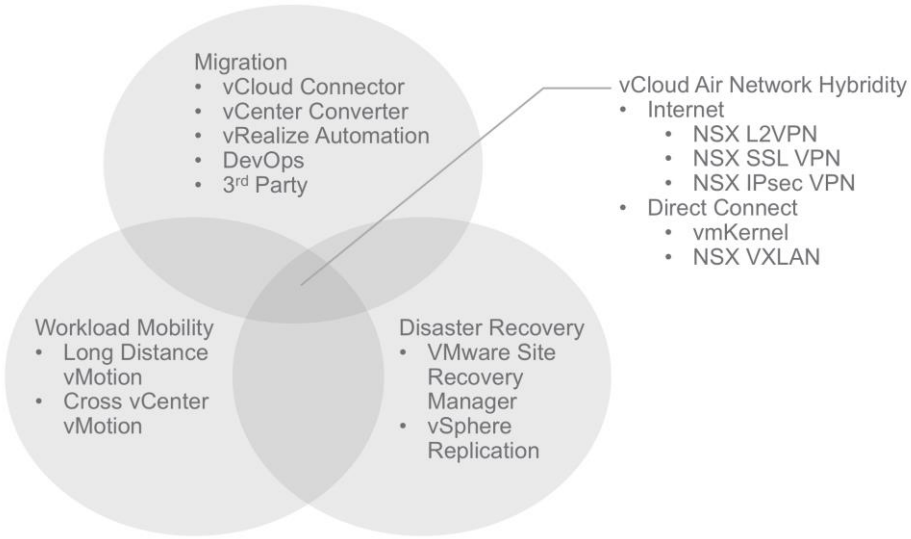

The components are used to deliver the core functionality of the hybrid cloud, which is mainly to connect the private cloud to VMware Cloud Provider Program partner data centers. This connectivity occurs at the physical layer, such as MPLS provisioned from a telecom carrier into the data center or through self-service software-defined networks, such as a Layer 2 IPsec VPN, established by the customer. These networks must have the throughput to support the number of active, concurrent migrations, and decisions must be made as to whether they will be temporary or long-lived.

Throughout the vCloud Architecture Toolkit, the phrase "workload mobility" is used. To clarify, while migration does have a component for moving workloads between environments, tackling the challenge through only workload mobility leaves a portion of the real opportunity unaddressed. When evaluating workload mobility against technologies like VMware vSphere® Replication, VMware Site Recovery Manager™, or even long distance VMware vSphere vMotion®, we assume that workloads are firmly

ensconced in the management of not only vSphere but other components, such as VMware NSX® or VMware vRealize® Operations™. The opportunity then involves capturing as much information as possible about workloads, regardless of their platform and dependencies, to gain adequate control of both their migration and ongoing management post-migration in the hybrid cloud.

Note The previous figure calls out third-party technologies that are not explicitly discussed in this section of vCAT-SP. There are VMware Technology Partners with solutions that cover use cases that are not covered in this document. Look for future vCAT blogs (https://blogs.vmware.com/vcat/) focused on migration partners.

Migration might be a one-time only event, where often many choices about which technologies, configurations and processes can be made within the scope of the migration window itself. It might also be a one-way journey where you are leaving physical systems or virtualization platforms behind, which might demand an even more utilitarian approach from the networking and tools perspective. When planning for migration, service provider partners and customers must do an analysis of future states based on application needs and opportunities, as well as the VMware Cloud Provider Program roadmap. This upfront analysis phase helps to make sure that the appropriate resources and tools are chosen. This analysis also instructs as to the sustainability and desired transformational capabilities of those resources and tools for enterprise customers to leverage future VMware Cloud Provider Program offerings. Examples of these include hybrid cloud, disaster recovery to cloud, database replication to cloud, and so on. Many of these examples are outlined in the vCAT-SP.

No choices can be made future-proof. VMware intends to establish the appropriate set of evaluation criteria for understanding how to transform to a future state so that any barriers for customers consuming VMware Cloud Provider Program hybrid cloud become greatly diminished. All parties involved in the migration to VMware Cloud Providers must have the correct set of working information. This data is used to plan and to remove some of the unknowns that tend to occur when information is being gathered by parties to make mutual decisions concerning application migration. The critical difference between a workload and application is that a workload is simply a container for an operating system and other opaque contents perhaps related to an application. The application residing in that workload, in a migration context, requires understanding and planning to accommodate the application layer dependencies that exist within workloads.

After this method of migrating and consuming VMware cloud technologies is established, whether on-premises or through the VMware Cloud Provider Program, you can leverage all of the benefits of a public cloud provider market while addressing the long-term potential risks of migrating to the public cloud without possessing true hybrid cloud capabilities. Much like renting your house as you migrate to a new location to preserve the potential for a return home, a good migration plan allows for an evolution of application placement that matches your overall enterprise goals and objectives. This document describes the thinking, the approach, and examples of different tools used to accommodate migration to the cloud while identifying and addressing the critical path criteria for success. It also empowers VMware service provider partners and customers to leverage the correct capabilities at the correct time using the appropriate price/cost structures.

1.1 Migration Center of Excellence

In many ways, there is no single requirement for a solution provider or customer to choose a particular combination of tools, whether VMware or third-party. VMware recommends using a Center of Excellence (COE) within the VMware Cloud Provider Program partner circle, a system integrator, or anyone who wants to own this concept of migration to the cloud for their customers, whether they are internal or external. The migration COE must be replete with potential migration use cases and adopted core service provider architectures presented within the vCAT-SP. The top-level migration use cases are P2V (physical-to-virtual), V2V (virtual-to-virtual), V2C (virtualization-to-cloud), and C2C (cloud-to-cloud).

Due to the varied combinations of tools available, this document provides a list of attributes for evaluating possibilities with the understanding that VMware customers, VMware service providers and their partners, as well as VMware itself will have differing recommendations depending on criteria. Within the VMware network of larger service provider partners, many permutations of a migration COE might exist. These permutations are used to best fit the potential combinations of tools, customers, partners, application types and so on, better enabling them to effectively "digest" the massive potential that is the coming years' migration to the cloud. In the following figure, the migration COE is made up of the people, processes and technologies involved in the migration of customer workloads to the cloud.

Figure 2. Migration Center of Excellence

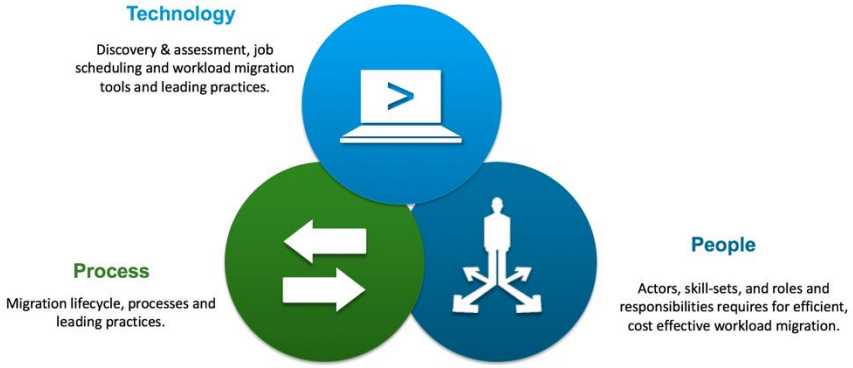

Migration Conceptual Architecture

The stages of the migration journey begin with an alignment of the four types of tools shown in the following figure. While each of these four types of tools might be required to complete a migration, the degree of support for each might vary for each set of requirements. Deployment of the tools might occur on a per-tenant basis, might be self-service, or might be operated by an ecosystem of participants, including third parties such as system integrators. There are a number of criteria that must be streamlined across architecture to operations, as choices made in either area may impact choices in the other. Adding traditional service provider capabilities, such as infrastructure and network provisioning, capacity monitoring of these dependencies and the overall timing for all of this to take place, it becomes evident that detailed planning must be undertaken. The level of detail has several variables that this document attempts to illustrate, and in many cases, it is a matter of where to perform which tasks for which there is no “right” answer.

Figure 3. Migration Functional Categories

2.1 Discovery and Assessment

To plan a migration “journey,” you must first collect information about applications running in the enterprise to support analyses of their suitability for transformation. Ideally, this enterprise architecture knowledge base exists in most enterprises today and is a good place to start to understand the lifecycle of an application and its dependencies. Some of the criteria and dependencies are shown in the following figure. There are many considerations to be made that will have an impact on decisions about where these workloads might live and how they might best be served by new paradigms, such as hybrid cloud afforded by VMware Cloud Providers. You can think of the items in the following figure as requirements you might give a real estate agent for a good picture of what your applications’ new home will look like. Of course, in this case that new home might be a VMware Cloud Provider. We want the best way to match tenants and their applications with their new home and Discovery and Assessment is the first step in reconnaissance to execute a successful migration.

Figure 4. Migration Workload Considerations

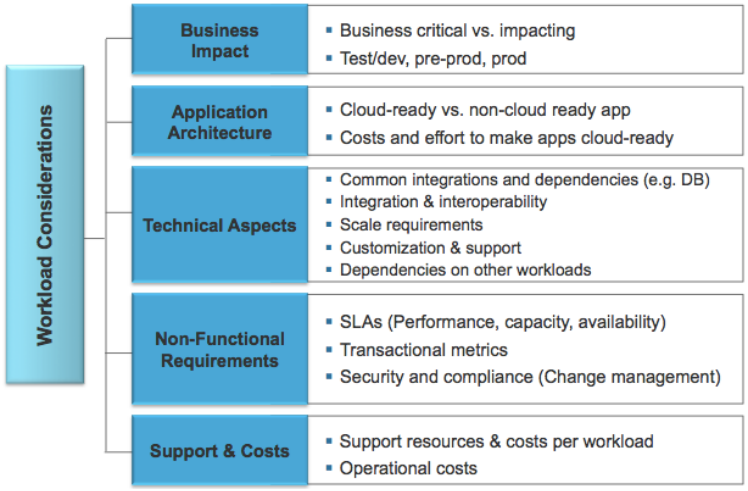

The following table describes the first of several attributes in what becomes a multi-dimensional analysis, although there are not an overwhelming number of values in any dimension. An application can be in one of many states in terms of its lifecycle. In many cases, there might be a lack of information or a lack of understanding as to how available information is parsed so that enterprises can decide whether applications are migrated beyond virtualization and into some type of cloud environment. The values in the following table are desired future state values, which will require some sort of transformation. (Except in the case of numbers 1 and 2, which is a conscious decision to either end the application lifecycle, maintain status quo, or at least forego any migration to virtualization or cloud.)

Table 1. Application Lifecycle Stages

Stage	Description
1.	End Of Life
2.	Remain on Physical (or non x86)
3.	Virtualized
4.	Private Cloud
5.	Public Cloud

Dependencies are important to glean in the Discovery and Assessment phase because they aid in characterizing applications in a different dimension from that of the Lifecycle Stages shown in the previous table. The new dimension in the following table addresses the overall complexity of the application with regard to its dependencies. It is critical to map and understand how applications leverage different tiers such as a traditional three-tier application with database, business logic, and Web facets.

While these discovered migration candidates might be defined local to a single instance of an operating system (a virtual machine in the case of migration), they might also be located in more than one virtual machine grouped together in what is known as a vApp, or a form of container to create startup/shutdown sequence and other elements such as a Network Address Translation (NAT) service from a VMware vCloud Director® managed VMware NSX Edge™.

Table 2. Application Characterizations

Number	Application Characterization
1.	Single Virtual Machine (VM)
2.	Multiple VMs (vApp)
3.	VM or vApp with Dependencies
4.	Distributed Application
5.	Packaged Application (SAP, etc.)

While numbers 1 and 2 in the table have inter-dependencies that are self-contained, as you transition from number 2 and beyond, you start to rely on services such as PKI, DNS, Internet routing, Active Directory, or LDAP, to name a few, that either have to be made available in the new environment or have a network configuration that allows for the components to communicate across sites. Just as whether or not you bring your vehicle on your aforementioned human migration is determined by its suitability for conditions, cost to keep it available, and so on, similarly, you must decide whether to make, for example, Active Directory available in the cloud through a new instance or with some sort of federation, providing the authentication services required. The potential changes to be discovered are in the operating model of account provisioning, network availability, throughput for replication, and so on.

After these dependencies are well understood and these services are available in the new environment, either by new deployments or having the existing services delivered through a hybrid cloud network, migration of the critical services can begin. The capability to configure your virtualization and hybrid cloud topology with VMware NSX and vCloud Director is described in other sections within the vCAT-SP. Providing for these dependencies is a requirement for the migration of any application with a number 3 and above from the preceding table. It is also where the most value lies, in terms of leveraging a VMware Cloud Provider, because of the ability for hosting enterprise applications in a hybrid cloud architecture when defined by architectures in the vCAT-SP.

In the previous table, the attributes reflect a maturity model, where number 4 is not only a VM or vApp with dependencies, but literally a distributed application architecture. This is the Zen-like apex of the hybrid cloud phenomenon, where each component in an application and its dependencies are located wherever they are best suited to serve their purpose, and dynamically providing the network connectivity and operational control to be able to achieve that goal. As an example, consider an e-commerce application that has Web and business logic tiers running elastically in a VMware Cloud Provider implementation. This application leverages an Active Directory server for authentication and a customer master database that both remain on premises and connected to the VMware Cloud Provider Program data center. It also posts transactions for orders to an SAP system hosted in Virtustream or other VMware Cloud Provider. If there is a well-defined set of network requirements such as those provided in the SAP HANA Network Requirements white paper, you can gain tremendous agility and reusability by leveraging the VMware software-defined components to support the application wherever it runs across the VMware Cloud Provider Program hybrid cloud environment.

In the previous table, number 5, the Packaged Application (such as SAP used in the example), can be a distributed application or not, but it also has additional considerations. To get an idea of why SAP or packaged applications might require their own migration category, see the VMware Adapter for SAP LVM document that describes how the adapter allows for automation of the packaged application along with SDDC constructs. The automation of the collective APIs from VMware and the packaged application vendor creates tremendous out-of-the-box opportunities to host large enterprise applications on the only SDDC platform able to provide the kind of performance required.

