

VMware vCloud® Architecture Toolkit™
for Service Providers

Workload Mobility and Disaster Recovery to VMware Cloud IaaS Providers

Version 2.9
January 2018

Adrian Roberts





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

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

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



Contents

- Introduction 5**
 - 1.2 Document Purpose and Scope 5
- Service Definition 6**
 - 2.1 Service Offering Overview 6
 - 2.2 Workload Mobility and Live Migration Services (Optional) 9
- Conceptual Architecture 10**
 - 3.1 Business Drivers 10
 - 3.2 Conceptual Architecture Solution Overview 10
- Designing the Solution 11**
 - 4.1 Logical Architecture Solution Overview 11
 - 4.2 Solution Architecture Bill of Materials 11
 - 4.3 Management Components Design 12
 - 4.4 Network Design 18
 - 4.5 Security Design 22
 - 4.6 Workload Mobility Design 23
 - 4.7 Disaster Recovery Design 24
 - 4.8 Failure Scenarios 30
- Operational Considerations 32**
 - 5.1 Cloud Service Provider Operations 32
 - 5.2 Cloud Service Tenant Operations 33
- Conclusion 36**



List of Tables

Table 1. Hosted Infrastructure Service Examples	6
Table 2. Direct Connect Service Examples	7
Table 3. VPN Service Examples	7
Table 4. Architecture Bill of Materials	11
Table 5. Example vCenter Server 6.0 Resource Requirements.....	13
Table 6. NSX Manager Resource Requirements.....	15
Table 7. vSphere Replication Resource Requirements.....	17
Table 8. vSphere Replication Resource Requirements.....	21
Table 9. Data Change Rate Compared with Link Speed and Potential RPO.....	25
Table 10. Site Recovery Manager Inventory Mapping Types.....	26
Table 11. Site Recovery Manager Operational Limits	29
Table 12. Provider Roles and Responsibilities	33
Table 13. End Customer Roles and Responsibilities.....	34

List of Figures

Figure 1. Conceptual Solution Overview.....	10
Figure 2. Logical Overview.....	11
Figure 3. vCenter Server Design	12
Figure 4. VMware NSX Management Services Design	14
Figure 5. Site Recovery Manager Server Topology.....	15
Figure 6. vSphere Replication Management Topology.....	17
Figure 7. Virtual Machine Routing Design	19
Figure 8. Active/Standby Data Center Architecture	20
Figure 9. NSX Controller Design.....	21
Figure 10. Security Design Example.....	23
Figure 11. Workload Mobility Design Overview	24
Figure 12. Example Inventory Mapping	27
Figure 13. Example Protection Groups.....	28



Introduction

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

VMware Cloud Providers are uniquely positioned to offer their services to the market and become a seamless extension of existing VMware enterprise customers' on-premises data centers. Having the capability to move workloads in and out of the customers' chosen cloud platform is a key factor for most enterprise customers in helping them maintain their existing investments with their on-premises applications, and to avoid lock-in to any one vendor or provider.

Disaster recovery and risk management is a key priority on all CIOs' agendas for 2016, and being able to offer a disaster recovery service between their own data center and a VMware Cloud Provider offers economies of scale, lower capital cost, and a lower operational cost model for protecting business critical applications and investments.

This document explores a potential solution where the end customer and the VMware Cloud Provider have common components running in their data centers to achieve a unified hybrid cloud experience offering seamless workload mobility and disaster recovery services to the end customers.

1.2 Document Purpose and Scope

The purpose of this document is to detail the solution architecture required for a VMware Cloud Provider Program managed hosting provider to offer a workload mobility and disaster recovery as a service offering to their customers.

The hosting solution architecture offers workload mobility and Disaster Recovery as a Service to the provider's end customers. A fictitious customer called Rainpole Inc. is used to demonstrate the end customer viewpoints of the solution.



Service Definition

The service is for a VMware Cloud Provider to offer a hosted Disaster Recovery as a Service and workload mobility solution to end customers. This enables the provider to build secure, reliable, cost-effective disaster recovery services for their customers to consume. The service consists of multiple, optional offerings that can be bundled together. This section details the typical offerings that a VMware Cloud Provider can offer to their end customers.

2.1 Service Offering Overview

The following section examines each area of the service offering and provides detailed examples of how the service can be presented to the VMware Cloud Provider's customers.

2.1.1 Target Virtual Infrastructure

Hosted, managed virtual infrastructure of varied sizes can be offered as a mobility and DR target for end customers to consume during planned migration/workload mobility (disaster avoidance) and disaster recovery scenarios. The hosted infrastructure can be sized by the provider based on the end customer's capacity and DR requirements. The target infrastructure will be dedicated to the end customer's requirements and must be sized appropriately.

2.1.1.1 Sizes and Specifications

The following sample virtual infrastructure sizing table describes three different service offerings—small, medium, and large. The service provider could also craft a custom offering which would go through a sizing exercise with the end customer to determine exact requirements.

In this example, there are three sizes of environment leveraging a single specification of host (2x 6 Core CPU @ 2.6 Ghz, 192 GB RAM and 1 TB of usable VMware vSAN™ local storage). The provider can offer different sizes of hosts to achieve economies of scale for the end customers.

Table 1. Hosted Infrastructure Service Examples

Size	CPU Cores and GHz	RAM GB	Storage GB
Small (2x hosts)	24 cores / 62 GHz	384 GB	2 TB
Medium (4x hosts)	48 cores / 124 GHz	768 GB	4 TB
Large (6x hosts)	72 cores / 186 GHz	1,152 GB	6 TB

Note Different hardware vendors offer different specifications. Make adjustments to suit your preferred hardware platform.

Note These specifications do not make adjustments for high availability (HA) reserved resources.

Where customers choose a custom option, you might need some sizing tools to determine the required amount of infrastructure.



2.1.2 Hybrid Networking Specification

When offering DR services to end customers who will be connecting from externally located data centers, it is important to verify that the provider can supply network connectivity for the end customer to manage the target infrastructure.

The VMware Cloud Provider can offer different services to their end customers with varying levels of service and cost.

2.1.2.1 Direct Connect Networks

Directly connected networks are provisioned by the end customer to the provider hosted data center / network. The provider can offer a connection point with a particular speed meeting specific criteria and the end customer can contract with a network service provider to provide the line.

The provider typically manages their side of the circuit, while the end customer is responsible for the its side of the circuit, along with the network provider. Existing telecommunications-based service providers might have offerings in this space that they can leverage end-to-end for their end customers by connecting the end customers to a local data center where the service provider has a point of presence.

The following table provides an example of two different direct connect service offerings.

Table 2. Direct Connect Service Examples

Circuit Bandwidth	Service Description
1 Gbps	1 Gbps connection from the end customer to the provider data center. This would typically be used for simple management traffic or replication of small quantities of data, and might be subject to larger RPOs.
10 Gbps	10 Gbps connection from the end customer to the provider data center. This would allow for larger amounts of data to be replicated to the target and would reduce the potential RPO available to the end customer.

Note The larger bandwidth allows for more data to be replicated simultaneously. However, latencies across the links are determined by the distances and hops required between sites, and must be considered when determining the RPOs and RTOs that the business requires and that the infrastructure can deliver.

2.1.2.2 VPN Services

The benefit of a VPN connection is that it is typically very quick to deploy over the Internet, allowing the provider to offer a connection in to their data centers immediately, without waiting for third-party network services companies to provision circuits. VPN services can also be used in conjunction with a direct network connection where required.

With VMware technologies, there are a number of different VPN services the provider can offer to their customers. The following table highlights the service offering examples.

Table 3. VPN Service Examples

VPN Service	Service Description
SSL VPN-Plus	With SSL VPN-Plus, remote users can connect securely to private networks behind a VMware NSX® Edge™ gateway. Remote users can access servers and applications in private networks.



VPN Service	Service Description
IPSec VPN	<p>NSX Edge supports site-to-site IPSec VPN between an NSX Edge instance and remote sites.</p> <p>NSX Edge supports certificate authentication, preshared key mode, IP unicast traffic, and no dynamic routing protocol between the NSX Edge instance and remote VPN routers. Behind each remote VPN router, you can configure multiple subnets to connect to the internal network behind an NSX Edge through IPSec tunnels. These subnets and the internal network behind a NSX Edge must have address ranges that do not overlap.</p>
L2VPN	<p>L2VPN allows you to configure a tunnel between two sites. Virtual machines remain on the same subnet in spite of being moved between these sites, which enables you to extend your data center. An NSX Edge at one site can provide all services to virtual machines on the other site.</p> <p>To create the L2VPN tunnel, you configure an L2VPN server and L2VPN client.</p>

For more information about the available VPN services from VMware, see the [VMware NSX documentation](#).

2.1.3 Virtual Machine Replication

Replication of virtual machines to the target data center is key for the disaster recovery solution. With a disaster recovery solution, there are typically two types of replication available:

- Storage-based replication – This type of replication enables the storage layer to replicate the data directly to another storage subsystem located on the target site. This typically requires the same specification of storage hardware at each side of replication.
- Virtual machine replication – VMware vSphere® Replication™ enabled through the vSphere hypervisor allows the provider to offer storage subsystem-agnostic replication between environments. The vSphere environments on each side must be vSphere Replication compatible.

With this solution provided to end customers who might have different storage architectures than those offered by the cloud provider, the best option is to leverage a vSphere Replication based architecture. In this reference architecture, Rainpole Inc. will use vSphere Replication as the replication service between the on-premises site and the VMware Cloud Provider.

The key replication services are as follows:

- VM protection granularity – The end-user can self-service configure replication of any virtual machine within their source VMware vCenter Server®.
- RPO definition – The replication engine (vSphere Replication) allows the end customer to self-service the required RPOs between 15 minutes and 24 hours for their source virtual machines.
- Point-in-time recovery – PITR allows the end customer to take periodic snapshots of virtual machines which can be recovered. To use this option, you must manually revert the machine back to the snapshot after recovery.

Note RPO policy compliance is dependent on the bandwidth available between the end customer source data center and the VMware Cloud Provider Program DRaaS hosting provider.



2.1.4 Automated Failover, Testing, and Migration Capabilities

The Disaster Recovery as a Service solution offers automated failover, testing, and planned migration capabilities to the end user's DR portfolio.