Categorizing migration candidates with these simple tags yields a subset of applications that can be suitable for one-time, self-service migration using one of the tools prescribed by the COE. As long as the target provider infrastructure and the dependent network connectivity or other services have been made available and the stakeholders involved agree, the opportunity to migrate workloads and perhaps even to later "stitch" in other dependencies to made available over the network can yield a successful first "wave". Lack of planning and understanding the totality of the migration, its waves and dependencies, however, can place the overall successful migration in jeopardy. Whether an application exists in a single VM or a distributed set of them, you can discover more about them for enhanced manageability in the hybrid cloud model. Methods to gather such information are covered in the following section.

2.2 Deeper Introspection

So far, this document has discussed virtual machines, their inter-relationships, and relationships to external dependencies. There is, however, another wealth of information that lies within the virtual machines themselves. That information is depicted in the following figure and can be leveraged to take control of facets internal to each virtual machine, such as operating system versions, application runtimes, patch level, and configuration. All present a value proposition brought about by the SDDC, hybrid cloud, and migration tools available for discovery and assessment.

Figure 5. Dissecting Application Containers

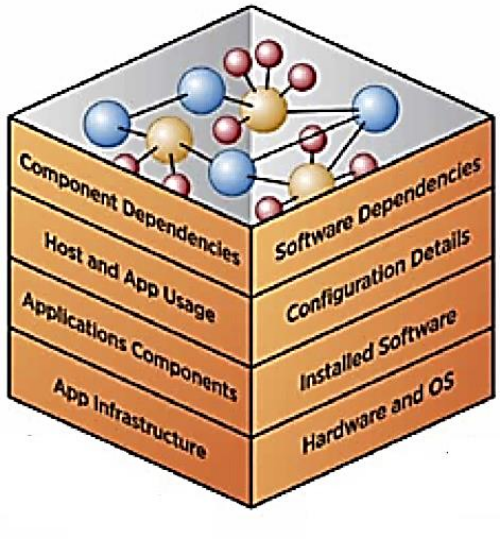

Looking at the application not only as a VM and network topology, but also a container of items that hosts application components, you can gather inventories on, for example, Apache Web servers, Java virtual machines, Windows or Linux OS, .Net or other runtimes and their respective configurations/versions. Additionally, information about the physical hardware (as applicable) along with host and application usage can now be captured. In most instances, you need to know about these items to gain control over facets of the application and its container to enable migration to a new environment. With this knowledge, you can plan the migration of many applications that share common components to standardize with regard to accepted versions and configurations. You can also accommodate a service and support model of shared responsibilities, where third parties might be required to execute tasks relative to application operations and maintenance. If you were considering a property management company to manage your remote properties, you would want very clear terms for what must be looked after, and they would want detailed information on how to do so and report back to help establish satisfaction and trust within the relationship.

This phase of Discovery and Assessment results in a repository of knowledge that drives the overall strategy for cloud migration. Developing and harvesting this body of knowledge is critical in creating a

foundation that supports many other VMware solutions represented within the vCAT-SP, such as VMware NSX, vRealize Operations, VMware vRealize Automation™, and VMware vRealize Business™. It also forms the foundation for embarking on a more DevOps-centric model which, when based on the VMware SDDC and VMware Cloud Providers running the SDDC, can be completely automated to not only migrate applications to the cloud but re-create them and their dependencies in total. Other content from the vCAT-SP addresses creation from standardized component libraries, including operating systems and application runtimes as well as infrastructure templates, such as the networking and security they require.

2.2.1 Job Scheduling

Job Scheduling might not be as complicated in a technology sense as Discovery and Assessment, but because it involves coordinating people and their functions, it can be more challenging. Job Scheduling during a migration is important because, according to our running metaphor, showing up to your new location without furniture having arrived is no fun. As a function of available tools for migration, Job Scheduling may or may not be present. Whether or not Job Scheduling requires an entirely separate tool, manual processes or a combination thereof, it must be governed by the migration COE. Job Scheduling begins the process of migration because, ideally, agreement has been reached with the application owner regarding the inventory collected in the Discovery and Assessment phase. Other criteria, such as cutover window, recovery point objective (RPO), and target performance expectations, must be mutually agreed on as well.

Many parties must coordinate to prepare the readiness of the migration infrastructure with the readiness of the target infrastructure that will host the application and facilitate its dependencies. Therefore, the migration COE must prescribe governance with entitlements and workflow. Because there are self-service capabilities for customers to migrate applications in an ad-hoc fashion, coordination is required for sustaining performance when a large number of jobs are requested so as not to over-utilize available resources during the migration window.

Typically, after provider provisioning of target and migration resources is optimized, migration is bound by human capacity. Therefore, coordinating labor-intensive tasks is critical for anyone directing or participating. Looking at the broader workflow discussed in this section, there is Application Verification, which can have automated components where available but likely consists of application owners and end-user populations that validate that the application is behaving as expected in the new environment. Because there can be Discovery and Assessment, Workload Migration, and/or Application Verification jobs that must be scheduled, the process is generally far from linear. Discreet steps require handoffs, notifications of completion, or exceptions, as well as some feedback into an aggregate set of management reports or dashboards. With both business and technical requirements and a diversity of environments and parties, Job Scheduling is a reflection of both time and resources and understanding capacities and capabilities of all types. Mastering the inputs and outputs of these tasks along with their measurements makes certain that the migration trains continue to run on time.

2.2.2 Workload Migration

In its simplest form, Workload Migration is about moving bytes from place to place in a way that maps to the prescribed overall process, with manual steps for completion either before or after the relocation itself. This section has described many of the items that must be reproduced in the target environment for the migrated application to be served there. Items such as authentication and authorization methods, public key encryption, and DNS must be made available as needed. Because most of these services are delivered over the network, in this case the software-defined network, to the VMware Cloud Provider's SDDC tenant construct, automation of the deployment of those services and networks is available. This topic is discussed in further detail later in this document and in other vCAT-SP documents and blogs.

In many cases where this type of work is front-loaded or automated, the Workload Migration phase can be the simplest portion of the migration exercise, while for others the actual workload migration might be a very sensitive operation. In any case, coordination with the more discreet movement of individual workloads is of paramount importance. In the case of regulatory compliance or other privacy considerations, for example, you must provide careful management of staging data, at rest and over the wire encryption and privileged credentials to execute certain migration tasks and their log evidence. Think

of it as bubble wrap and “white glove” handling for your valuable items, the heirloom china, or grand piano. While this undertaking can be intricate and required to provide evidence, there are methods, tools, and partners who can help customers achieve these most sensitive of workload migrations.

2.2.3 Application Verification

After migration is complete, the remaining steps are to verify that the application is working properly in its new home. As mentioned previously, this typically consists of regular users of the application or other service, such as an administrative function, undergoing user acceptance testing of some kind to make certain they are behaving as expected. However, application verification might often involve more detailed exercises, such as automated functional or load-based testing, an audit of the entire environment to meet some regulatory requirements, and so on. Because this is a software-defined data center environment, you must also consider the Software Development Life Cycle which introduces Application Verification of “stubs,” using tools such as cURL, that can be run from within the tenant to verify network and other application-level connectivity required by components yet to be migrated.

Application Verification is the penultimate step in the migration of an application. In most cases, VMware or VMware Cloud Provider Program technology and service delivery partners offer solutions to accommodate more advanced use cases for Application Verification where such solutions are not available directly from VMware. These can include an auditor for a specified regulation, SaaS-based testing tools for simulating the heaviest loads, or protection from potential penetration from cyber-attack in the new environment. Whatever the combination or rigors applied, after sign-off by the application owner and other stakeholders, the process is considered to be complete.

Migration Logical Architecture

This section describes some tools that can be used in developing migration capabilities for VMware Cloud Providers, their partners, and customers. While VMware provides the tools discussed in this section, there are a wealth of third-party tools available from VMware Technology Partners, such as ATAdata, ThinkOn, Rackware, and RiverMeadow to name a few, that provide additional features and migration use cases not available with the tools provided by VMware. Stay tuned to the vCAT-SP Blog for more information on these partners and how they can be leveraged to provide additional migration services to our customers.

3.1 Migration Use Cases

The following figure illustrates how migration technology operates in principal.

Figure 6. Migration Logical Architecture Diagram

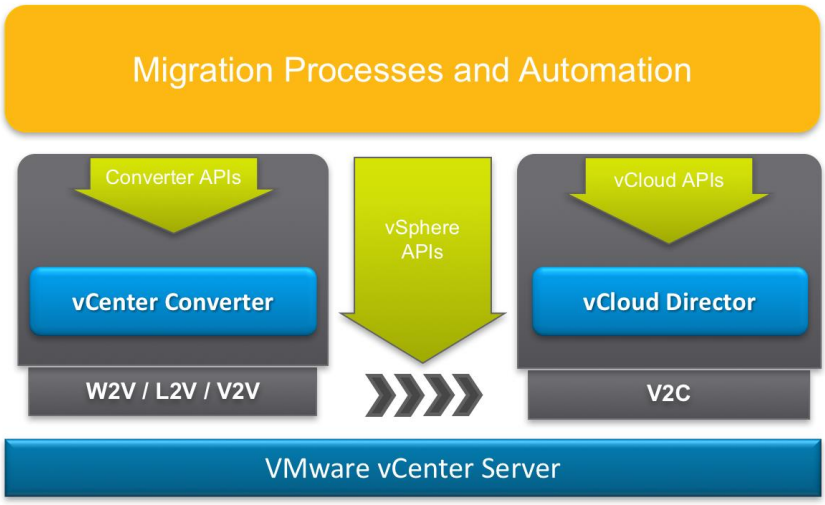

The orange box at the top of the figure represents migration process and consists of tools and human tasks the must be completed. While there might be single tools that support a full migration from physical or any other environment to the cloud, for the VMware tools described here, there is no use case called P2C. However, it is achievable as a two-phase process through combining P2V and V2C together. The P2V use cases supported by VMware vCenter® Converter™ are W2V (Windows-to-virtual), L2V (Linux-to-virtual), V2V (virtual-to-virtual), and V2C (virtual-to-cloud). Technologies in play for each type of infrastructure include vCenter Converter APIs for physical hosts (Linux and Windows) and vSphere APIs for writing new VMs or extracting to OVF format for vCloud APIs.

3.1.1 Self Service

As a service provider, the ideal use case involves deploying some tools for each new tenant and allowing them to, in an on-demand fashion, accomplish their required tasks. However, due to the overall potential complexities involved in migration, this might actually do more harm than good. VMware technologies offer many ways to throttle the consumption of resources. However, trying to employ them to manage a “free for all” situation where each tenant can perform migrations at their discretion might leave a subpar user experience for all involved. Because certain resources might be shared, such as physical network bandwidth or shared transfer storage, careful stewardship of parsing out these resources to consumers is required, and likely some form of Job Scheduling is needed as well.

3.1.1.1 Discovery with VMware vRealize Infrastructure Navigator

As previously mentioned, typically the first step in a migration is Discovery and Assessment. While there are many tools that can be used in this process, both VMware and third-party, this section discusses only VMware vRealize Infrastructure Navigator™. vRealize Infrastructure Navigator is a virtual appliance that you can deploy on VMware vCenter Server®. With the components of vRealize Infrastructure Navigator, you can map services running in your virtual environment, examine the application discovery status, view and analyze the dependency map, and have a centralized view of the entire application environment relative to vSphere. The following figure illustrates various components of vRealize Infrastructure Navigator and their dependencies.

Figure 7. vRealize Infrastructure Navigator Architecture

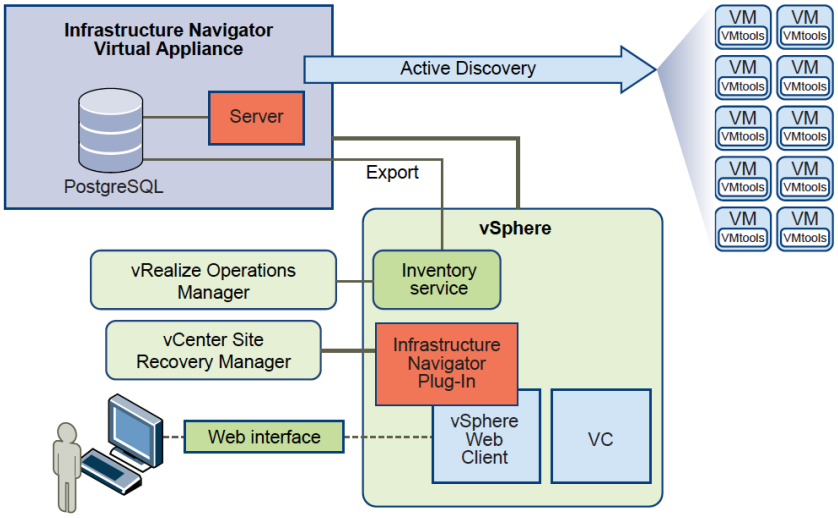

While tools provided by VMware, such as Application Dependency Planner, can aid in Discovery and Assessment of applications running in physical and virtual environments, vRealize Infrastructure Navigator is used to provide similar capabilities within a vSphere virtual environment. Applications can certainly undergo migration from a physical environment to the cloud, but it is often a logical first step to perform a migration to a vSphere based virtual environment. This is because the SDDC platform that on-premises vSphere environments are comprised of is compatible with the SDDC used to implement VMware Cloud Provider environments. Having a hybrid approach to cloud architecture provides many benefits, some of which are described in this document. For this document's purpose, the ability to prepare an application as a hybrid cloud migration candidate involves collecting adequate information about the application and its dependencies. vRealize Infrastructure Navigator is a great tool for doing just that.

Because vRealize Infrastructure Navigator supports the discovery of operating system and application components, including networking and the virtual machine hardware construct (which exposes the underlying components running in a virtual machine), a great deal of pertinent information can be gained by running applications on the SDDC. While not all applications have associated support within vRealize Infrastructure Navigator, many platforms used to build and deliver business applications are supported. The following figure provides an example of supported applications and forms the foundation for an initial analysis of components running in the SDDC and their dependencies to plan a migration for them. There are up-to-date spreadsheets on the product download page at https://my.vmware.com for operating systems and application components that are supported as well as the methods used by vRealize Infrastructure Navigator to perform discovery on virtual machines and their application-specific contents.

Figure 8. vRealize Infrastructure Navigator Natively Supported Applications

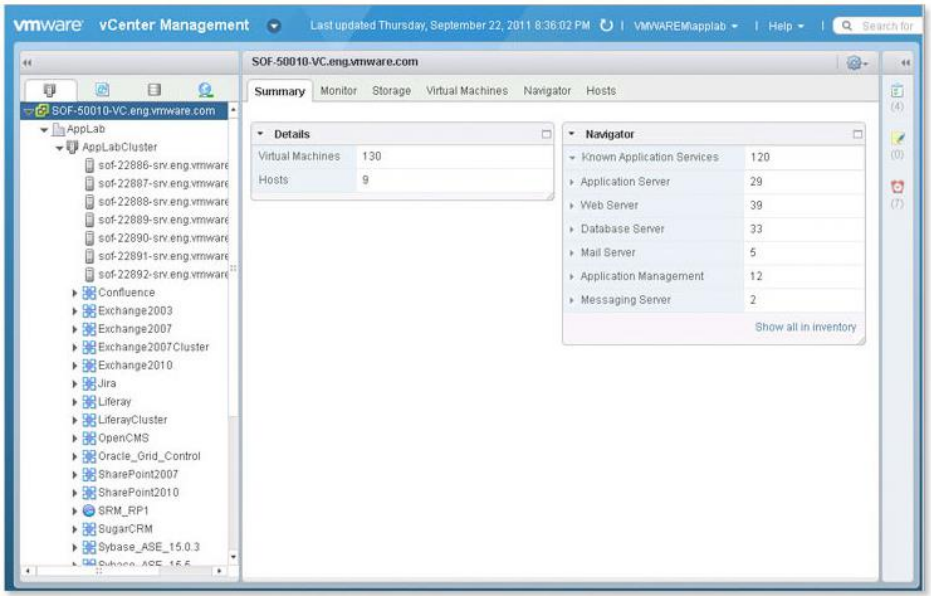

After the Discovery is complete (it perpetually discovers new/changed virtual machines, their components and dependencies), vRealize Infrastructure Navigator can be used for the Assessment phase. It supports this in a variety of ways starting with what are known as Application Definitions. Application Definitions allow administrators to group patterns of the ways in which services (ports and processes) are normally grouped together to essentially create boundaries around groups of servers. When matched to an Application Definition, virtual machines and their contents are mapped together in applications that show communication paths between them, direction, ports, and processes that support the communication. In the following figure, this collection is depicted in the VMware vSphere Web Client providing visual aids for assessing the application landscape.

Figure 9. vRealize Infrastructure Navigator Application

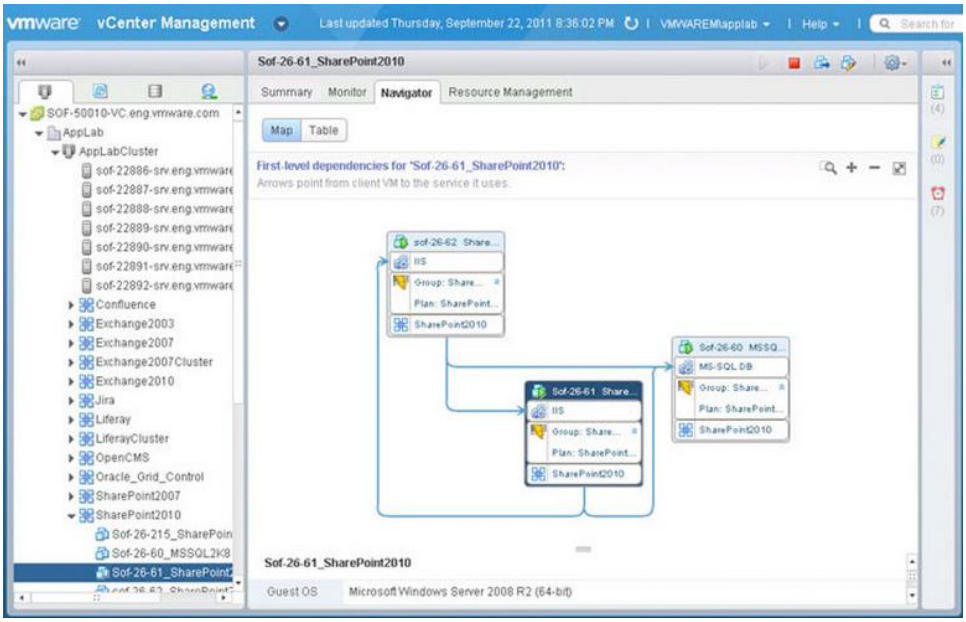

vRealize Infrastructure Navigator is extensible in a few different ways. For one, tabular data for Application Definitions and Inventory can be exported to CSV. To make certain that you are not limited to the Application Definitions items shown in the previous two figures, there are user-defined services where a port and/or process can be provided. vRealize Infrastructure Navigator also collects information about dependency "peers," whether they are virtual or physical, that are not located in the same vCenter Server realm. This information is rendered as a map of IP addresses and DNS names for the dependencies that interact with application VMs under the management of the vCenter Server associated with vRealize Infrastructure Navigator. Because of the extensions provided by vRealize Infrastructure Navigator to the metadata within the Inventory Service, as shown in Figure 7, there are a number of enhanced use cases within the broader VMware vRealize Suite that are not covered in this document but can be found in other sections of the vCAT-SP as well as this VMware TAM Blog.

3.1.1.2 Manual Migration with vCenter Converter and OVF Tool

Migration can sometimes be of the more involved enterprise architecture variety described in the previous sections. For applications that match numbers 1 or 2 in Table 2, it is best to use tools to perform the migration that do not require a great deal of Discovery and Assessment. One method to do this is to utilize vCenter Converter and the OVF Tool. These tools are pervasive within VMware environments and both are free to download and use. They are also fairly well-known within the VMware administrator community. In the following figure, vCenter Converter supports many types of sources and targets and also is a Windows-based GUI tool (with a Windows Service to manage vCenter Converter instances remotely). For information about the latest release of vCenter Converter, see the VMware vCenter Converter Standalone User's Guide.

Figure 10. vCenter Converter Sources and Targets

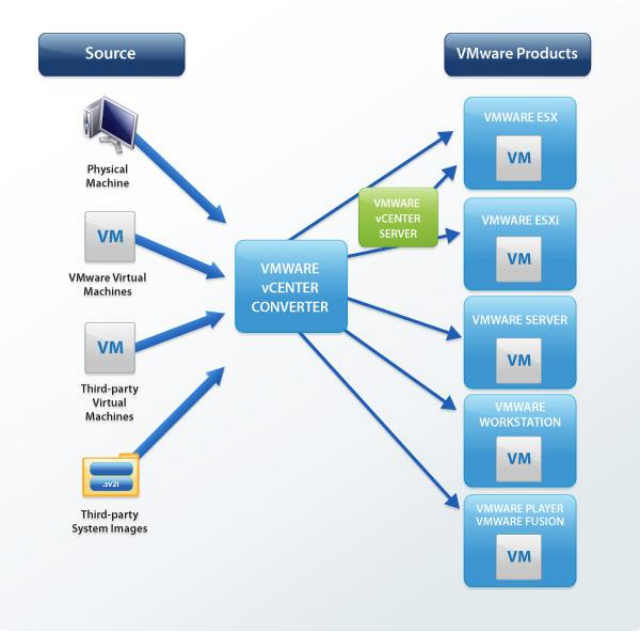

OVF Tool is based on the standard Open Virtualization Format developed by DMTF (committee chaired by VMware along with other vendors) to promote compatibility for accepting virtual machines into different cloud environments. OVF Tool is a command line tool that begins with a simple syntax although more advanced flags are available for specific requirements. For more information, see the VMware OVF Tool User's Guide. In the case of both tools, it is about putting power in the hands of the virtualization administrator to prepare the right package and ship those bits to the target environment, be it from physical to virtualization or to cloud.

3.1.1.3 vCloud Connector 2.7

While OVF Tool is a great command-line utility, VMware customers also need a GUI-based tool that provides the ability to migrate from on-premises virtualization to the cloud. To meet that need, VMware introduced vCloud Connector, now in version 2.7.2. As shown in the following figure, vCloud Connector consists of the vCloud Connector Server for the "Internal" components (on the left side of the diagram), along with additional copies of the vCloud Connector Nodes running at the VMware Cloud Provider locations. This service provider-side vCloud Connector node is typically made available from the vCloud Director Public Catalog to make it easy for customers to deploy to their environments.

Figure 11. vCloud Connector Logical Architecture

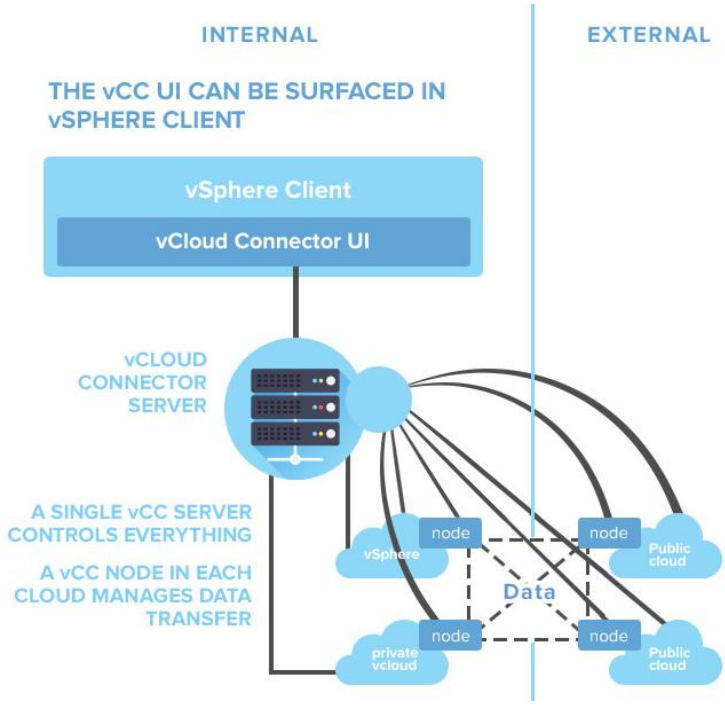

vCloud Connector is a free tool and performs offline migration of applications to the VMware Cloud Provider. The key difference from OVF Tool is that vCloud Connector orchestrates the V2C use case and allows a single place to execute application workload migration that is available from the vSphere C# client as shown in the following figure (visible in the bottom left of the C# client after the vCloud Connector solution is registered). vCloud Connector supports the creation of stretched Layer 2 networks that can be used for common subnets across clouds, easing the requirement to provide networking in the target tenant during migration. Those capabilities are not covered in this section of vCAT-SP.

Figure 12. vCloud Connector in vSphere C# Client

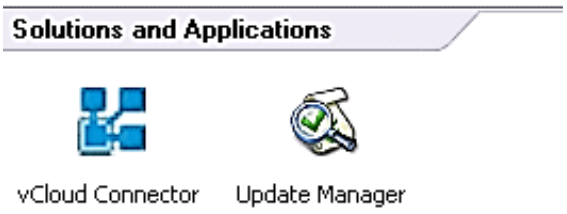

One of the features that previously came only in an Advanced Edition of vCloud Connector (all features are bundled and free as of vCloud Connector version 2.6), is Content Sync. Content Sync enables you to

synchronize content across multiple clouds. It allows you to create a master library of templates and to keep the templates synchronized among users across different clouds. Users subscribe to content that is published to the master library, known as the Content Library. Any changes to items in the Content Library are automatically reflected in the subscribers' folders or catalogs. For example, if a template is added to a folder or catalog that is published to the Content Library, it is added to each subscriber's folder or catalog as well at a preset interval. As shown in the following figure, users can subscribe to a Content Library that syncs to prescribed Catalogs inside each of their mapped VMware Cloud Provider or other vCloud Director based organizations.

Figure 13. vCloud Connector Content Sync

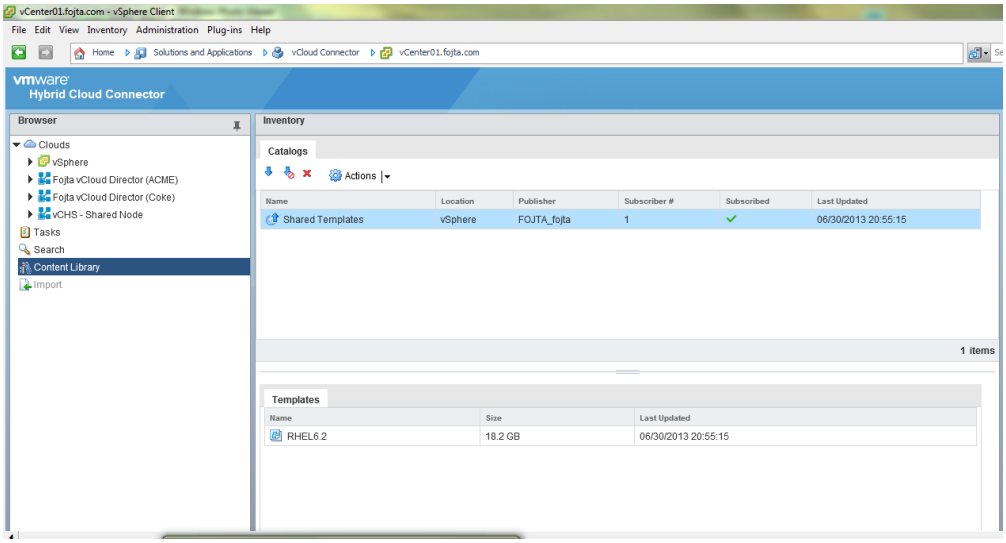

While the self-service model for vCloud Connector is attractive for end users to perform their own migration of application workloads to the VMware Cloud Provider cloud, often times the bandwidth or reliability of the network connection can cause frustration for users trying to upload larger virtual machine disk files. Configuring vCloud Connector for multi-tenancy as well as handling large numbers and sizes of files is described in the next section. Another way in which vCloud Connector can be deployed is for Offline Data Transfer (ODT), which accommodates large file transfers (with latency that is bound by how quickly you can copy data to a drive). You can then ship it to the VMware Cloud Provider and have it loaded into the Organization Catalog. The following figure illustrates this process where the “External Storage Disk” plays the role of the network in a typical vCloud Connector migration.