2.1.4.1 Disaster Recovery Scenario

In this scenario, automated disaster recovery is executed from the target site to recover protected workloads in the event of a disaster. The end user logs in to the recovery site and executes the DR recovery plan which brings up the virtual machines in write mode.

2.1.4.2 Planned Migration Scenario

The end user can execute a planned migration failover of workloads to the target hosted platform. This might be required due to planned on-premises data center maintenance to mitigate risk to business applications. This option is executed by the end user logging in to the recovery site and executing planned migration mode, which gracefully shuts down the source virtual machines, stops replication, and brings up the virtual machines in the recovery site.

2.1.4.3 Testing Scenario

Workloads can be failed over to the recovery site in test mode, which isolates the recovered workloads on to an isolated network, and maintains the existing replication schedule. Source virtual machines are unimpeached by the test.

2.2 Workload Mobility and Live Migration Services (Optional)

The end customer's vSphere administrator should have the ability to execute a long-distance VMware vSphere vMotion® action that migrates the virtual machine's running state between the on-premises data center and the VMware Cloud Provider's data center with limited or no disruption to the application running within the virtual machine.



Conceptual Architecture

3.1 Business Drivers

The key business drivers for a service provider to implement this solution are:

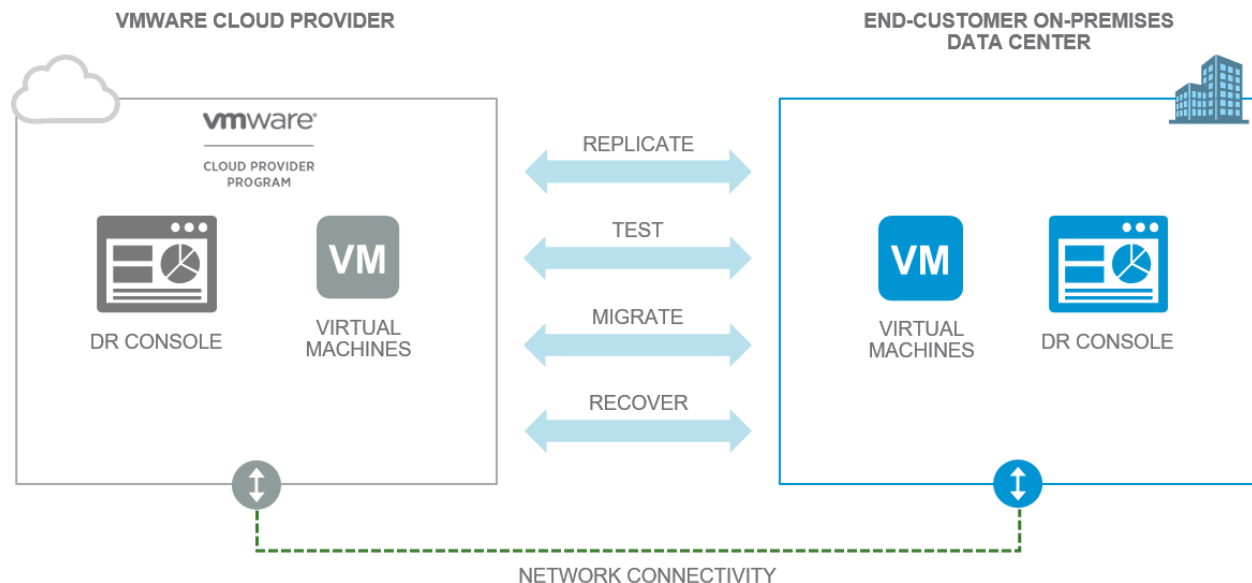
- Active/active data centers to run business application workloads.
- Workload mobility or running application services between availability domains.
- Business continuity of business critical applications for enterprise customers.
- Planned migration capabilities to offer disaster avoidance.
- Enterprise cost reduction in building and operating a disaster recovery solution for enterprise applications.

3.2 Conceptual Architecture Solution Overview

The solution must provide the following key functions to the business:

- Live migration of running virtual workloads from the end customer on-premises data center to the VMware Cloud Provider Program managed hosting provider's software-defined data center.
- Planned migration of multiple virtual workloads from the end customer on-premises data center to the VMware Cloud Provider Program managed hosting provider's software-defined infrastructure.
- Disaster recovery of multiple virtual workloads from the end customer on-premises data center to the VMware Cloud Provider Program managed hosting provider's software-defined data center and the reverse.
- Non-disruption testing of the disaster recovery scenarios between the end customer on-premises data center and the VMware Cloud Provider's software-defined data center.

Figure 1. Conceptual Solution Overview



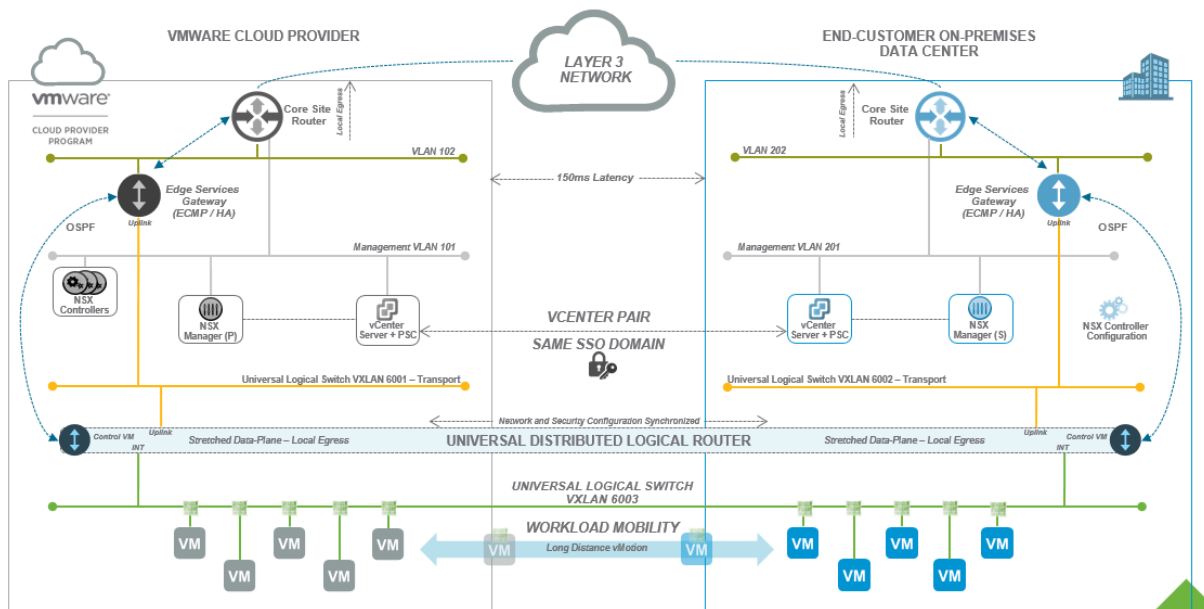


Designing the Solution

4.1 Logical Architecture Solution Overview

The following figure shows a logical view of a VMware Cloud Provider offering both workload mobility and Disaster Recovery as a Service to their end customer. The logical architecture highlights the minimum management virtual machine footprint and a simple network topology. The actual implementation of this solution might vary slightly depending on how the provider and end customer have architected their data centers.

Figure 2. Logical Overview



4.2 Solution Architecture Bill of Materials

In this architecture example, the end customer has VMware NSX implemented within their own data center. This could be either outsourced to the VMware Cloud Provider or managed internally by the end customer's own operations teams.

Table 4. Architecture Bill of Materials

Software Component – Provider	Software Component – Customer
<ul style="list-style-type: none"> VMware vSphere 6.0 vSphere Cluster(s) VMware vCenter Server + VMware Platform Services Controller™ 	<ul style="list-style-type: none"> vSphere 6.0 vSphere Cluster(s) vCenter Server + Platform Services Controller
VMware NSX® 6.2	VMware NSX 6.2
vSphere Replication 6.0	vSphere Replication 6.0
VMware Site Recovery Manager™ 6.1	Site Recovery Manager 6.1



4.3 Management Components Design

This solution has multiple management components that the provider must deploy and configure between their data center and the end customer data centers.

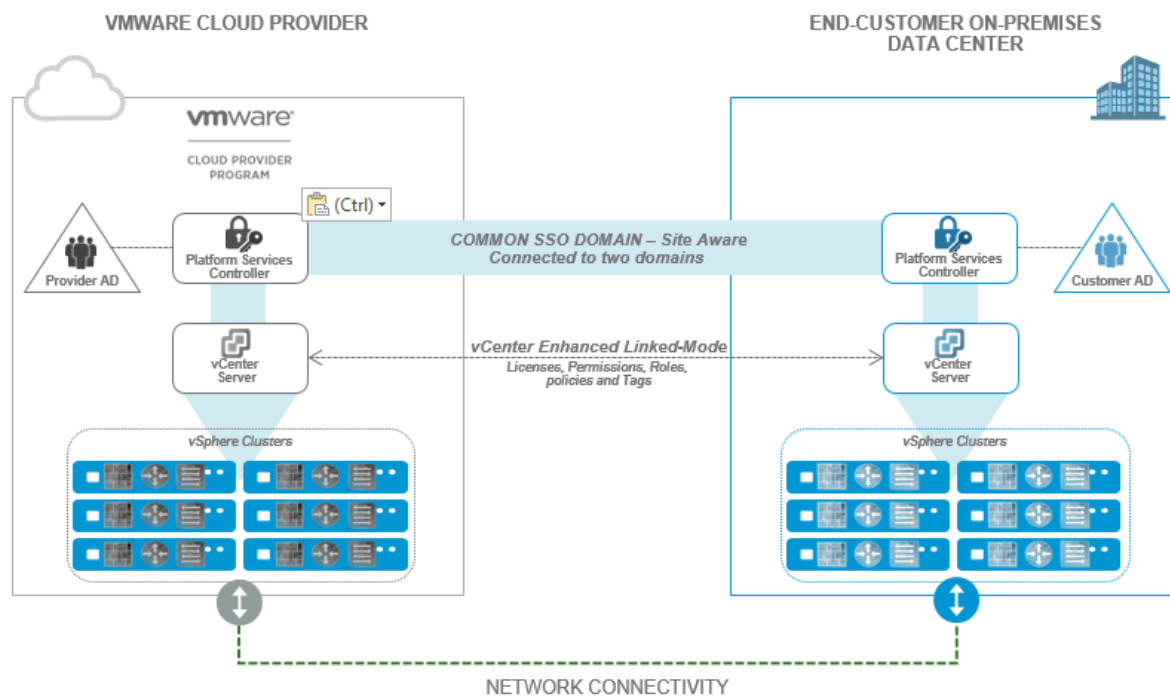
4.3.1 VMware vCenter Server Management Services

vCenter Server nodes provide the management plane for the VMware ESXi™ hosts within each data center. When deploying vCenter Server nodes, you have many options. This solution architecture will not address all design options for vCenter Server. For more information about VMware Cloud Provider Program deployment options for vCenter Server, see [Architecting a VMware vSphere Compute Platform for the VMware Cloud Provider Program](#).