Figure 14. Offline Data Transfer with vCloud Connector

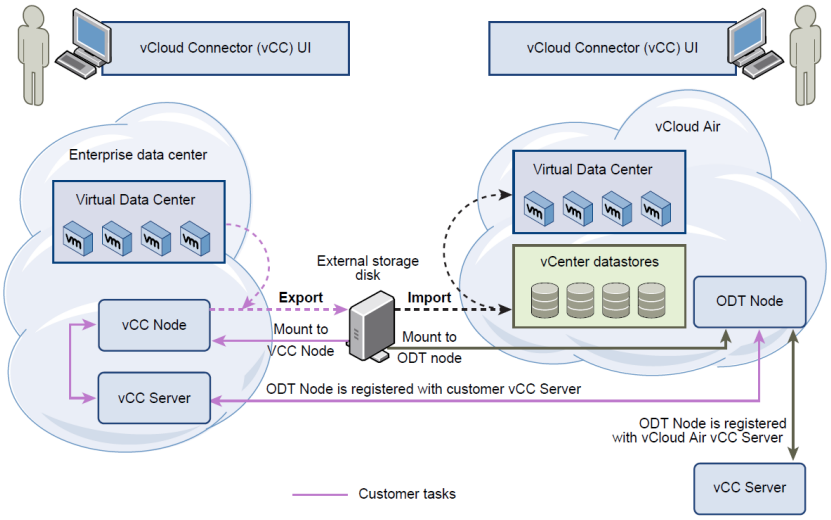

In this case, the customer is still responsible for the migration but the bytes are transferred to storage devices mounted as an NFS file system. The metadata for the vApps are posted to the Organization Catalog in vCloud Director and the VMware Cloud Provider mounts the drive to complete the process, copying the bytes to the data store for the vApps. As shown in the following figure, this operation is available for vApps that are in a "Stopped" state, and a special button is provided after the NFS storage has been configured. More information about configuring vCloud Connector for ODT is provided in the following section.

Figure 15. vCloud Connector Offline Data Transfer

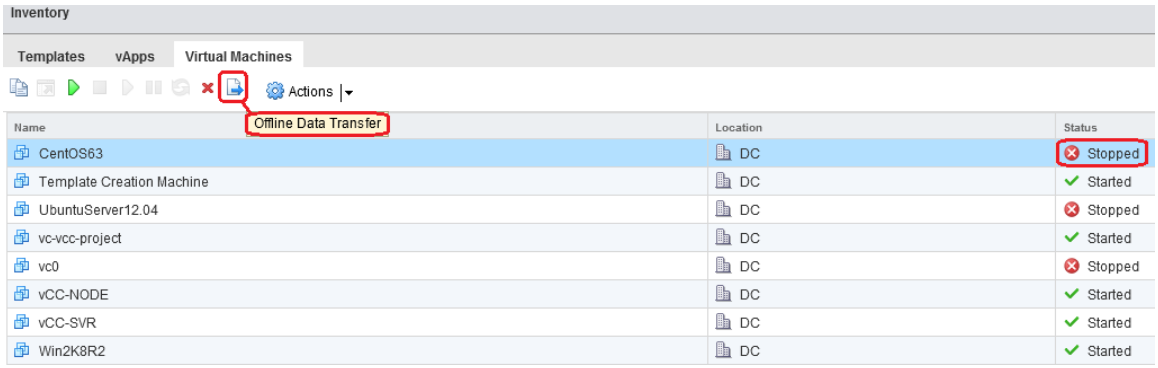

Migration Physical Architecture

In this section, the physical architecture considerations with regard to migration are discussed. While focused on the service provider, in most cases there are requirements for dealing with physical components within the customer data center. As previously discussed, this might consist of migrating physical servers to vSphere with vCenter Converter, provisioning network connectivity to the VMware Cloud Provider Program partner to be used for migration, or other services as well as measurement of required SDDC capacity for the target the VMware Cloud Provider Program tenant configuration.

4.1 vRealize Infrastructure Navigator

vRealize Infrastructure Navigator 5.8.4 is an on-premises tool used in the migration process for Discovery and Assessment. Future vCAT-SP blogs will illustrate how this tool and others can be used to gather adequate information to perform the assessment of migration candidates. The physical architecture considerations for vRealize Infrastructure Navigator are not complex and involve deploying an OVA to your vSphere environment with requirements outlined in the following tables. You can find more information about installation and configuration of vRealize Infrastructure Navigator in the VMware vRealize Infrastructure Navigator Installation and Configuration Guide. Of particular interest beyond the physical requirements listed in the following tables is the configuration of the appliance for secured access. There might be times when accessing the console for data extraction is required and files like `vadm.keystore`, found under `/opt/vadm-engine/conf/`, must be tightly controlled and monitored. Also note the PostgreSQL instance on the vRealize Infrastructure Navigator appliance cannot be run externally from the appliance. However, data extraction can be configured for remote execution.

Table 3. vRealize Infrastructure Navigator Minimum Requirements

Component	Minimum Requirement
CPI	2 vCPU
Memory	4 GB
Disk Size	24 GB
Network	1 Gbps

Table 4. vRealize Infrastructure Navigator Port Requirements

Port Number	Description
From your PC to vRealize Infrastructure Navigator	
5280	Appliance Web console
22	SSH for appliance console access
From vCenter Server to vRealize Infrastructure Navigator	
2868	Plug-in download
6969	vSphere Web Client
From vRealize Infrastructure Navigator to vCenter	

Port Number	Description
80	vSphere Web Service API
443	vSphere Web Service API
10109	vSphere Inventory Service
From vRealize Infrastructure Navigator to Hosts and Virtual Machines	
443	VIX Protocol to perform discovery
902	VIX Protocol to perform discovery

Note Ports 5280, 22, 80, 2868, and 6969 on the vRealize infrastructure Navigator appliance are controlled by the SUSE firewall located in `/etc/init.d/SuSEfirewall2_setup`.

vRealize Infrastructure Navigator is part of the vRealize Suite and as such has some meaningful integration with other solutions like vRealize Operations. The vRealize Operations Management Pack for vRealize Infrastructure Navigator, which facilitates sharing of discovered facets of virtual machines, creates a rich environment for management of operating system and application layer components. When coupled with other integrated solutions, such as VMware vRealize Hyperic®, a fully managed operational model can be created for applications discovered and monitored by this combination. While out of scope for this document, more information can be found in other sections of the vCAT-SP.

4.2 Capacity Planner

As a service provider, one of the most important aspects of planning for the Physical Architecture portion of a migration is capturing information about the required capacity to support new tenants along with their migrated and newly deployed workloads. To gather this information from customer premises running physical or virtual workloads, VMware provides a tool called Capacity Planner. The data captured from Capacity Planner is useful in understanding consolidation ratios when the source workloads are running on physical hosts, and potential over-commitment scenarios that can be used for a subset of applications depending on tenant resource allocation models available in the VMware Cloud Provider offerings.

Capacity Planner has a SaaS portal at https://optimize.vmware.com that presents findings from infrastructure utilization measurements. VMware PSO offers a Virtualization Assessment package built around Capacity Planner that can help provide optimal readiness for tenant migration efforts. VMware Solution Providers can also use Capacity Planner and this service model to deliver a similar service. More information can be found at the Capacity Planner product page.

4.3 V2C Self-Service

This section focuses on the vCloud Connector implementation of the V2C migration use case to accommodate self-service migration of vApps to the cloud. It also covers items related to preparing for the support of Offline Data Transfer with vCloud Connector.

4.3.1 vCloud Connector Deployment

The size requirements for deploying vCloud Connector components are listed in the following table.

Table 5. vCloud Connector Size Requirements

VM	Datastore Space	Total Space Needed
vCloud Connector Server	10 GB + 3 GB + 0.1 GB = 13.1 GB	13.1 GB + 52.1 GB = 65.2 GB
vCloud Connector Node	10GB + 40 GB + 2 GB + 0.1 GB = 52.1 GB	

The network connectivity requirements for deploying vCloud Connector components are listed in the following table.

Table 6. vCloud Connector Network Requirements

Purpose	Source	Destination	Port
Management Traffic for Customer's vCloud Connector Server and Node	Customer's Management Access Source	Customer's vCloud Connector Server and Node	TCP 5480
Communication between Customer's vCloud Connector Server and Node	Customer's vCloud Connector Server and Node	Customer's vCloud Connector Server and Node	TCP 80 and 443
Communication from Customer's vCloud Connector Server and Node to Customer's vCenter Server	Customer's vCloud Connector Server and Node	Customer's vCenter Server	TCP 80 and 443
Communication from Customer's vCloud Connector Node to Customer's ESXi Hosts	Customer's vCloud Connector Node	Customer's ESXi hosts	TCP 443, 902, and 903
Communication from Customer's vCloud Connector Node to the VMware Cloud Provider Program vCloud Connector ODT Node URL	Customer's vCloud Connector Node	VMware Cloud Provider Program vCloud Connector ODT Node URL	TCP 80, 443, and 9443

Purpose	Source	Destination	Port
Communication from Customer's vCloud Connector Node To VMware Cloud Provider Program vCloud Director	Customer's vCloud Connector Node	VMware Cloud Provider Program vCloud Director Public URL	TCP 80 and 443

4.3.2 vCloud Connector Data Flow

The data flow in the following figure indicates the vCloud Connector Server is located on the customer premises. As shown in the bottom right hand corner of Figure 14, there is a second vCloud Connector Server deployed by the VMware Cloud Provider in the case of ODT.

Figure 16. vCloud Connector Ports and Data Flow

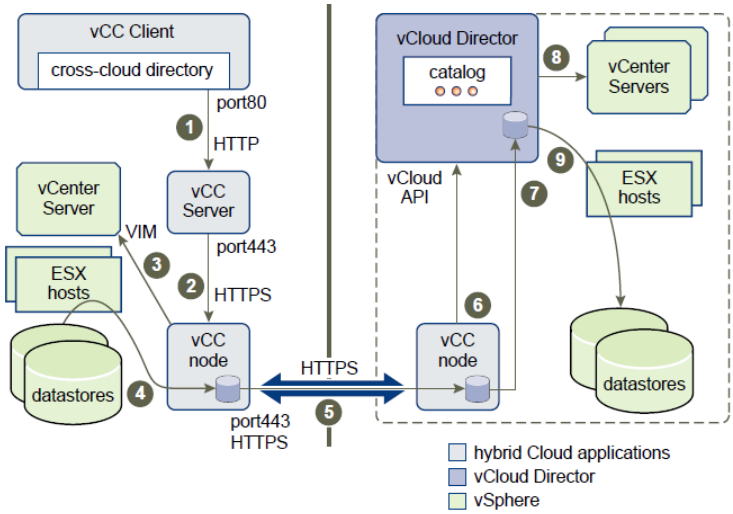

As shown in the previous figure:

1. The customer premises vCloud Connector Server is responsible for presenting the vCloud Connector client interface for the initial migration.
2. The customer premises vCloud Connector Node is instructed to begin an export of the vApp.
3. That export is invoked from vSphere with the VIM API.
4. On the vSphere datastores, parallel operations are started with regard to the activities of streaming data from the datastores through the on-premises vCloud Connector Node. These operations occur over the VMkernel port.
5. More activities are used to pass data over to the VMware Cloud Provider Program based vCloud Connector Node using MIME multipart type transfer over HTTP(S) to allow for check pointing of progress.
6. After the transfer is complete, the vCloud Connector Node in the tenant Org signals the vCloud Director APIs to perform an import.
7. The workload is transferred to vCloud Director transfer server storage.
8. vCloud Director then instructs the vCenter Server to import the vApp.

9. After making the vApp available in the datastore, the vApp appears in the vCloud Director Organization Catalog where it can be imported, further configured, and powered on.

4.3.3 vCloud Connector Production Configuration and Tuning

vCloud Director transfer server storage is out of scope for this document but can be found in other portions of the vCAT-SP. By default, a vCloud Connector node offers a 40-GB staging area, which is available as a LVM Volume Group in the node (shown from the Linux console with a `df -h` shell command). You can increase the staging area by following this procedure.

1. Log in to the vSphere Client.
2. In the hierarchy tree, select the vCloud Connector Node virtual appliance.
3. Right-click and select **Edit Settings**.
4. The **Virtual Machine Properties** window opens to the **Hardware** tab.
5. Select **Hard disk 2** in the **Hardware** column.
6. Modify the size, based on the size of the resources you are going to be transferring, and click **OK**.
7. Open the console for the vCloud Connector Node.
8. Run the following shell command to resize the disk:
   `sudo /opt/vmware/hcagent/scripts/resize_disk.sh`

For vCloud Director based vCloud Connector nodes, you cannot increase the virtual hardware disk size (because the disk resize routine is not a feature of vCloud Director). You have to add an extra virtual disk, and add this disk to the staging area. The procedure is described in the “Configure vCloud Connector Nodes” section of the VMware Installing and Configuring vCloud Connector guide.

Rather than making each tenant deploy a vCloud Connector node, the service provider can provide each tenant with a URL to a vCloud Connector node that is managed for items like disk capacity. These vCloud Connector nodes are called multitenant nodes and can support up to 20 tenants each. It is possible to alleviate both tenant storage and network bandwidth consumption when undergoing migration with vCloud Connector as opposed to private vCloud Connector Node deployed inside of the Org VDC. VMware Cloud Provider administrators can access the console of the multitenant node for storage expansion or access the log files as shown in the “Accessing Log Files from the Console” section of the installation and configuration guide. Other items, such as installation of production certificates and configuring maximum concurrent tasks are also detailed in the “Preparing vCloud Connector for Production Use” chapter of the same document. Multitenant vCloud Connector nodes run in a headless fashion with regard to vCloud Connector functionality (requiring port 8443 open for communication as well as vCloud Director on TCP 443).

Network throughput of the migration itself can also be tuned in the form of enabling UDT. UDT is a reliable, high-speed data transfer protocol based on UDP (User Datagram Protocol). UDT offers significantly higher speeds for transfer over high-latency, high-bandwidth networks. By default, data is transferred as plain text with the UDT protocol but you can choose to encrypt data. Select data encryption by choosing the Enable Encryption option on the destination vCloud Connector node. Note that selecting this option is the only way to enable encryption with UDT. The SSL setting only applies to HTTP(s) transfer and has no effect on UDT transfer. Be aware that there is some performance degradation associated with encryption.

With UDT, data transfer occurs over a dynamically-generated port on the source node and port 8190 on the destination node. Any firewall rules must allow for this type of connection for UDT-based data transfer. (When you copy data between a private cloud and a public cloud, data transfer is between a dynamically-generated port on the private cloud node and port 8190 on the public cloud node. Port 8190 must be open in the public cloud.) If you use a proxy server with UDT, communication between the local node and the proxy server occurs with two separate connections. For more information on tuning UDT and running UDT with a SOCKS5 proxy server, see the “Set UDT Properties” and “Using Proxy Servers” sections of the same guide. Instructions on how to configure NAT and firewall settings for individual

components are also found throughout the guide. In the interest of "unit testing" software components, refer to the 2105292 Knowledge Base article from the VMware Cloud Provider Program about potential testing of vCloud Connector connections.

4.3.4 vCloud Connector Offline Data Transfer Particulars

While tuning both storage and networking can get you so far, there are times the VMs and their disk files to be transferred online are too large for available resources. For this reason, vCloud Connector also provides the Offline Data Transfer method shown in Figure 14.

By executing the following mandatory steps, a normal vCloud Connector node is converted to an ODT Node:

1. Log in to vSphere Client.
2. Right-click the deployed ODT Node VM.
3. Click **Open Console**.
4. Press Enter to get the login prompt.
5. Log in as **root** with **vmware** as password.
6. Change the directory with the command `cd /opt/vmware/hcagent/scripts`.
7. Run the script `./configureSneakernetNode.sh`.

To verify that the vCloud Connector Node is configured as an ODT Node:

1. SSH to the ODT Node as **admin** using the password **vmware**.
2. Run the following Postgres query and verify it returns true:

```
psql hcs postgres
select config_value from hcs_config where config_key='is_sneakernet_node';
config_value
--------------
true
(1)
```

4.3.4.1 vCloud Connector ODT Customer Steps

To make the ODT process as smooth as possible, VMware recommends choosing and testing one or more Network Attached Storage (NAS) devices that support NFS. In this way, you can prepare the device so that it supports only the NFS mount at the customer data center with very specific instructions on how to mount and copy data. From the customer premises vCloud Connector Node run the following command:

```
mount -t nfs NFSDeviceIP:/nfs/ODT /data/ODT -o nolock
```

As shown in Figure 15, the available VMs, vApps and templates from the associated vCenter Server (more on configuring that in the following section) are displayed. After the Offline Data Transfer is initiated, the data is written to the mounted share. The data that is written to the file system is encrypted for security and the encryption key along with other credentials are stored (also encrypted) in the service provider vCloud Connector instance.

The following snapshots illustrate this process.

Figure 17. NFS Mount Path Relative to vCloud Connector Node

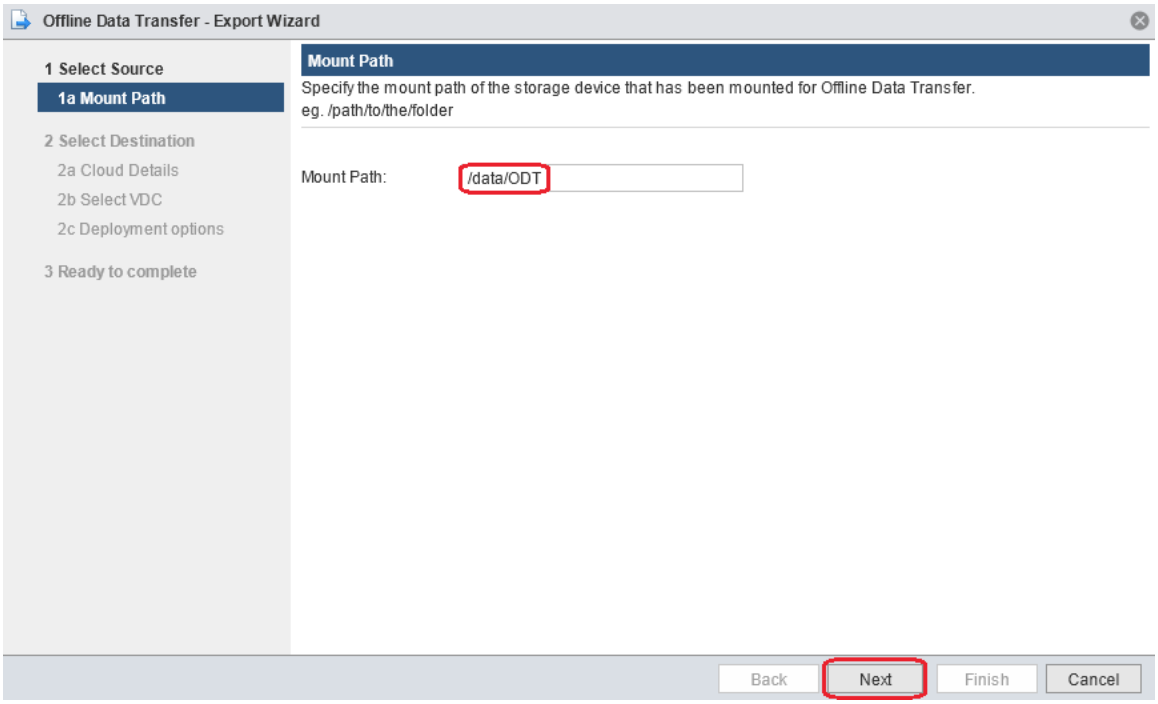

This NFS mount information does not impact the service provider mount or import process and can be decided by the customer. It is important that the NFS share be exported with read and write permission for EVERYONE.

Figure 18. VMware Cloud Provider vCloud Director Tenant Org URL and Credentials

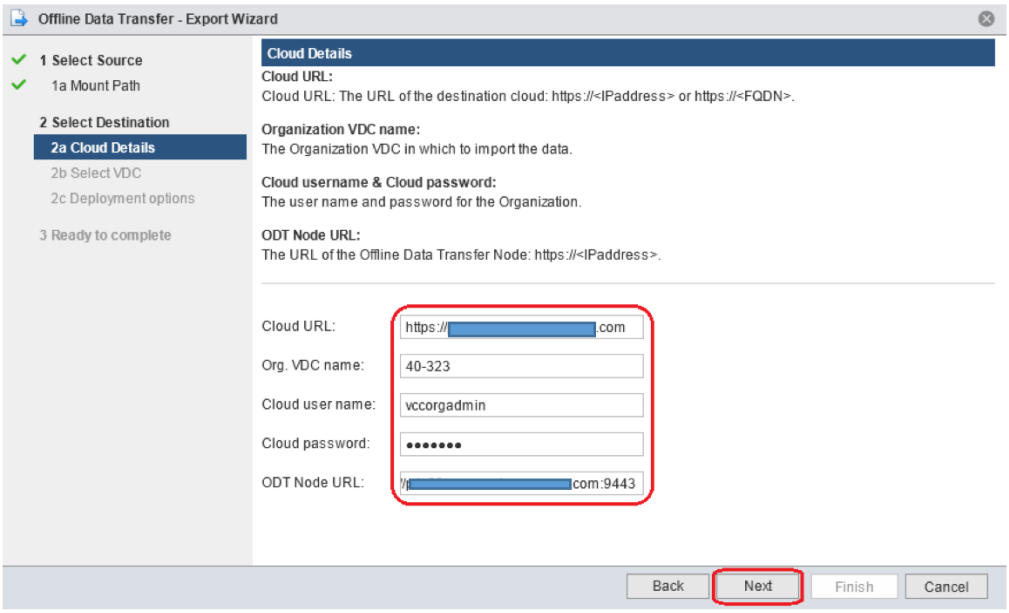

The credential must be an Organization Admin. This information is encrypted before being stored. The ODT Node URL comes from the VMware Cloud Provider.

Figure 19. VMware Cloud Provider vCloud Director VDC and Catalog

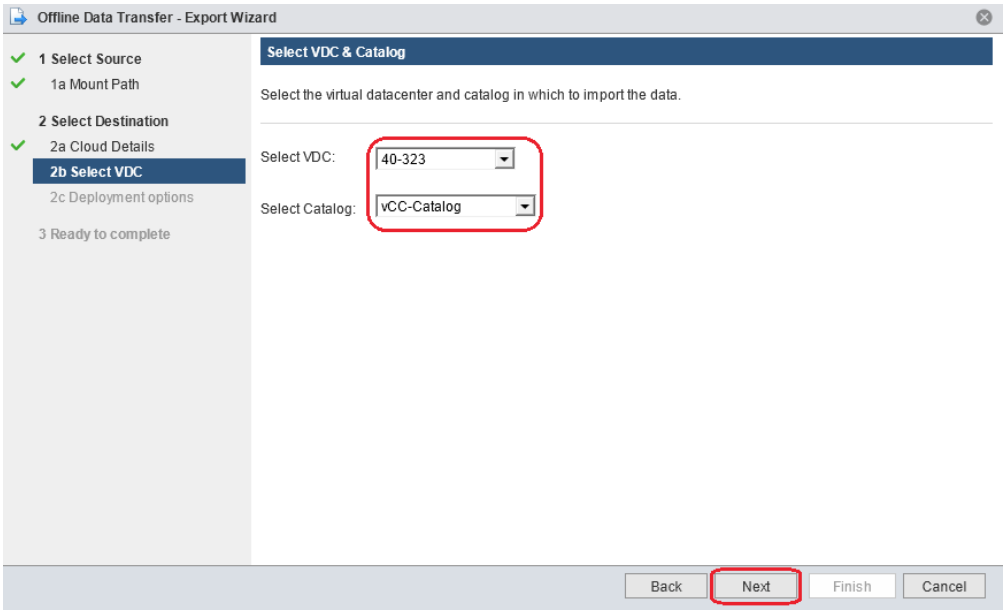

Caution Do not use the Public Catalog. Use the customer-specific Organization Catalog instead. If you use the Public Catalog all tenants will be able to download a copy of your VMs.

Figure 20. VMware Cloud Provider vCloud Director Network Pairing and Power State

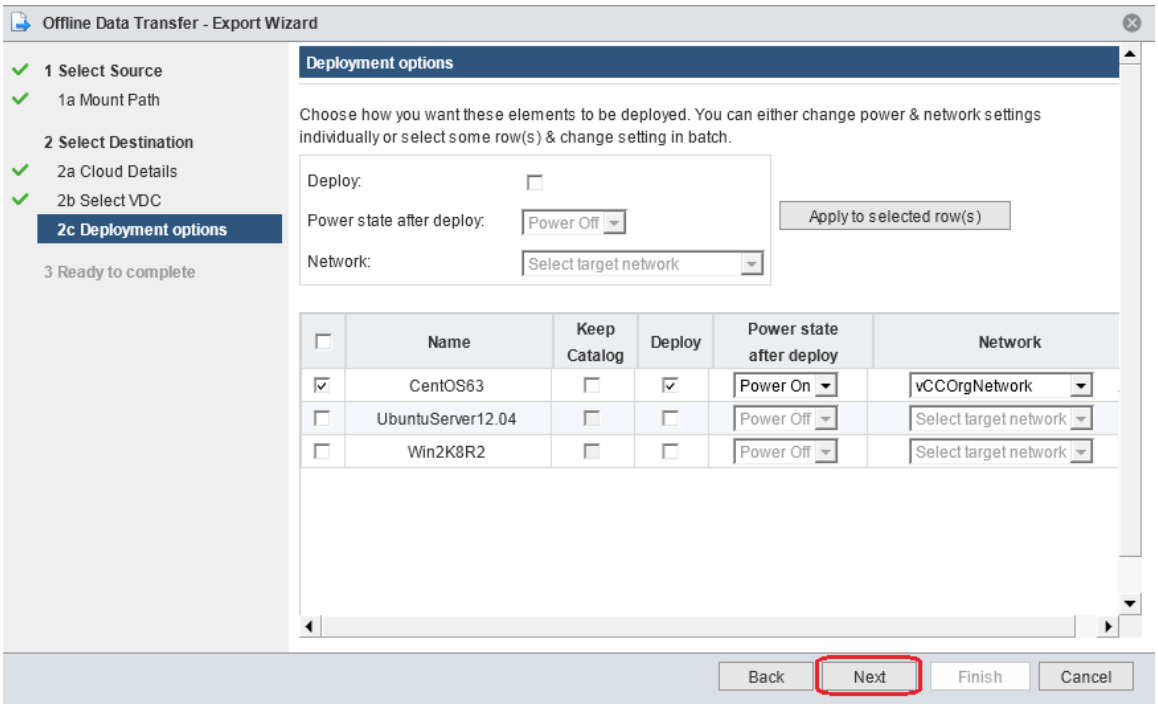

Note Any vApps that are exported without the **Deploy** check box selected are imported as multiple separate vApp templates, one for each VM that was part of the vApp. They are named SourcevAppName_RandomSequence_1, SourcevAppName_RandomSequence_2, and so on. ODT by design appends random numbers to the vApps and vApp templates names to avoid naming conflicts during the import process.

To avoid possible undesired situations on guest operating systems with non-present networks, the recommendation is to perform the following steps:

1. Before starting the export, create a copy (clone) of the VMs.
2. Disable and remove the vNICs from the cloned VMs guests and from vSphere.
3. Perform the export.
4. Re-create the networking to consume what is provided by vCloud Director.

Alternatively, creating a default vSphere port group such as “VM Network” on the target ESXi hosts in the target provider SDDC allows for a clean match.

If you choose to keep the template in the catalog by selecting the **Keep Catalog** check box, you are prompted with a warning that the template will consume disk space in the catalog. You can click **Next** to continue.