The following example highlights some of the key design considerations for this solution:

- A vCenter Server is to be deployed in both data centers—one to manage the end-customer vSphere clusters and another to manage the VMware Cloud Provider's clusters.
- The vCenter Server nodes are configured using Enhanced Linked Mode (see [Enhanced Linked Mode](#)) to allow sharing of licenses, permissions, roles, policies, and tags between vCenter Server systems.
- SSO is configured to connect to multiple active directory sources—to the provider AD for management and to the end-customer AD for customer management and access.

Figure 3. vCenter Server Design



Note Enhanced Linked Mode is a requirement for long-distance vSphere vMotion migration to be executed through the VMware vSphere Web Client user interface. If this raises security concerns, separate SSO domains can be leveraged and vSphere vMotion operations can be executed through API only.



4.3.1.1 vCenter Server Resource Requirements

When deploying the vCenter Server nodes for both on-premises and off-premises infrastructure management, adhere to the best practices described in the Knowledge Base article [Installing vCenter 6.0 best practices](#).

The following table provides an example of a “small” (see [vCenter Server Appliance Requirements](#)) configuration with embedded Platform Services Controller. This configuration is capable of managing 100 hosts and 1,000 virtual machines.

Table 5. Example vCenter Server 6.0 Resource Requirements

Resources	Virtual CPUs	Virtual RAM	Virtual Storage
Small vCenter Server with embedded Platform Services Controller	4	10 GB	150 GB

Note The VMware Cloud Provider might need to adjust the specification based on the number of workloads they plan to manage for the end customer.

4.3.1.2 Scripted Installation

For VMware Cloud Providers, time-to-market speed is essential, and therefore the ability to quickly instantiate a consistent build of the management components for any cloud or hosting solution is critically important. vCenter Server can be scripted for easy installation and configuration to suit most environment requirements. This [blog article from a VMware engineer](#) offers a capability for the provider to script/automate the deployment of the vCenter Server services.

4.3.1.3 vCenter Server Database Requirements

vCenter Server can be installed in multiple models with varying database requirements. The appliance has an embedded vPostgres database, which would be the default option for most scenarios. See [Architecting a VMware vSphere Compute Platform for the VMware Cloud Provider Program](#) for more information and design options.

4.3.2 NSX Management Services

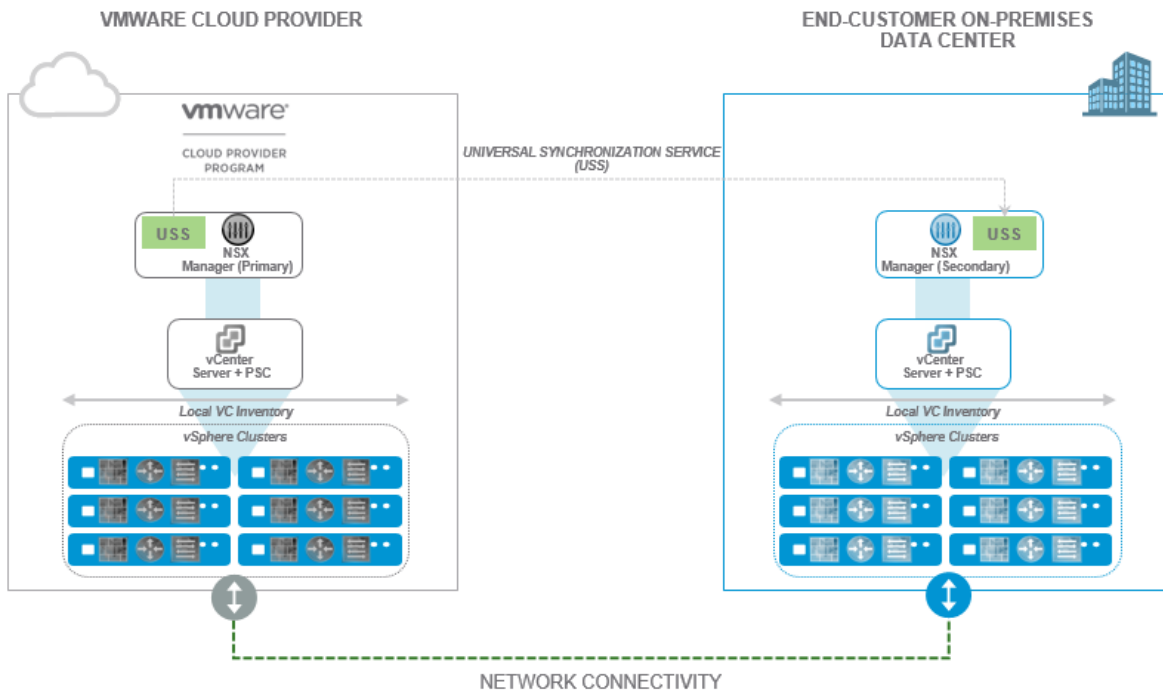
The VMware NSX Manager™ (see [NSX Manager](#)) is the centralized network management component of VMware NSX and is installed as a virtual appliance on any VMware ESX® host in your vCenter Server environment. It provides an aggregated system view.

One NSX Manager maps to a single vCenter Server environment and multiple NSX Edge, VMware vShield Endpoint™, and VMware NSX data security instances.

For the mobility and DR solution, there will be an NSX Manager that must be created at both the provider data center and the end customer’s data center. The following figure highlights an example design for the deployment of NSX Manager nodes.



Figure 4. VMware NSX Management Services Design





4.3.2.1 NSX Manager Resource Requirements

Each NSX Manager has a set specification as detailed in the following table.

Table 6. NSX Manager Resource Requirements

Resources	Virtual CPUs	Virtual RAM	Virtual Storage
NSX Manager	4	12 GB	60 GB

Note For VMware NSX system requirements specifications, see [System Requirements for NSX](#).

4.3.2.2 NSX Manager Scripted Installation

Installation of NSX Manager instances can be scripted to enable faster time to market and a repeatable process for onboarding new customers. While this topic is out of scope for this document, there are various solutions available from VMware and external blog posts.

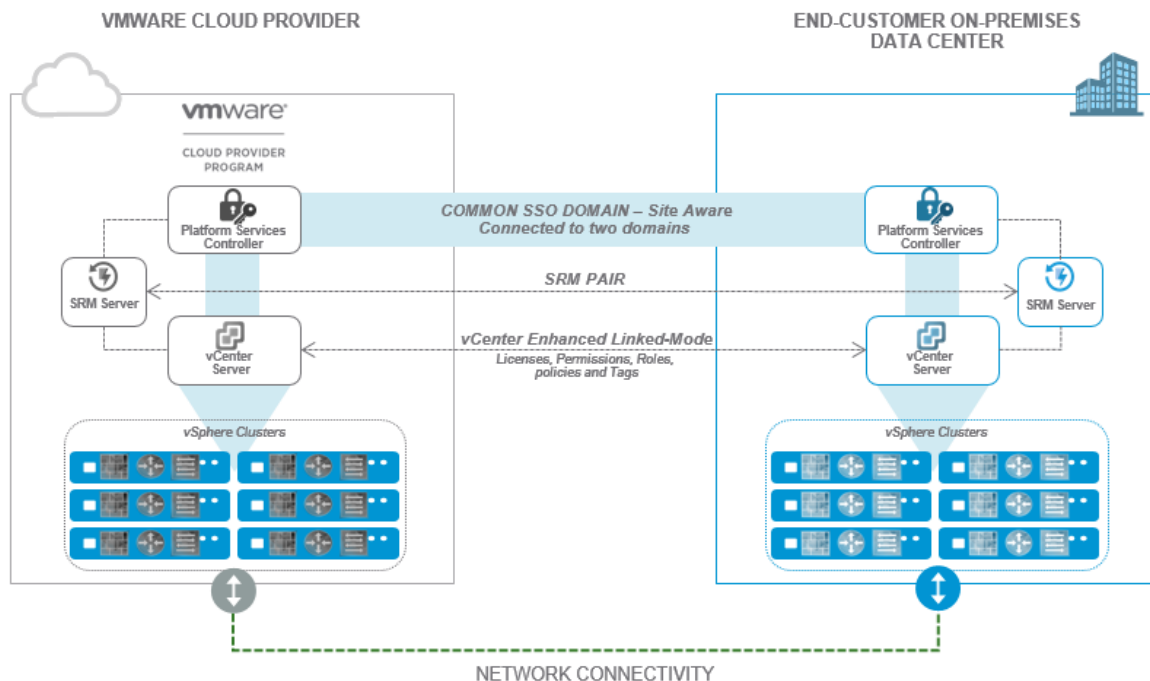
4.3.3 Site Recovery Manager Management Services

When deploying Site Recovery Manager, there are several models to consider based on the type of vCenter Server and Platform Services Controller deployment you have configured within your environment. For consideration of different deployment models, see [Site Recovery Manager and vCenter Server Deployment Models](#).

The architecture model that this document focuses on is a simple [Site Recovery Manager in a Two-Site Topology with One vCenter Server Instance per Platform Services Controller](#). In addition, the sites are configured for Enhanced Linked Mode (single SSO domain).

The following figure highlights the Site Recovery Manager server topology between the two sites.

Figure 5. Site Recovery Manager Server Topology





4.3.3.1 Site Recovery Manager Database Requirements

The Site Recovery Manager server requires its own database, which it uses to store data such as recovery plans and inventory information.

Site Recovery Manager provides an embedded vPostgres database that you can use with minimal configuration. You can select the option to use the embedded database when you install Site Recovery Manager. The Site Recovery Manager installer creates the embedded database and a database user account according to the information that you specify during installation.