Figure 21. Keep Catalog Warning

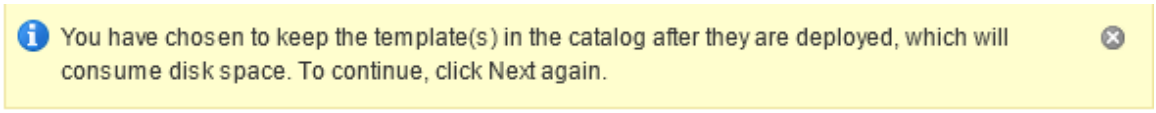

Figure 22. Confirmation of Details

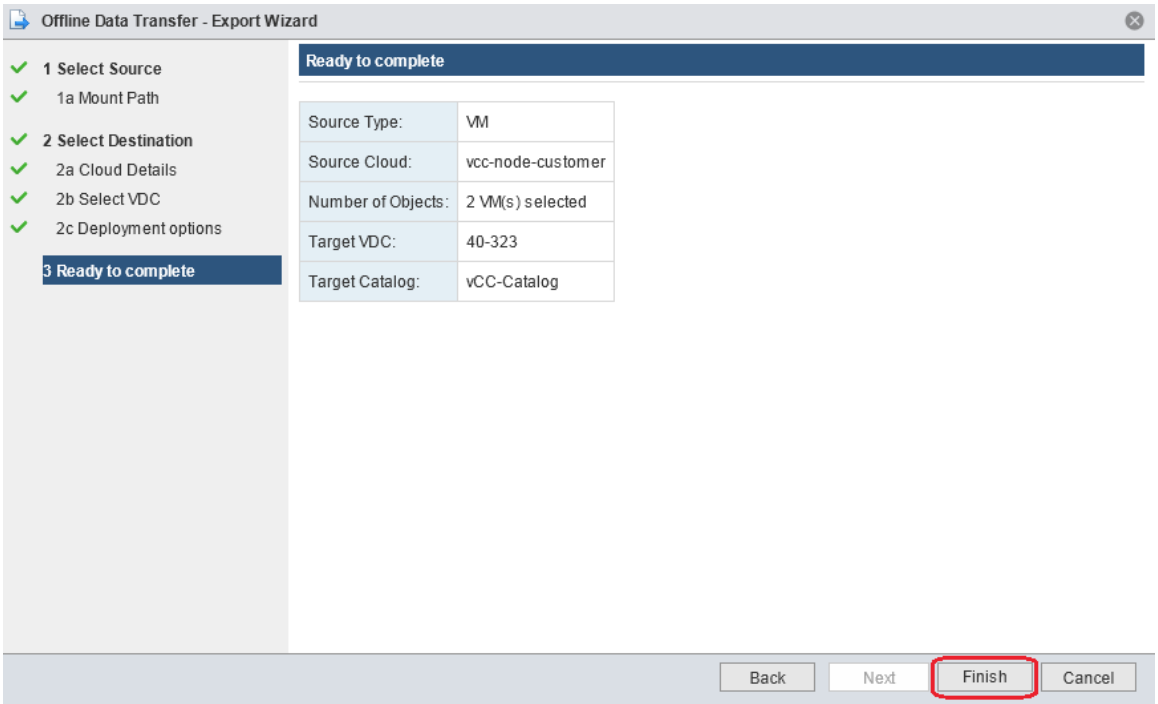

Figure 23. Export Progress

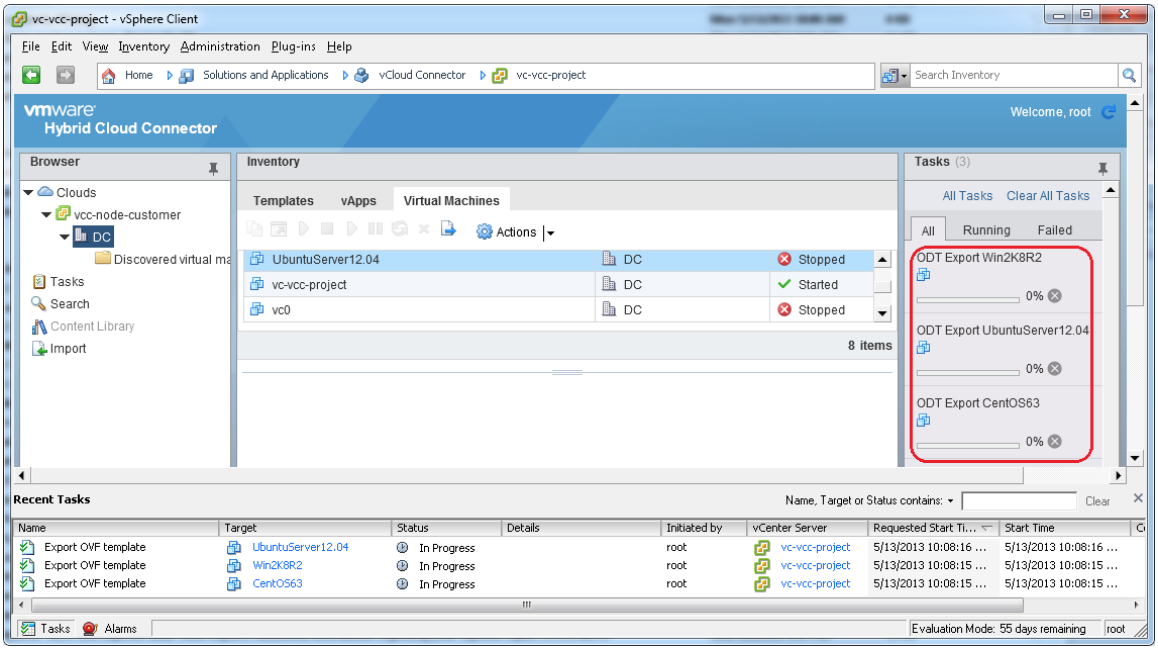

Figure 24. Export Completion

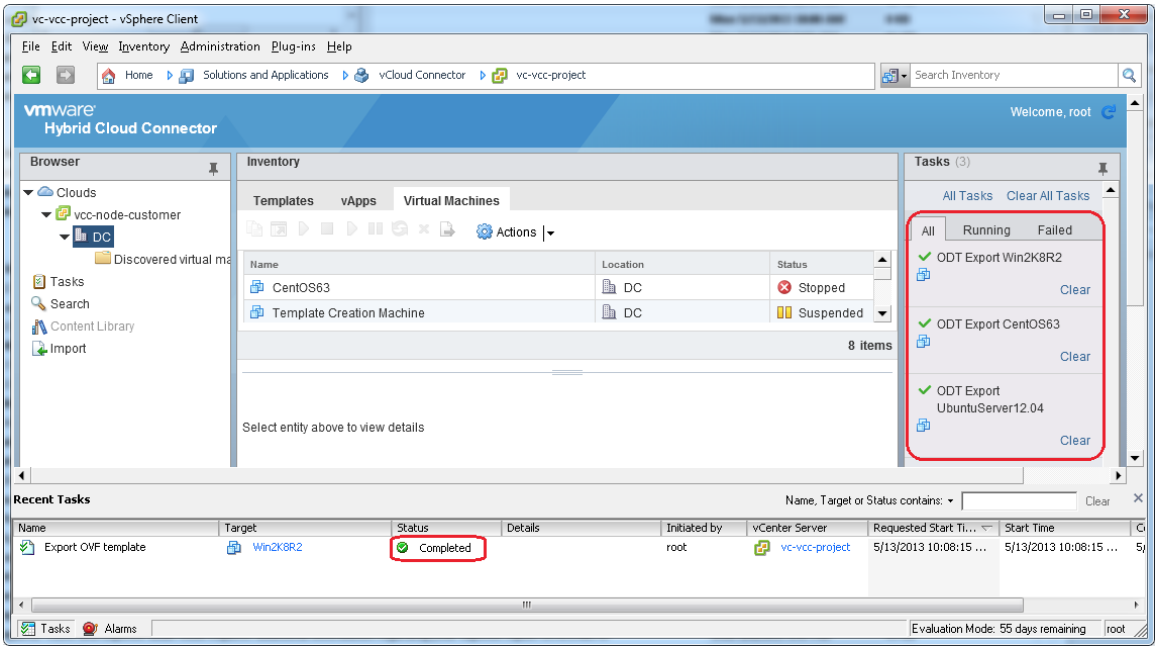

Figure 25. Export Verification on vCloud Connector Node NFS Mount

```
vccnode13218:/ # cd /data/ODT
vccnode13218:/data/ODT # ls
vCCOrg
vccnode13218:/data/ODT # cd vCCOrg/
vccnode13218:/data/ODT/vCCOrg # ls
CentOS63_628926252004894            Win2K8R2_581966314902725
UbuntuServer12.04_470210904445979
vccnode13218:/data/ODT/vCCOrg # cd CentOS63_628926252004894/
vccnode13218:/data/ODT/vCCOrg/CentOS63_628926252004894 # ls
_metadata  descriptor.ovf  disk-0.vmdk
vccnode13218:/data/ODT/vCCOrg/CentOS63_628926252004894 # cd ../Win2K8R2_58196631
4902725/
vccnode13218:/data/ODT/vCCOrg/Win2K8R2_581966314902725 # ls
_metadata  descriptor.ovf  disk-0.vmdk
vccnode13218:/data/ODT/vCCOrg/Win2K8R2_581966314902725 # cd ../UbuntuServer12.04
_470210904445979/
vccnode13218:/data/ODT/vCCOrg/UbuntuServer12.04_470210904445979 # ls
_metadata  descriptor.ovf  disk-0.vmdk
vccnode13218:/data/ODT/vCCOrg/UbuntuServer12.04_470210904445979 #
```

4.3.4.2 vCloud Connector ODT Service Provider Steps

The service provider instance of the vCloud Connector is used to initiate the copy process from the vCloud Connector Node mounted NFS drive and volume. vCloud Connector then posts the VMs to the specified vCenter Server datastore and uses the stored key to decrypt the data, and using the stored credentials, finally posts to the Org catalog. An exception to this data flow is that templates consisting of multiple VMs with the deploy flag unchecked (as shown in Figure 19) are posted directly to vCloud Director.

Figure 26. vCloud Connector Service Provider Import

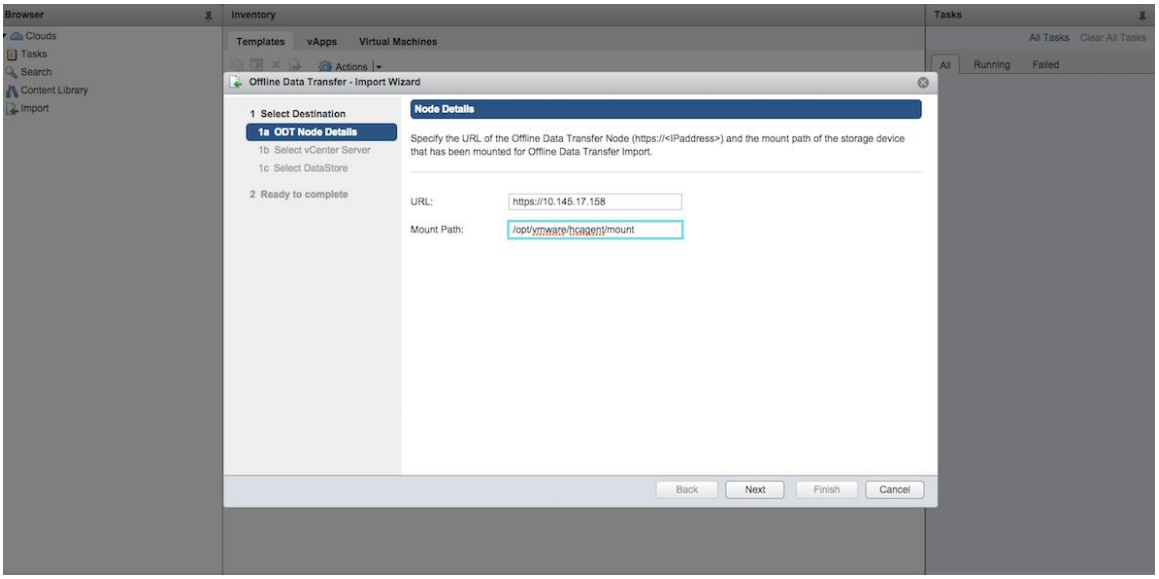

The service provider mounts the same NFS share on the shipped drive to the vCloud Connector node.

Figure 27. vCenter Server Credentials Supporting vCloud Director Org VDC

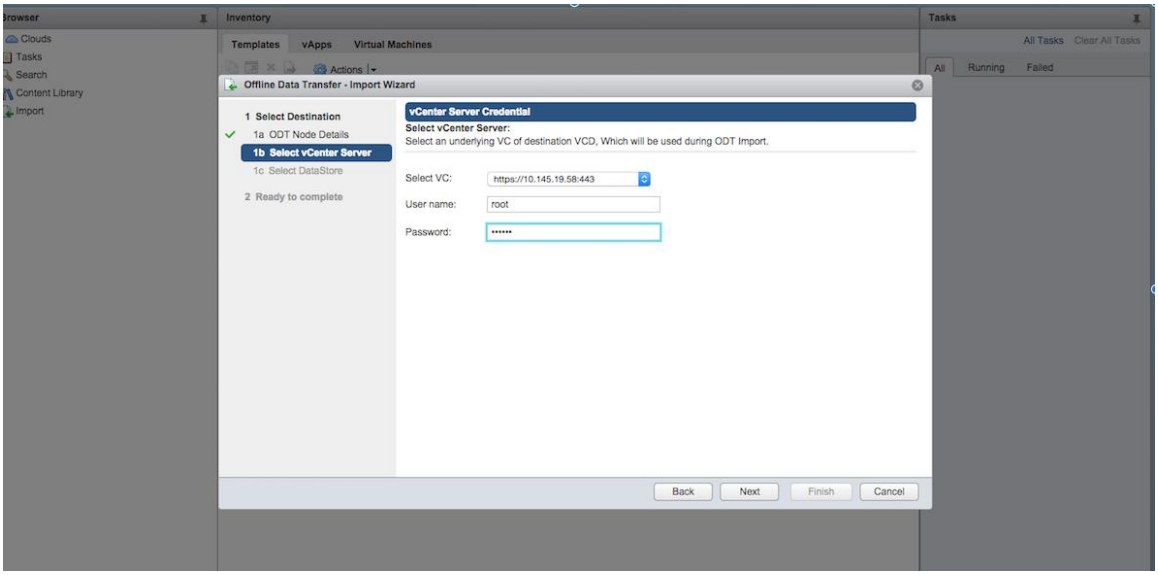

These credentials do not leave the vCloud Connector system and are encrypted before being stored.

Figure 28. vCenter Server Target Datastore

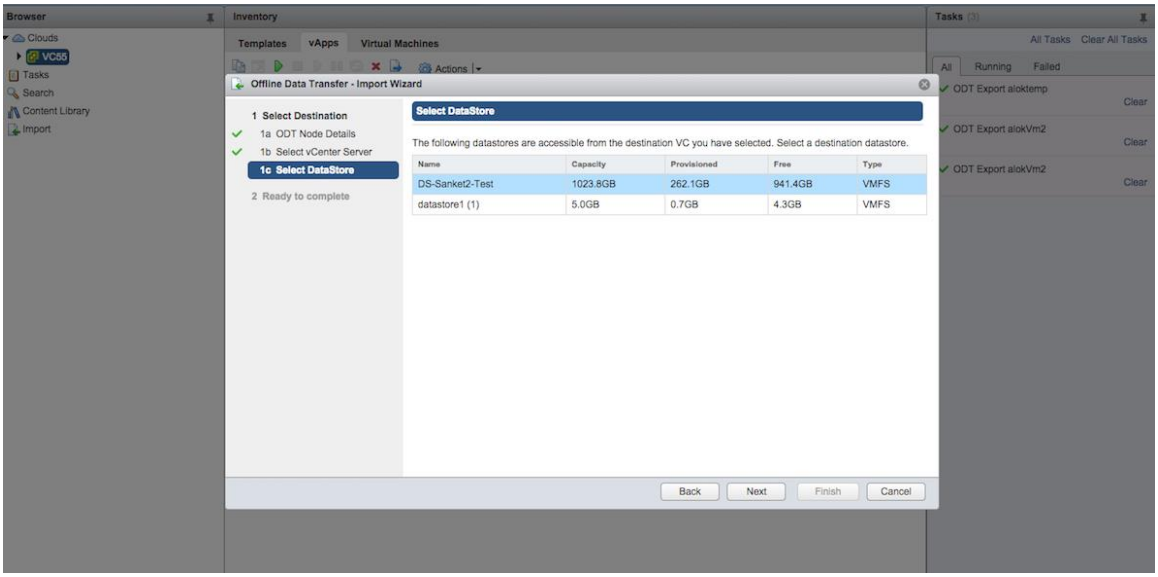

Figure 29. Verification of Details

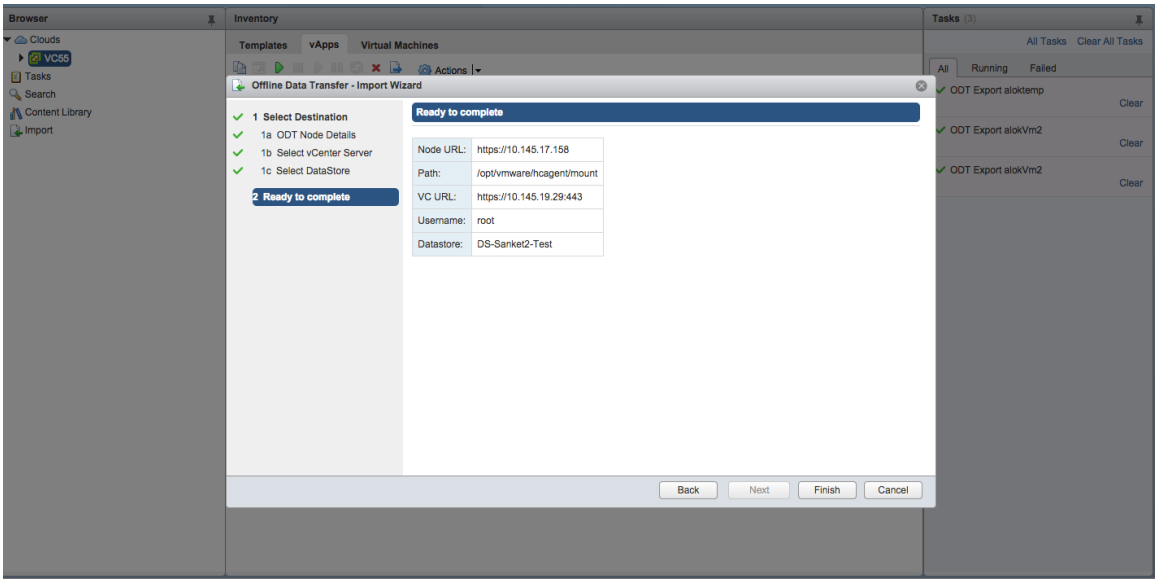

Figure 30. vCloud Connector Tasks View

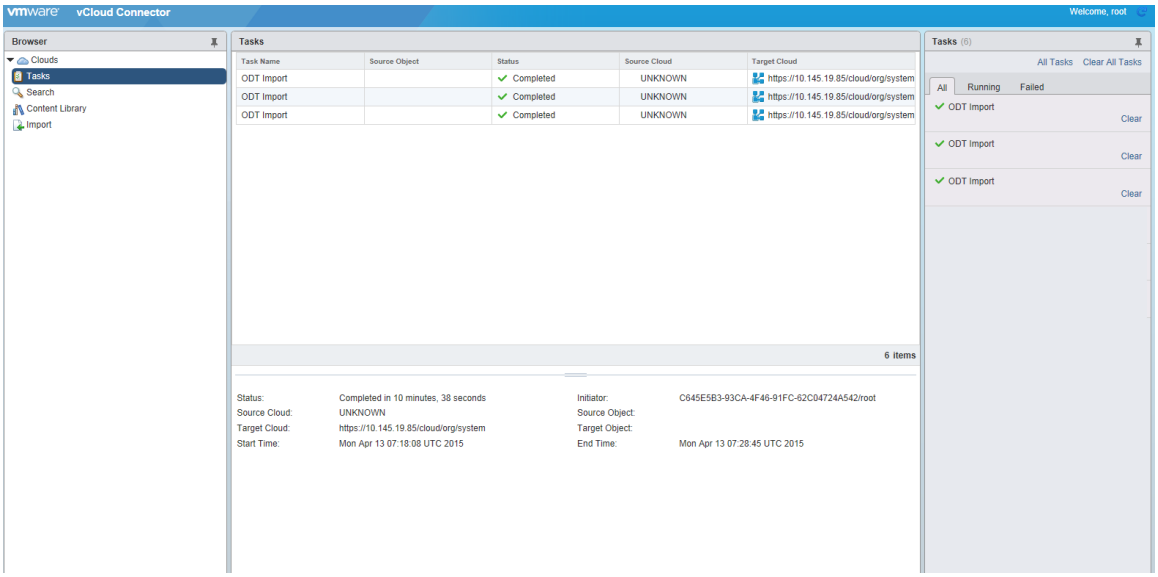

It is important that the ownership of resulting imported entities is changed to that of the Organization Administrator because they are created with the vCloud Director system administrator. After this step is done, notify the customer that their import is complete, and have them verify the import. After it is confirmed, recycle resources for use in another instance of the ODT process.

Migration Reference Implementation

This section describes an environment that has been configured to support a V2C migration leveraging vCloud Connector. This reference implementation is meant to portray the simplest configuration possible for each set of tools to perform the migration from the vSphere environment to the vCloud Director environment (shown in the following figure).

Figure 31. vCloud Director View of the Reference Implementation

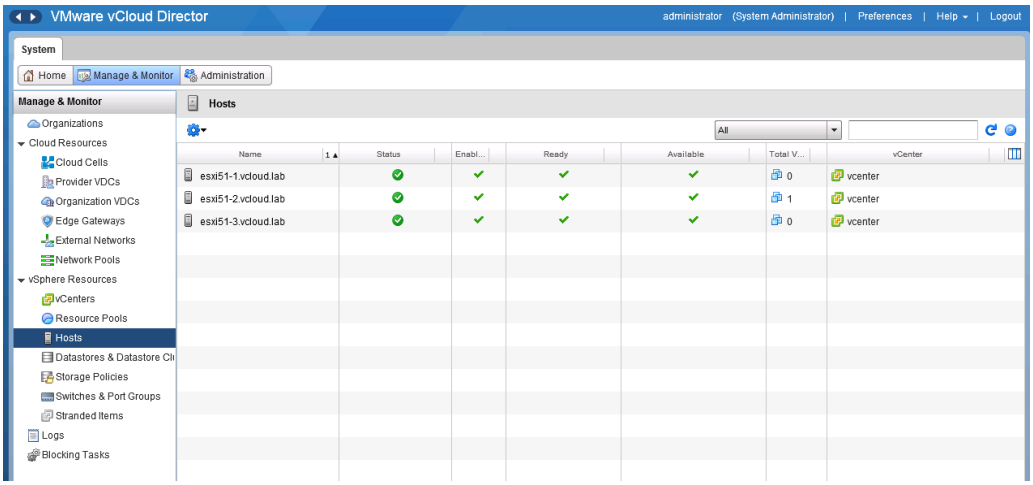

Certain components are required to reflect a production environment. Fully Qualified Domain Name (FQDN) through DNS resolution is one such requirement.

Figure 32. Active Directory DNS View of the Reference Implementation

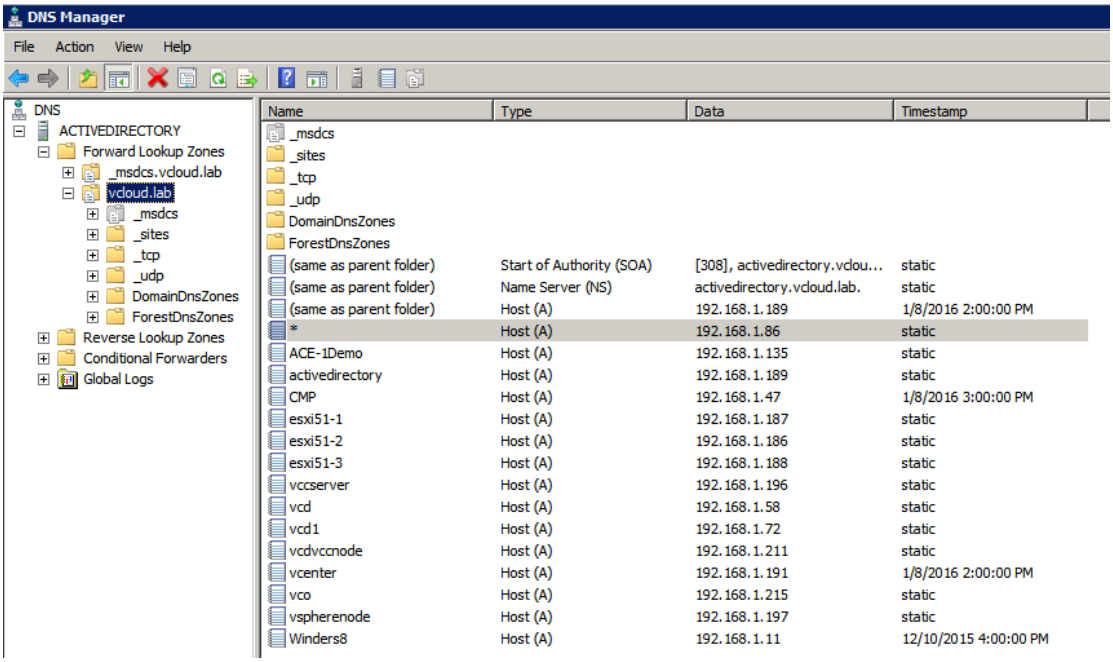

In this reference implementation, there are three ESXi hosts managed by a single vCenter Server that also hosts vCloud Director. We can provision Org VDCs from resource pools carved from the same vSphere cluster. The objective of the reference implementation is to make certain you understand how to

provide the required connectivity for the components to work in a simple environment prior to adding many of the complexities that might be required for production.

Figure 33. vSphere View of the Reference Implementation

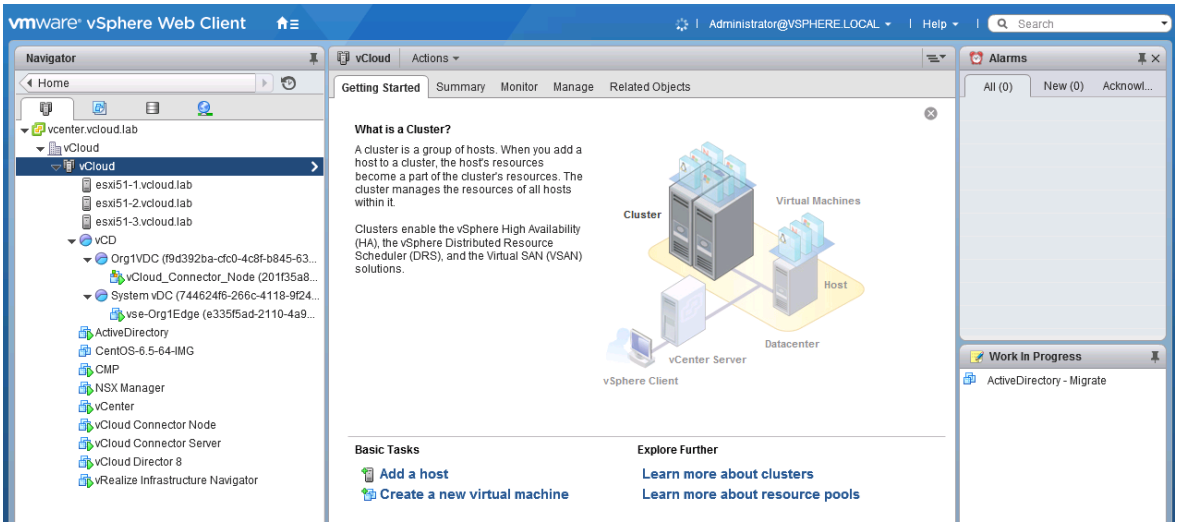

5.1 V2C Self-Service

The V2C self-service reference implementation consists of deploying vCloud Connector 2.7.1 so that it can be used to migrate a workload from a vSphere environment to a vCloud Director environment. This consists of deploying the vCloud Connector Server and two vCloud Connector nodes, one in the customer's data center where source VMs in vSphere exist and one in the target vCloud Director Organization where they will be migrated. After the appliances are deployed using OVF import (more information in the installation and configuration guide), you can reach the initial configuration screens on HTTP port 5480.