You can also use an external database. If you use an external database, you must create the database and establish a database connection before you can install Site Recovery Manager.

For more information, see [About VMware Site Recovery Manager Installation and Configuration](#).

4.3.4 vSphere Replication Management Services

vSphere Replication 6.0 has two components that must be deployed to the management domain—vSphere Replication Manager Appliance and vSphere Replication Server.

4.3.4.1 vSphere Replication Manager Appliance

The vSphere Replication Manager Appliance is the main component of vSphere replication management services and it registers with the vCenter Server as a solution.

Each Manager Appliance and, therefore, each vCenter Server can support up to 2,000 replications. This is achieved by leveraging 10 replication server instances.

Note Only one vSphere Replication Manager Appliance can be registered with a vCenter Server. They have a 1:1 relationship. Additional scale can be achieved by leveraging additional vSphere Replication Server Appliances—up to 10 per vSphere Replication Manager to give the scale of 2,000 replications.

4.3.4.2 vSphere Replication Server Appliance

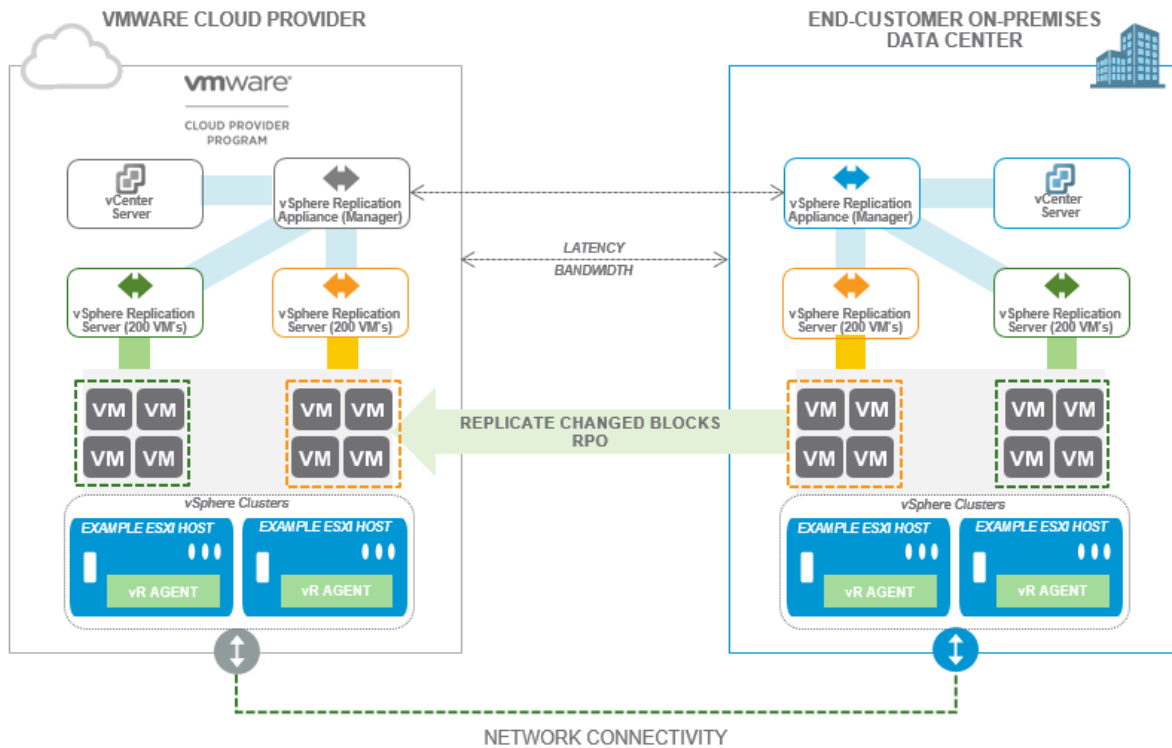
The vSphere Replication Server Appliance is responsible for the replication job, which includes gathering the data from the source ESXi host and transferring the data to the target infrastructure. The default installation of a vSphere Replication Server Appliance contains a vSphere Replication Server, which can effectively support 200 replications.

Note If the design only requires replication of less than 200 VMs, the simplest architecture to use is the embedded vSphere replication service within the Manager Appliance.

The following figure provides an example vSphere replication management architecture between an end customer data center and a VMware Cloud Provider Program data center. The architecture highlights a separate vSphere Replication Manager Appliance and vSphere Replication Server Appliance for larger scale.



Figure 6. vSphere Replication Management Topology



The following points are shown in this example:

- The vSphere Replication Manager Appliance is registered with the vCenter Server at each site. The manager can support up to 200 replications.
- Each additional vSphere Replication Server Appliance is registered with the vCenter Server and Manager Appliance.
- The vSphere Replication Server Appliances share the replication load of the virtual machines—200 replications per appliance, with maximum of 10 for 2,000 VM replications per vCenter Server.

Note Take into account other factors, such as network bandwidth and latency, when calculating the maximum number of replications and the impact on the potential RPO. You can leverage the vSphere Replication Calculator [here](#).

4.3.5 vSphere Replication Resource Requirements

The following table highlights the minimum recommended resource requirements for both the vSphere Replication Manager and Server Appliances.

Table 7. vSphere Replication Resource Requirements

Resources	Virtual CPUs	Virtual RAM	Virtual Storage	Scale
vSphere Replication Manager Appliance	2	16 GB	2 GB	200 VMs



Resources	Virtual CPUs	Virtual RAM	Virtual Storage	Scale
Additional vSphere Replication Server Appliance	2	768 MB	2 GB	200 VMs / Server Maximum 10 (9+Manager) per vCenter Server

4.4 Network Design

This section of the document examines the network design elements that are required for a successful deployment of this solution. This is not meant to be a design guide for VMware NSX or physical networking topologies, but rather to determine how the products could be implemented to achieve the business drivers identified in Section 3.1, Business Drivers.

4.4.1 Data Center Connectivity

The connectivity between the end customer and the VMware Cloud Provider is important to be able to offer a consistent service level.

The provider can either offer their direct connected services or VPN services to the end customers. This solution example assumes that the data centers are connected by an established WAN/MPLS connection offered between the VMware Cloud Provider and end customer.

It is important that the following services can communicate across the Layer 3 network with their counterparts on the target sites (provider and end customer):

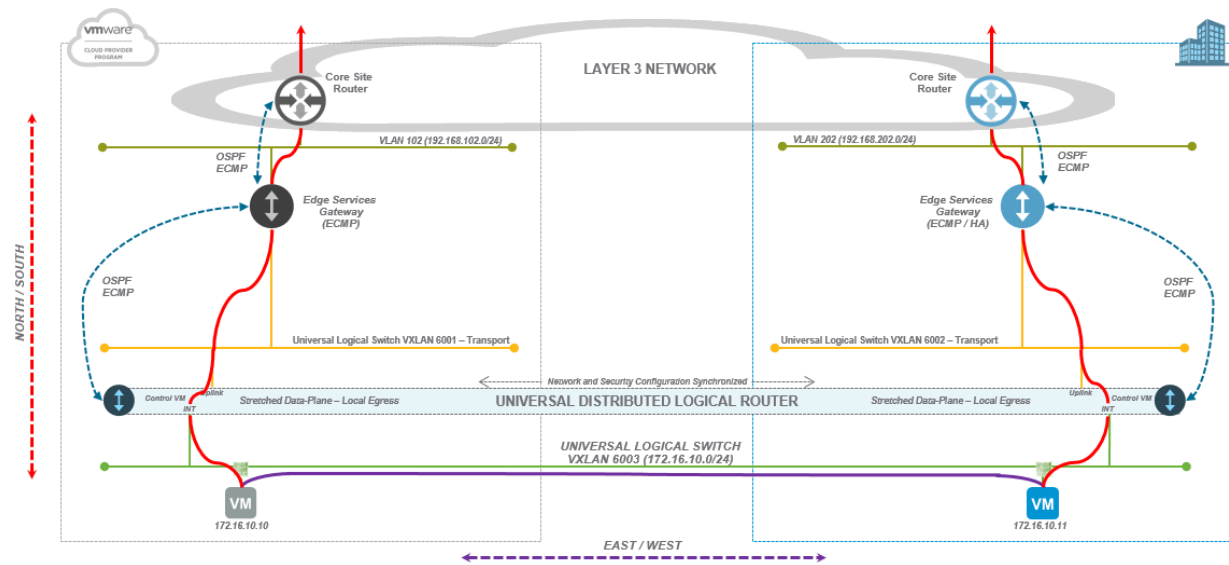
- NSX Manager
- VMware NSX Controller™ nodes
- vCenter Server nodes
- Platform Services Controller instances
- vSphere Replication Managers and Servers
- Site Recovery Manager servers

Note It is also important that the MTU is at least 1,600 end-to-end between data centers to support the additional packet size for VXLAN.

4.4.2 Data Center Routing Design

The routing design for cross-vCenter NSX enables virtual machines in both data centers to reside on the same L2 subnet. The traffic flow between virtual machines that are connected to the same logical switch will traverse the Layer 3 data center interconnect network, but maintain the encapsulation within the L2 subnet. The architecture shown in the following figure represents an active/active data center, where traffic will egress the local data center only.

Note The specification of the Layer 3 network is out of scope for this solution design.


Figure 7. Virtual Machine Routing Design


4.4.2.1 North-South Routing

All new and existing networks that are connected to the Universal Distributed Logical Router (UDLR) will be connected and propagated to the edge services gateway devices within each data center. This is achieved through the use of the Open Shortest Path First (OSPF) dynamic routing protocol.

North-South routing for virtual machine traffic will traverse the edge services gateway devices which are connected to the core network through a VLAN connection. This gives complete routed connections in and out of the data center core network.

The default gateway for virtual machine traffic will always be within the site that the virtual machine is located in. This is due to the LocalID that is presented to the ESXi hosts managed by the respective NSX Manager. With local egress turned on, there is a controller VM for the UDLR in each data center that contains the default gateway for the logical switch traffic.

4.4.2.2 East-West Routing

East-West routing is done within the Layer 2 domain for the logical distributed router, which is configured to span across the two data centers. This means that although traffic does traverse the data center connections, it is contained within the same broadcast domain and passed over the VXLAN transport network.

4.4.2.3 Local Egress

With VMware NSX 6.2 and UDLR, there is an option to enable egress optimization to avoid the potential for hairpinning of the WAN network for the default gateway. For more information about egress optimization, see [Local Egress](#).

This design leverages local egress so that traffic exits the data center through the local routing infrastructure even after a migration. This can cause some potential issues with client server communication if the physical network cannot manage the ingress routing to the workloads.



4.4.2.4 Ingress Optimization

Ingress of traffic to the virtual machines when they fail across data centers is extremely important to consider, because without a solution in place, traffic destined for virtual machines on Site A can be returned through Site B. This is often referred to as asymmetric routing.

Route ingress optimization can be achieved in a number of ways and can be dependent on the physical networking capabilities. One solution is to leverage route injection through APIs during an event, such as a long-distance vSphere vMotion migration. See [Ingress Optimisation with NSX for vSphere](#).

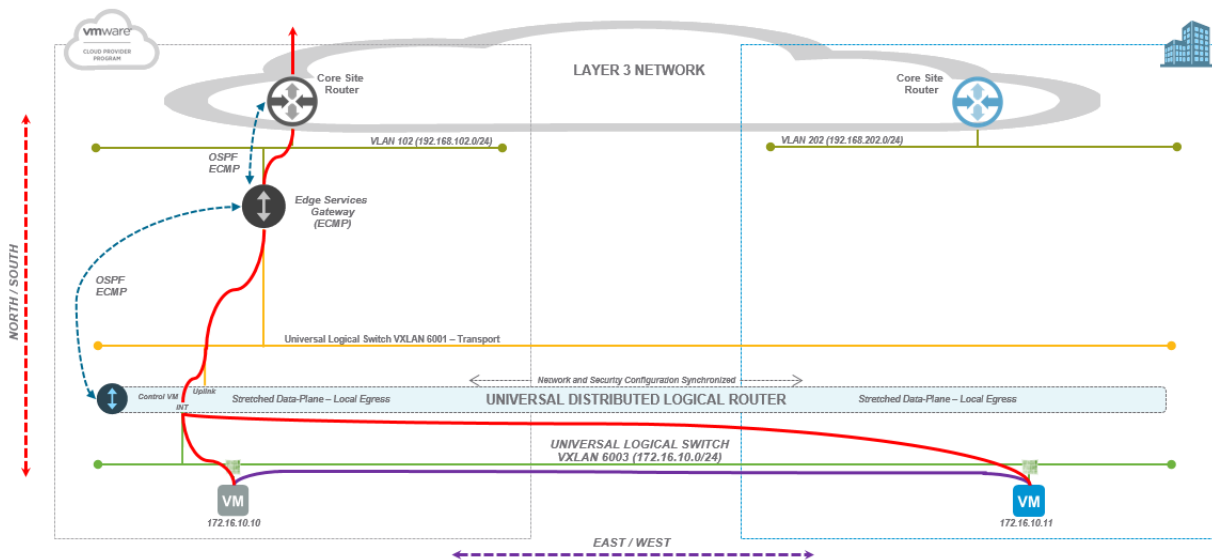
4.4.2.5 Stateful Services Routing

If you are running stateful services within your data center on the NSX Edge services gateway appliances, it is important to understand the impact to these services when you move workloads from one data center to another. For example, if you run a NAT service on an edge services gateway to control access to an application service in data center A, and that workload is migrated to data center B, the inbound traffic from the client will go through the same NAT service in data center A, traverse the WAN, and egress through data center B. This will break the NAT service.

4.4.3 Alternative Network Architectures

An alternative architecture is to set up the data center to be active/standby in terms of networking capabilities. This means all traffic will egress from one data center only (where the edge services are placed), which could lead to excessive hairpinning of the WAN network. If the WAN is reliable enough, bandwidth is sufficient, and latency is low enough, this might be a better option which reduces complexity.

Figure 8. Active/Standby Data Center Architecture



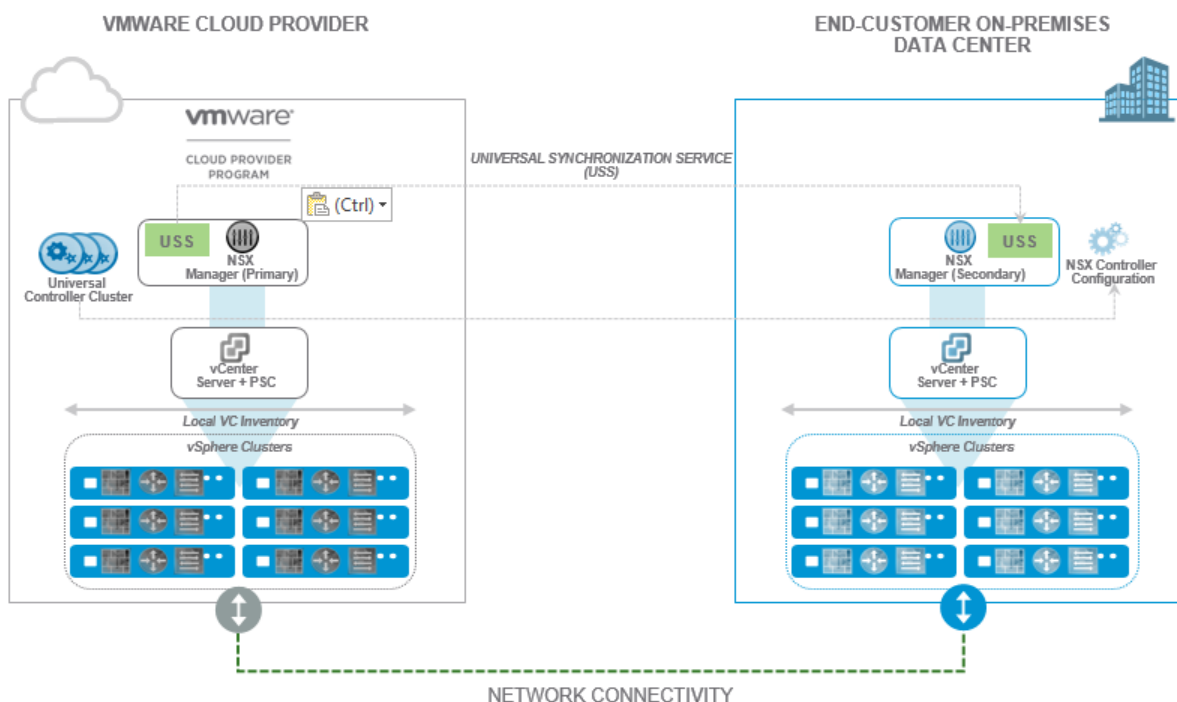
4.4.4 Control-Plane Design

The VMware NSX control plane runs in the NSX Controller cluster. More information, see [Control Plane](#).

The NSX Controller nodes for this solution reside within the VMware Cloud Provider's data center, and the design state of the primary NSX Manager is replicated to the secondary NSX Manager, which contains the NSX Controller configuration.



Figure 9. NSX Controller Design



4.4.4.1 NSX Controller Resource Requirements

NSX Controller nodes are deployed only at the provider side of the solution. To maintain a majority within the cluster, deploy three controllers.

Table 8. vSphere Replication Resource Requirements

Resources	Virtual CPUs	Virtual RAM	Virtual Storage
NSX Controller nodes	4	4 GB	20 GB

4.4.5 Data-Plane Design

The VMware NSX data plane consists of the VMware NSX Virtual Switch™, which is based on the VMware vSphere Distributed Switch™ (VDS) with additional components to enable services. VMware NSX kernel modules, user space agents, configuration files, and install scripts are packaged in VIBs and run within the hypervisor kernel to provide services, such as distributed routing and logical firewall, and to enable VXLAN bridging capabilities.

4.4.5.1 Universal Transport Zone Design

The universal transport zone is configured from the primary NSX Manager and is configured to span the clusters across both the provider and end customer vCenter Server nodes. For more information about adding a transport zone, see [Transport Zones](#).

Note You can create only one universal transport zone within a cross-vCenter NSX implementation.



4.4.5.2 Universal Logical Distributed Router Design

The Universal Logical Distributed Router can be deployed from the primary NSX Manager and it offers centralized administration and a routing configuration that can be customized at the universal logical router, cluster, or host level.

When you create a universal logical router, you must choose whether to enable local egress, because this cannot be changed after creation. Local egress allows you to control what routes are provided to ESXi hosts based on an identifier, the locale ID.

Each NSX Manager is assigned a locale ID, which is set to the NSX Manager UUID by default. You can override the locale ID at the following levels:

- Universal logical router
- Cluster
- ESXi host

If you do not enable local egress, the locale ID is ignored and all ESXi hosts connected to the universal logical router will receive the same routes. Whether or not to enable local egress in a cross-vCenter NSX environment is a design consideration, but enabling local egress is not required for all cross-vCenter NSX configurations.

For this design, local egress is enabled for deploying the universal logical router from the primary NSX Manager so that only routes that are destined for the local site ESXi hosts will be received.

4.4.5.3 Universal Logical Switch Design

The universal logical switches within this architecture are deployed from the primary NSX Manager, and these logical switches are where the workloads will be placed.

Within this design, each logical switch will be configured with a unique subnet to avoid any potential IP address conflicts, because they are connected upstream to a common routing layer.

Note If overlapping IP addresses are required, this setup can be achieved with VMware NSX isolation and segmentation, but that is out of scope for this document.