Figure 34. vCloud Connector Server on vSphere

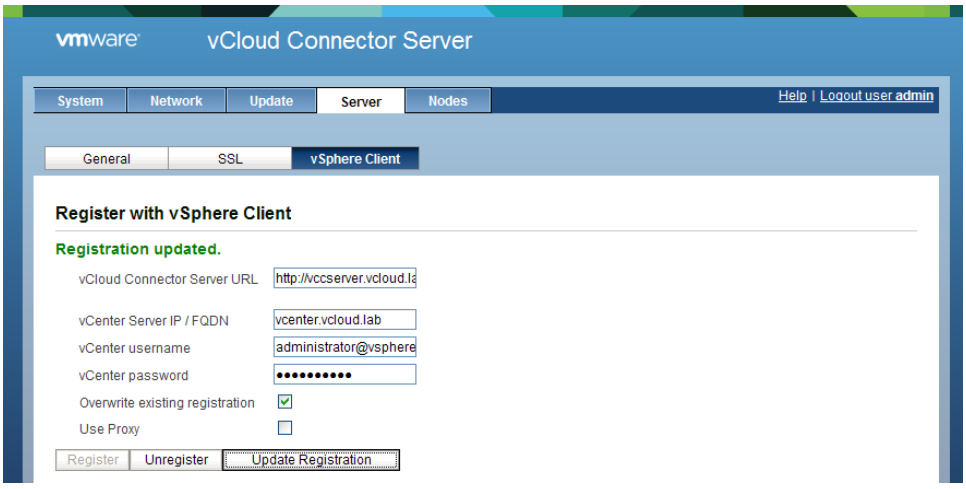

The vCloud Connector Server must be registered with vSphere to appear in the C# vSphere client as a "Solution".

Figure 35. vCloud Connector Node on vSphere

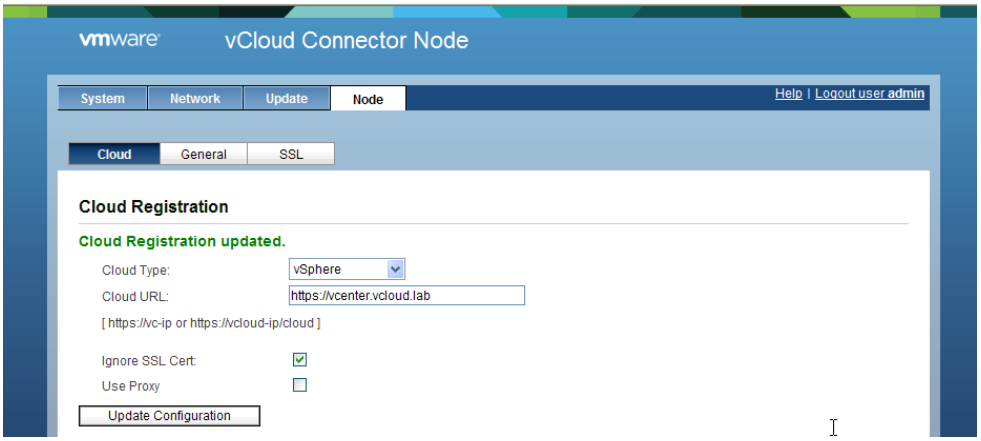

The customer then deploys and registers the same vSphere environment with their vCloud Connector node.

Figure 36. vCloud Director vCloud Connector Node

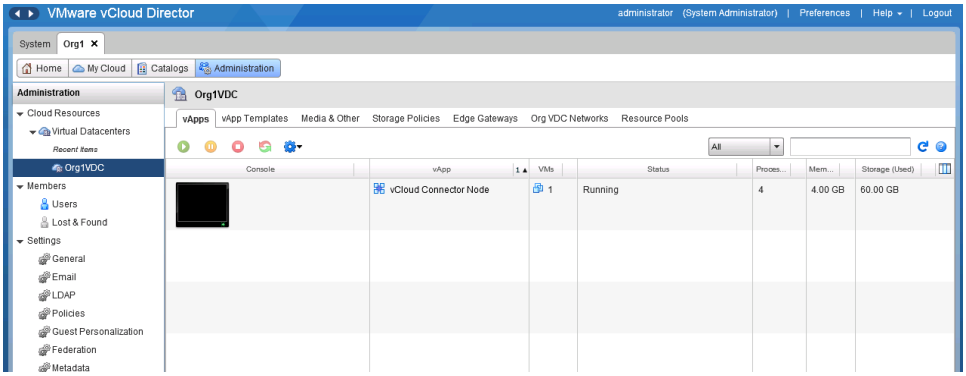

Within the vCloud Director environment, there must be another vCloud Connector node. In this reference implementation, the vCloud Connector node is placed inside the tenant (Org1).

Figure 37. vCloud Connector Node on vCloud Director

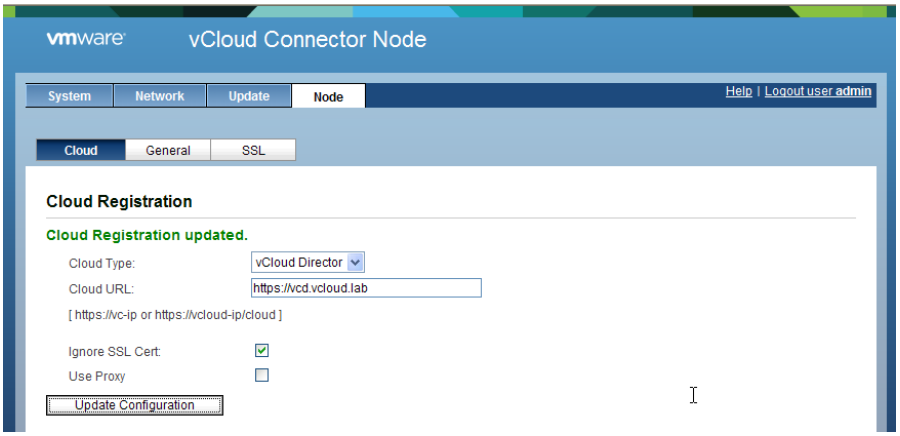

That node must be able to communicate with the vCloud Director URL.

Figure 38. Linking vCloud Connector Nodes to vCloud Connector Server

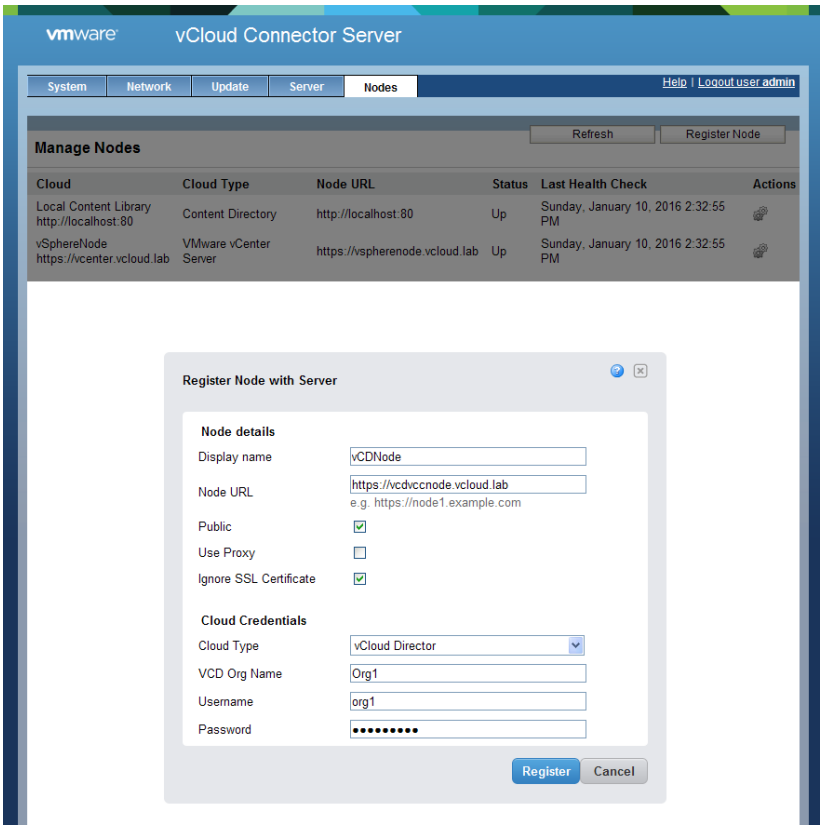

Finally, you must make both nodes available over the network to be registered in vCloud Connector Server.

Figure 39. vCloud Connector View in vSphere Client

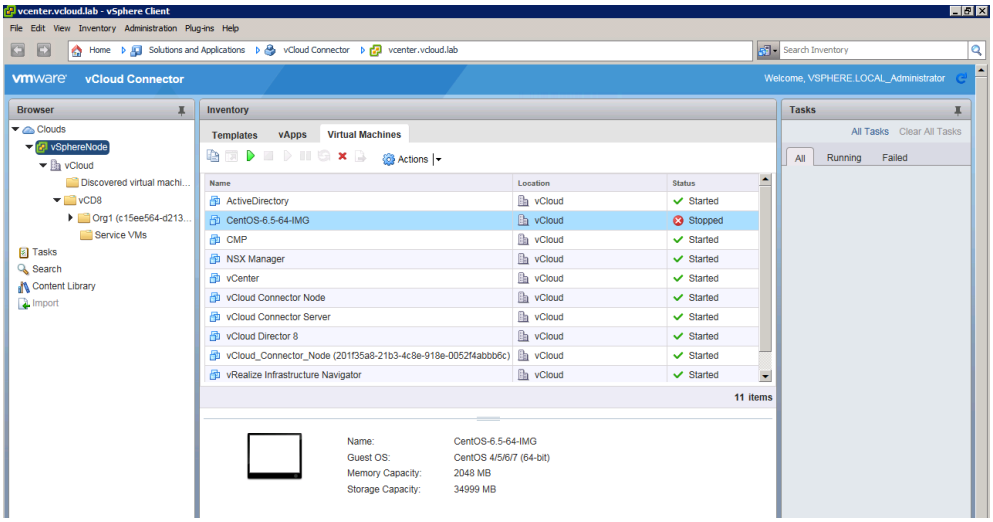

Now you can migrate workloads from vSphere to vCloud Director leveraging vCloud Connector.

Migration Operational Considerations

VMware recommends the migration Center of Excellence to plan these critical services so they can be delivered effectively and repeatedly throughout the VMware customer base with as little deviation from normal as possible. Useful training tools can be developed in a Migration COE, determining the best fit for each phase of the migration for each customer and application. As described by the reference implementation in this document, starting simple and planning functionality from end-to-end while expanding the reference environment is a wise approach. This allows you to match individual migration use cases mapped to particular source application attributes while allowing for infrastructure preparation. Because of the complexities involved with the tools, provisioning of the tools, infrastructure, and the migration process itself, enabling those stakeholders involved is also a critical feature of the COE.

After the COE is established with migration tools and use cases, you must map physical and logical dependencies across all of the layers and facets. This is the precursor for process analyses to reveal the limiting factors and is the first order in planning the migration operational models. Often one of the most prominent factors is the cutover time windows, which drive Recovery Time Objectives (how long before an application must be available in the new environment) and Recovery Point Objectives (how much time worth of transactions can be missed during the outage). While not exactly disaster recovery in nature, (where concepts such as RTO and RPO are typically addressed), in a migration it is possible to capture a delta of transactions from the source instance to the target, often times at a block level that is transparent to the application and database. The bandwidth and latency of the migration network as well as the size and frequency of potential changes to those datasets require careful analysis and testing. Verify that the cutover can meet the required timing objectives along with any mechanisms that are used for data integrity.

Like any other service, migration needs a clearing house, a funnel that looks at the types of migrations required and creates orders to both provision the tools and infrastructure necessary to accommodate them. It also requires detailed operational plans on how to support different types of migration with tiers of support for different migration tools as well as discreet areas of infrastructure supporting the migration. These must be “on call” during scheduled migration windows of pertinent types so that roadblocks that might otherwise back up other migrations and their dependencies can be resolved.

Updates for the vCloud Connector appliances can be retrieved from the administrative web interface on port 5480 and scheduled for download (with adequate connectivity). VMware has released several new versions of vCloud Connector to address critical vulnerabilities in both Java and glibc. Details on those releases are located in the following links:

https://www.vmware.com/support/hybridcloud/doc/hybridcloud_271_rel_notes.html

https://www.vmware.com/support/hybridcloud/doc/hybridcloud_272_rel_notes.html

Conclusion

The subject area of migration to the cloud is vast. Throughout the history of information technology, parties that migrated to new platforms in haste, without gaining a true understanding of how that new platform could actually change the life of their applications for the better, often never received the value those platforms promised. Cloud is no different.

This document has described several ways to leverage VMware tools to embark on that migration journey. Planning a Center of Excellence around those and other VMware Technology Partner tools is the first step in making sure everything has been addressed when asking our customers to start that journey. Make sure every customer and their corresponding application profiles has a “tailor made” plan for migrating to the cloud. This plan includes all of the hybridity features and control offered by the VMware SDDC running in the VMware Cloud Provider environments so they are free to run applications, or portions thereof, where it makes the most sense, even as that changes with time.

Looking at the market trends and opportunities, the potential to capitalize on the hybrid version of cloud rests in the balance of VMware delivering adequate migration services for customers. These services are not simply a way to get customers bits and bytes to run on our clouds. They offer a superior experience in the manageability, security, and value-add to IT operations. Getting migration right, every time, unlocks the true potential of hybrid cloud for all stakeholders involved.