4.5 Security Design

With VMware NSX, there are two points where the end customer and provider can manage security services:

- Edge services gateways
- Distributed firewall

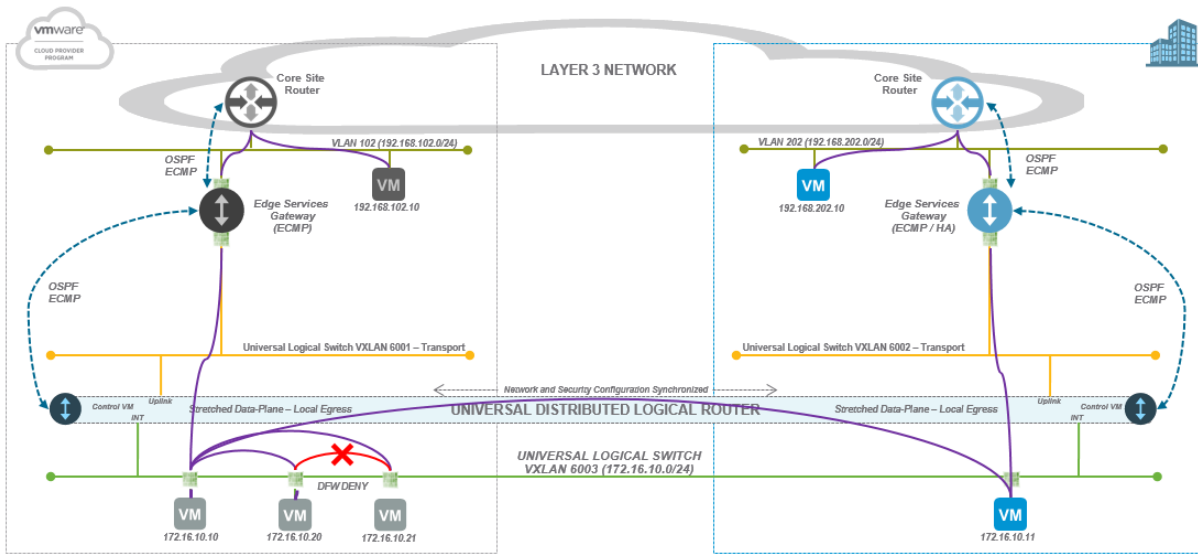
The edge services gateways at each location can be used to control North-South bound firewalling to the physical world. For example, access to networks northbound of the edge services gateway can be controlled through the edge firewall.

The distributed firewall allows definition of security rules that can control access between virtual machines on the same logical network, or within the same universal logical switch. This is commonly known as micro-segmentation. For cross-vCenter security, you can only leverage MAC/IP sets for control, but you could create some dynamic group membership based on server name which would allow for policy to follow the virtual machine.

The following figure highlights how both security solutions can be used in a typical environment.



Figure 10. Security Design Example



For more information about security design, see the [VMware NSX for vSphere Documentation Center](#).

4.6 Workload Mobility Design

One key feature of this solution is the ability to migrate an active state of a virtual machine from one location to another—in this case between the end customer's on-premises data center and the VMware Cloud Provider's data center.

4.6.1 Long-Distance vSphere vMotion

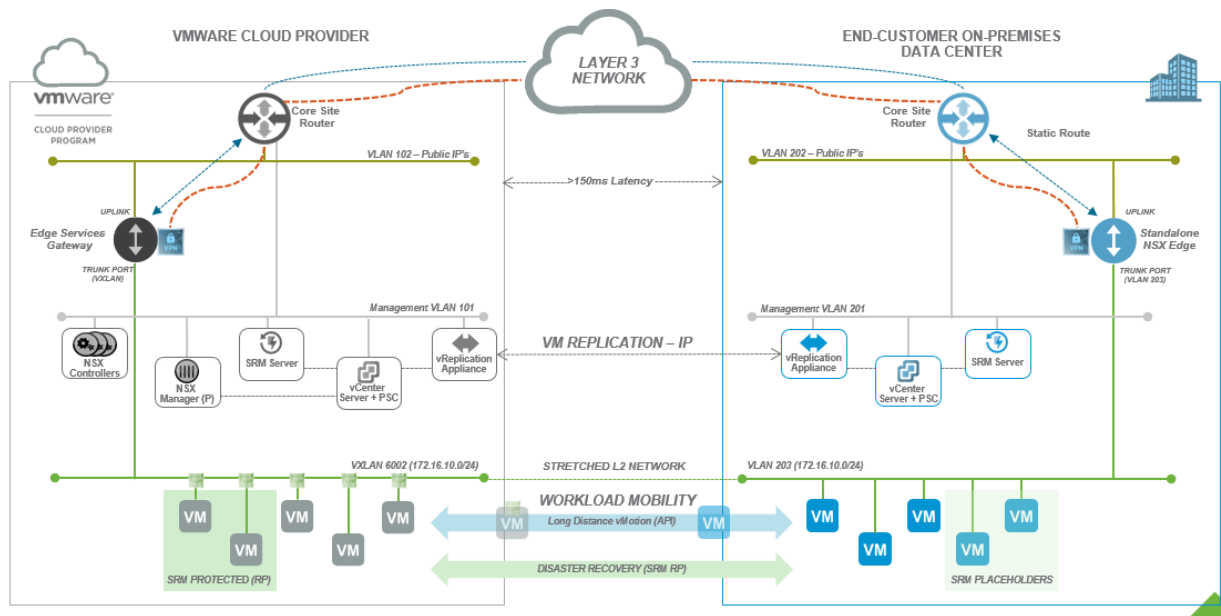
The enabling technology to provide seamless workload mobility to running virtual machines is long distance vSphere vMotion. This capability became available in vSphere 6.0 and enables virtual machines to be migrated across vCenter Server nodes, virtual switches, and physical data centers. The key requirements for long-distance vSphere vMotion are described in the Knowledge Base article [Long-distance vMotion requirements in vSphere 6.0](#).

But the key requirements are:

- Less than 150 ms latency (RTT) between data centers
- At least 250 Mbps bandwidth available for each long-distance vSphere vMotion operation

So the long-distance vSphere vMotion will not be available to everyone based on physics and geographical dispersion.

Note Long-distance vSphere vMotion operations can be executed either through the web client UI or through the API. Long-distance vSphere vMotion operations executed through the vSphere API present an option to authenticate against separate SSO domains.


Figure 11. Workload Mobility Design Overview


4.7 Disaster Recovery Design

This section describes an example disaster recovery design, adding functionality to the existing data center design. Within any disaster recovery design, external factors can drive potential Recovery Point Objective (RPO) and Recovery Time Objectives (RTOs) that the customers can achieve. This includes the following key considerations:

- Latency between the end customer and provider data centers.
- Bandwidth available between the end customer and provider data centers.
- Data change rate of protected virtual machines.
- Complexity of protection groups and recovery plan configuration in Site Recovery Manager.

It is important to understand this before defining a service level agreement for RPO and RTO.

4.7.1 Virtual Machine Replication

vSphere Replication will be used to replicate virtual machines from the end customer's data center to the VMware Cloud Provider's data center. With vSphere Replication, there are a number of considerations when defining the replication schedule.

4.7.1.1 Calculate Bandwidth for vSphere Replication

The text in this section is taken from the [VMware vSphere Replication Administration](#) guide.

“To determine the bandwidth that vSphere Replication requires to replicate virtual machines efficiently, you calculate the average data change rate within an RPO period divided by the link speed.

If you have groups of virtual machines that have different RPO periods, you can determine the replication time for each group of virtual machines. For example, you might have four groups with RPO of 15 minutes, one hour, four hours, and 24 hours. Factor in all the different RPOs in the environment, the subset of virtual machines in your environment that is replicated, the change rate of the data within that subset, the amount of data changes within each configured RPO, and the link speeds in your network.



Prerequisites:

Examine how data change rate, traffic rates, and the link speed meet the RPO. Then look at the aggregate of each group.

Procedure:

1. Identify the average data change rate within the RPO by calculating the average change rate over a longer period then dividing it by the RPO.
2. Calculate how much traffic this data change rate generates in each RPO period.
3. Measure the traffic against your link speed.

For example, a data change rate of 100 GB requires approximately 200 hours to replicate on a T1 network, 30 hours to replicate on a 10-Mbps network, 3 hours on a 100 Mbps network.”

4.7.1.2 Defining an Appropriate RPO

The Recovery Point Objective set by the business has a direct impact on the bandwidth available for replication traffic. For example, if you have an RPO set to one hour, and have an hourly data change rate of 20 GB across your protected virtual machines, that means that vSphere Replication will need to be able to transmit 20 GB of changed data within each one-hour period.

A connection between the end customer and the VMware Cloud Provider of 100 Mbps will allow for transmission of 100 GB each hour, assuming the link is dedicated for replication traffic. If the network is leveraged for additional services, you must take in to account the amount of saturation of the link for other services.

With this RPO example, we could have 80 percent saturation for other services and still meet the RPO of transmitting 20 GB within the hour period.

The following table highlights how the data change rate and the network link speed have a direct impact on the potential RPO that could be offered to end customers.

Table 9. Data Change Rate Compared with Link Speed and Potential RPO

Hourly Data Change Rate	Link Speed (No saturation)	Potential RPO
5 GB	100 Mbps	15 mins (7 mins)*
10 GB	100 Mbps	15 mins (13 mins)
20 GB	100 Mbps	30 mins (25.5 mins)
40 GB	100 Mbps	1 hour (50.1 mins)
80 GB	100 Mbps	2 hours (101.8 mins)

Note These calculations round up and do not take in to account any saturation on the links. Take this in to consideration when making your calculations.

*With vSAN as the storage layer, it is possible to configure an RPO of 5 minutes where applicable.

It is important to set the correct RPO because violations will create a snowball effect where the RPO falls further and further behind the schedule and this will have an impact on the ability to perform recovery.



4.7.1.3 Multiple Point-In-Time Instances

When configuring replication of a virtual machine, you can configure multiple point-in-time (PIT) instances, which are basically a series of timestamped snapshots at the target site that you can leverage post-failover. vSphere Replication supports up to 24 multiple PIT instances per virtual machine. See [Replicating a Virtual Machine and Enabling Point-in-Time Instances](#) for more information.

The benefit of this is that the end customer can revert to a known state of their virtual workloads after a disaster recovery scenario.

See also the [VMware vSphere Replication Administration](#) guide.

4.7.2 Site Recovery Manager Inventory Mappings

Inventory mappings provide default objects in the inventory on the recovery site for the recovered virtual machines to use when you run recovery.

For array-based protection and vSphere Replication protection, if you configure site-wide inventory mappings before you create protection groups, you do not have to configure protection individually on each virtual machine when you create a protection group. Site Recovery Manager applies the site-wide mappings to all virtual machines in an array-based replication protection group or a vSphere Replication protection group at the moment that you create the protection group.

When you use storage policy protection, Site Recovery Manager applies inventory mappings at the moment that a recovery plan runs. You cannot configure protection individually on the virtual machines in a storage policy protection group. As a consequence, you must configure site-wide inventory mappings if you use storage policy protection.

When creating a vCenter Server pair, the administrator must map resources, such as virtual port groups (logical switches), and so on. The following table lists the typical resources that must be mapped for Site Recovery Manager.

Table 10. Site Recovery Manager Inventory Mapping Types

Inventory Category	Description
Network mapping	Map networks on the protected site to networks on the recovery site.
Folder mapping	Map data centers or virtual machine folders on the protected site to data centers or virtual machine folders on the recovery site.
Resource Mapping	Map resource pools, standalone hosts, vApps, or clusters on the protected site to resource pools, standalone hosts, vApps, or clusters on the recovery site. You can map any type of resource on one site to any type of resource on the other site. Note: You cannot map individual hosts that are part of clusters to other resource objects.

4.7.2.1 Automatic or Manual Mapping

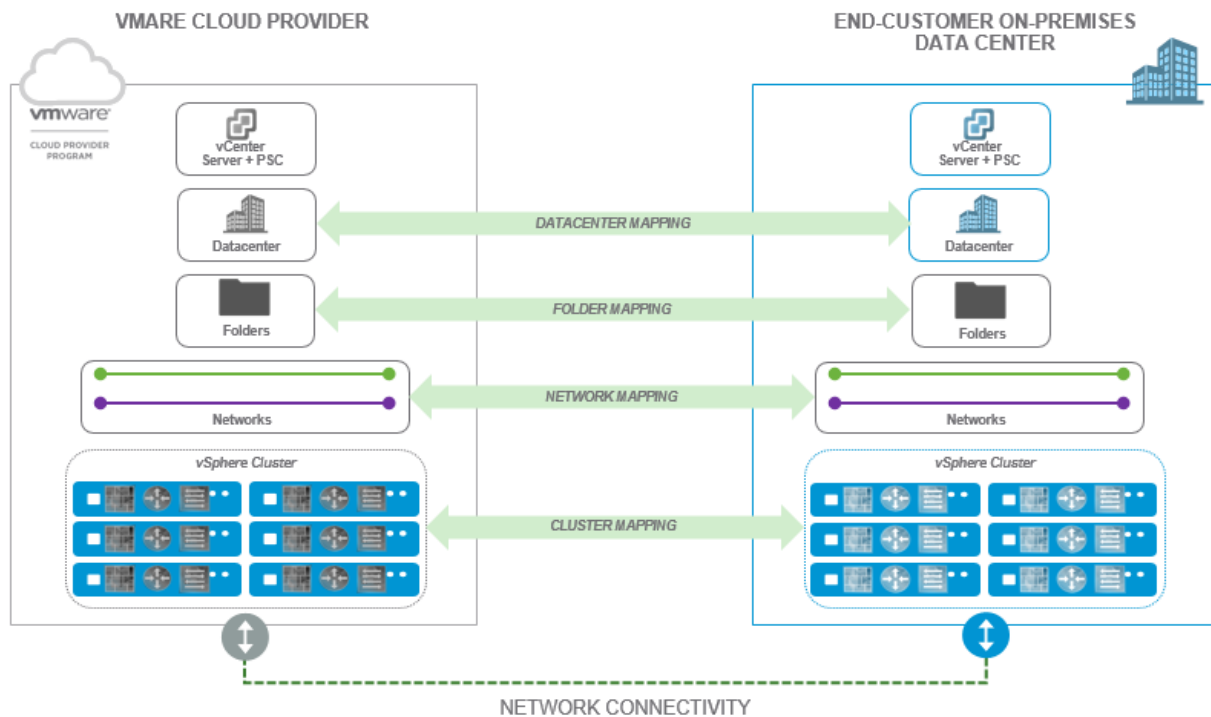
Site Recovery Manager automatically maps networks and folders on the protected site to networks and folders on the recovery site that have the same name. This is quite effective where universal objects are defined, such as universal logical switches.

Manual mapping requires manual definition of mappings between objects.



The following figure shows an example inventory mapping between end customer and VMware Cloud Provider data center objects.

Figure 12. Example Inventory Mapping



4.7.3 Site Recovery Manager Protection Groups

You can include virtual machines that you configured for vSphere Replication in vSphere Replication protection groups. Virtual machines in the vCenter Server inventory that are configured for vSphere Replication are available for selection when you create or edit a vSphere Replication protection group.

You select a target location on a data store on the remote site when you configure vSphere Replication on a virtual machine. When you include a virtual machine with vSphere Replication in a protection group, Site Recovery Manager creates a placeholder virtual machine for recovery. It is possible for the replication target for vSphere Replication and the placeholder virtual machine that Site Recovery Manager creates to both be on the same data store on the recovery site because they are created in different data store folders. When the replication target and the placeholder virtual machines are in the same data store, Site Recovery Manager creates the placeholder virtual machine name by using the replication target name with the suffix (1). To avoid confusion, the best practice is to use different data stores for the vSphere Replication target and for the Site Recovery Manager placeholder virtual machines. Site Recovery Manager applies the inventory mappings to the placeholder virtual machine on the recovery site.

vSphere Replication synchronizes the disk files of the replication target virtual machine according to the Recovery Point Objective that you set when you configured vSphere Replication on the virtual machine. When you perform a recovery with Site Recovery Manager, Site Recovery Manager powers on the replication target virtual machine and registers it with vCenter Server on the recovery site in the place of the placeholder virtual machine.

When using vSphere Replication protection groups, Site Recovery Manager is dependent on vSphere Replication, but vSphere Replication is not dependent on Site Recovery Manager. You can use vSphere Replication independently of Site Recovery Manager. For example, you can use vSphere Replication to



replicate all of the virtual machines in the vCenter Server inventory, but only include a subset of those virtual machines in protection groups. Changes that you make to vSphere Replication configuration can affect the Site Recovery Manager protection of the virtual machines that you do include in protection groups.

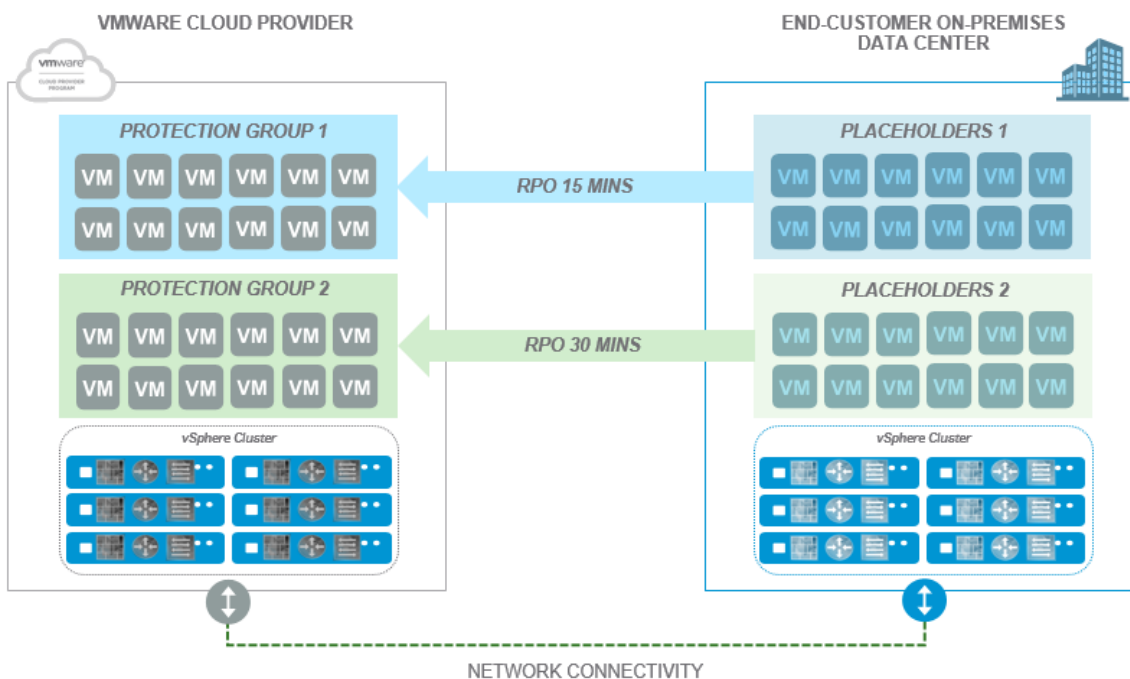
Site Recovery Manager monitors the vSphere Replication status of the virtual machines in vSphere Replication protection groups. If replication is not functioning for a virtual machine in a protection group, Site Recovery Manager cannot recover the virtual machine.

If you unconfigure vSphere Replication on a virtual machine, Site Recovery Manager continues to include that virtual machine in protection groups in which you included it. Site Recovery Manager cannot recover that virtual machine until you reconfigure replication. If you unconfigure vSphere Replication on a virtual machine, you can remove it from the protection group manually.

If you configured vSphere Replication on a virtual machine that resides on a data store that Site Recovery Manager already protects with array-based replication, Site Recovery Manager reports an error if you try to include that virtual machine in a vSphere Replication protection group.

If you remove a virtual machine with vSphere Replication from a protection group, vSphere Replication continues to replicate the virtual machine to the recovery site. The virtual machine does not recover with the rest of the virtual machines in the protection group if you run an associated recovery plan.

Figure 13. Example Protection Groups



Note As the image illustrates, different protection groups can be configured with different RPO policies, allowing the provider to differentiate their service offerings into tiers.

4.7.4 Site Recovery Manager Recovery Plans

Site Recovery Manager recovery plans are the runbooks that are executed by the end customer or cloud operator in the event of a disaster. The definition of the recovery plans must be made with the failure scenarios in mind. For example, the user might require a recovery plan to recover a certain business critical application that is contained within a single protection group. This would be a recovery plan to



recover those workloads only. The user can also create application dependencies within the recovery plan, such as the fact that the database server is recovered before the application and web servers. This is done by leveraging the recovery priority options:

- All priority 1 virtual machines are started before priority 2 virtual machines.
- All priority 2 virtual machines are started before priority 3 virtual machines.
- All priority 3 virtual machines are started before priority 4 virtual machines.
- All priority 4 virtual machines are started before priority 5 virtual machines.
- Startup order of virtual machines within each priority group can also be specified.

The user might also want to create a complete DR recovery plan where all workloads are recovered. This would be a recovery plan that would contain multiple protection groups.

For more information on definition of recovery plans, see the [VMware Site Recovery Manager Documentation Center](#).

4.7.5 API Consumption

Site Recovery Manager has a comprehensive API. For information on how to leverage the API for configuration and execution, see [VMware Site Recovery Manager API Documentation](#).

4.7.6 Site Recovery Manager Operational Limits

Table 11. Site Recovery Manager Operational Limits

Operation Limits	Site Recovery Manager 6.x
Total protected virtual machines	5,000*
Simultaneously recoverable VMs	2,000
Protected virtual machines in a single protection group	500
Protection groups	250
Simultaneous running recovery plans	10
vSphere replicated virtual machines	2,000*

Note Because this solution leverages vSphere Replication, the limit of 2,000 vSphere Replication machines is a limiting factor ahead of the total protected virtual machines.

4.7.7 Site Recovery Manager Planned Migration

One of the key capabilities of this solution is the ability to perform a planned migration to the VMware Cloud Provider. This allows the end customer to leverage the replication schedule and RPOs associated with this solution to move virtual workloads in a controlled manner with low impact and downtime of the workloads.



Planned migration differs from disaster recovery because it assumes the source site is still active and available, and then the process cleanly shuts down the source virtual machines in the correctly defined order, and brings the virtual machines online in the target VMware Cloud Provider Program data center.

4.7.8 Site Recovery Manager Non-Disruptive Disaster Recovery Testing

The ability to perform non-disruptive testing can be a key requirement when defining any BCDR service offering. Site Recovery Manager enables this functionality by allowing the failover to happen without affecting the source virtual machines. The VMs are brought online in an isolated test network, which can be defined by Site Recovery Manager or can be created and mapped in advance.

Note This is a great opportunity for micro-segmentation with VMware NSX 6.2 to create an isolated VXLAN-based network with limited access to services to avoid any potential clashes on the network.

4.7.9 Site Recovery Manager Disaster Recovery

Disaster recovery with Site Recovery Manager assumes that the source data center is no longer available. DR is executed from the target data center and breaks the replication, and forces the target site to bring up the virtual machines in the defined order (set by the recovery plan).

The recovery plan can be configured with additional scripts and external interfaces so that all the required environment variables are correct when the virtual machines are brought up in the target site.

4.7.10 Re-protection

After the virtual machines are failed over to the target data center, the end customer might want to re-protect the workloads back to their on-premises data center as a DR target.

This can be achieved with vSphere Replication and Site Recovery Manager by simply reversing the process for protection with the end customer on-premises data center as the new target.

4.8 Failure Scenarios

This section examines the failure scenarios and the impact to the management solution and the end customer workloads.

4.8.1 Complete Site Failure

In the event that a whole site fails, the recovery steps are as follows:

1. Virtual machines are recovered leveraging Site Recovery Manager.
2. The secondary NSX Manager is promoted to the primary manager.
3. NSX Controller nodes are redeployed in the recovery site.
4. Normal operations resume.

4.8.2 Application Component Failure

If an application component fails and must be recovered to the target site, the procedure is as follows:

1. Virtual machine is recovered by Site Recovery Manager or failed over using long-distance vSphere vMotion.
2. Virtual machine egresses the target data center (if local egress option is enabled).
3. Layer 2 adjacent virtual machines maintain connectivity through VXLAN across the WAN connection.



4.8.3 Edge Cluster Failure

In the unlikely event that the edge cluster fails in one data center, North-South traffic no longer flows and potentially the NSX Controller nodes are offline, which will interrupt the ability to create new networking components within the VMware NSX solution.

The recovery from this scenario is as follows:

1. Instigate DR process to failover workloads from failed site.
2. Promote secondary NSX Manager as primary.
3. Redeploy NSX Controller nodes in recovery site.
4. Redeploy NSX Edge services gateways in the recovery site.
5. Normal operations resume.

An alternative approach is to change the locale ID to the target site, which will send virtual machine traffic across the WAN and egress the target data center until the edge cluster is recovered.

Note If you do not have local egress enabled, you must change the locale ID so that the traffic egresses the correct data center during a failure scenario.



Operational Considerations

5.1 Cloud Service Provider Operations

Because this document is focused on provider managed based solutions, it is assumed that the VMware Cloud Provider will manage both sides of the solution—end customer data center and the VMware Cloud Provider Program based data center. Therefore, operational support is part of the cloud provider's existing operational support systems (OSS). This section describes where VMware software and solutions aid operational support by providing holistic management of the solution.

5.1.1 Platform Monitoring and Alerting

Monitoring of the platform and alerting in the event of a component failure can be achieved through existing tools. However, if there are no existing solutions in place, VMware recommends that the provider bring this solution under the management of the VMware vRealize® Operations™ ecosystem which gives the provider the following key functions:

- Centralized monitoring and alerting (performance and fault data)
 - Compute
 - Networking (VMware NSX MP)
 - Storage (Storage MP)
 - Applications
- Centralized capacity management and planning
- Root cause analysis of identified faults
- Log management and analysis to support root cause
- Operational reporting capabilities

5.1.2 Platform Billing Integration

The VMware Cloud Provider has a requirement to integrate the solution with their existing billing systems or to report on platform usage effectively to generate billing data to send to their end customers.

The billing integration of this platform can be done in many ways ranging from the following:

- VMware vRealize Business™ – vRealize Business offers the VMware Cloud Provider the ability to meter cloud platform usage for end customers by way of the vSphere platform. This enables the provider to generate billing data that can be entered in an existing billing system or to generate bills from the data based on cloud costing models.
- vCenter Server data – the VMware Cloud Provider can also query the vCenter Server platforms to get all the relevant usage data to populate their billing systems which will generate the appropriate bills for their end customers.



5.1.3 VMware Cloud Provider Program Usage Metering

VMware vCloud Usage Meter is an appliance that is installed on VMware vSphere Client™. It meters and collects consumption data on products that are part of the VMware Cloud Provider Program (formerly called VMware Service Provider Program or VSPP) bundles and also generates reports.

Service providers send the monthly data reports to their VMware Cloud Provider Program aggregator for diagnostic purposes.

The VMware Cloud Provider will need to install and configure the usage meter for this solution. The provider may already have a centralized usage meter implementation that they want to leverage for this solution.

The current version of VMware Cloud Provider Program usage meter supports up to:

- 20 vCenter Server instances
- 25,000 virtual machines

Usage meter appliances can be scaled out accordingly based on workload scalability limits and demands.

5.1.4 Service Provider Roles and Responsibilities

Table 12. Provider Roles and Responsibilities

Role	Product	Responsibilities
Site Recovery Administrator	Site Recovery Manager vSphere	Responsible for configuration of Site Recovery Manager and definition of protection groups, recovery plans, and inventory mappings.
vSphere Administrator	vSphere	All vSphere configurations required by the VMware Cloud Provider.
vSphere Replication Administrator	vSphere Replication	All vSphere Replication operations required.
VMware NSX Enterprise Administrator	VMware NSX	All VMware NSX operations.

5.2 Cloud Service Tenant Operations

5.2.1 Executing Long-Distance vSphere vMotion Operations

Long-distance vSphere vMotion can be executed through the vSphere Web Client user interface of the vCenter Server instances that are members of the same PSC domain. The process is very similar to the normal vSphere vMotion operations:

1. Right-click the selected virtual machine.
2. Select **Migrate**.
3. Select **Change both compute resources and storage**.
4. Drill down to select the correct target vCenter Server and vSphere cluster.



5. Select storage and disk format.
6. Select the network to attach to in the target site. (If universal objects are configured, this will be the same universal logical switch.)
7. Schedule high-priority or normal vSphere vMotion operation.

The vSphere vMotion operation will then start to migrate the workload across to the provider data center and vCenter Server.

If the vCenter Server instances are not part of the same PSC domain, long-distance vSphere vMotion operations must be executed through the API. See this [example script](#) written by a VMware engineer.

5.2.2 Executing Planned Migration with Site Recovery Manager

Planned migration within Site Recovery Manager enables the operator to fail over protected workloads in a migration style, which cleanly shuts down the source virtual machines and registers the virtual machines on the target site in a clean and consistent fashion.

Planned migration can also leverage vSphere vMotion where possible to reduce the downtime required while migrating the workloads. For more information, see the [Site Recovery Manager Documentation Center](#).

5.2.3 Executing Disaster Recovery Test Scenario

Site Recovery Manager allows the operator to perform test migrations in line with their testing requirements. This effectively brings up the workloads in an isolated bubble in the recovery site. The isolated networks can be preconfigured, or Site Recovery Manager can create an isolated network that will be leveraged.

For more information, see the *Site Recovery Manager Documentation Center*.

5.2.4 Executing Disaster Recovery Scenario

When you run a recovery plan, Site Recovery Manager migrates all virtual machines in the recovery plan to the recovery site. Site Recovery Manager attempts to shut down the corresponding virtual machines on the protected site.

When it cannot shut down the virtual machines cleanly, Site Recovery Manager assumes that the site is unavailable and forces the DR process to recover the workloads on the target recovery site.

For more information, see the *Site Recovery Manager Documentation Center*.

5.2.5 End Customer Roles and Responsibilities

Table 13. End Customer Roles and Responsibilities

Role	Product	Responsibilities
Protection Groups Administrator Recovery Administrator Recovery Plan Administrator Test Administrator	Site Recovery Manager vSphere	Responsible for operation of Site Recovery Manager where appropriate.
vSphere Administrator	vSphere	All vSphere configurations required by the VMware Cloud Provider Program user.



Role	Product	Responsibilities
Virtual Machine Replication User	vSphere Replication	Managed replications.
VMware NSX Administrator	VMware NSX	VMware NSX operations only. For example, installing virtual appliances and configuring port groups.

Note See the product documentation for complete roles and responsibilities and assign the roles that meet your service requirements.



Conclusion

As this document outlines, by including the VMware NSX 6.2 and Site Recovery Manager 6.1 into a VMware Cloud Provider's IaaS powered hosting portfolio, the service provider can offer a unified hybrid platform that enables the provider to become a strategic extension of their end customer's data center. This also enables the provider to layer disaster recovery and business continuity services on top of the solution with ease and reduced complexity to help solve the business challenges of their end customers.

This solution highlights a single deployment model, and might not be suitable for all providers to offer to their end customers. VMware will follow up this document with more deployment models that the service provider can consider to offer workload mobility or disaster recovery services to their end customers.