WHITE PAPER–NOV 2019

Workload Migration Strategies with VMware Cloud Foundation

Table of contents

Introduction 4

Migration Strategy 4

Migration from Existing vSphere to VCF 5

Migration Types 5

Live Migration 5

Warm Migration 5

Cold Migration 6

Migration Direction 6

Cloud Foundation Workload Domains 6

Cloud Foundation Network Topology 7

Migration Tool Options 7

HCX 9

Overview 9

Architecture 9

HCX Cloud Gateway appliance 9

WAN Optimization appliance 10

Layer 2 Concentrator appliance 10

Multi-Site Connectivity 10

Installation Overview 10

Appliance Installation Location 11

Network Connectivity 11

Migrations 12

HCX vMotion 12

Bulk Migration 12

Cold Migration 12

HCX Enterprise 12

Additional migration services are available with HCX Enterprise edition. 12

Replication Assisted vMotion (RAV) 12

OS Assisted Migration (OSAM) 13

Migration time 13

Best Practices 13

HCX Summary 13

vSphere Replication with SRM 14

Architecture 14

Installation 14

How to Use 16

Additional Network Configuration 25

Other Migration Methods 26

Using Third-Party Migration Software 26

Overview 26

Recommendation 27

Appendix: VM Guest–Setting Considerations 27

Guest Settings 27

About the Authors 28

Workload Migration Strategies with VMware Cloud Foundation™

Introduction

We live in a hybrid-cloud world where application workloads require a highly resilient and flexible infrastructure: the kind of infrastructure that enables users to move those workloads freely between data centers, as well as between private and public clouds, with minimal downtime and without having to replatform. Business today demands the flexibility to choose where to run workloads as well as the option of the private or public cloud. VMware believes that the key to realizing this kind of highly resilient and flexible environment is in the adoption of a software-defined approach to IT infrastructure. We call this the Software-Defined Data Center (SDDC). The SDDC is the architecture for the modern hybrid cloud. The VMware SDDC is extensively tested, validated, and automated through VMware Cloud Foundation™ for a turnkey product approach. Cloud Foundation can be both deployed on premises, as part of a private cloud, and hosted by cloud providers, as a public cloud offering.

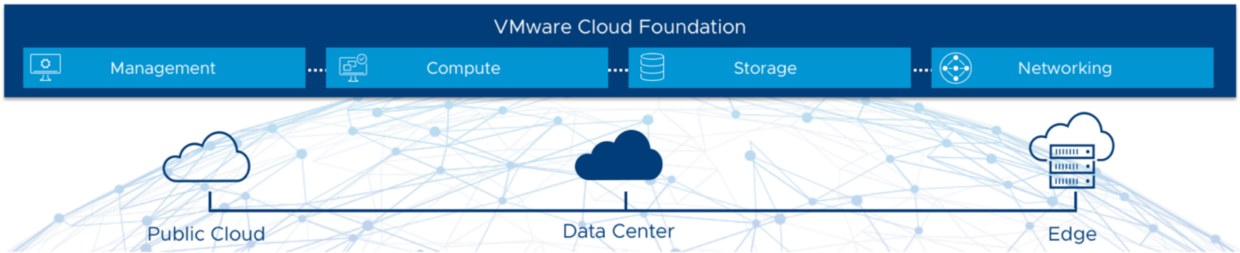

A Cloud Foundation deployment is only the start. Customers have large numbers of existing, mission-critical workloads already running in their data centers. As they work to transition their data centers into private clouds and possibly to look to adopt public clouds, they must be able to migrate these existing workloads off their legacy infrastructure and into the new SDDC. This white paper provides an overview of the available migration options and enables readers to understand the pros and cons of each option to help determine which is best suited to pursue the fastest path from legacy to a modern private cloud and from there to the hybrid cloud .

Scope

The workload migration examples in this document are based on a migrating-to-hybrid cloud operational model using the components of the VMware SDDC. This model enables users to run application workloads in any cloud based on the VMware SDDC, whether that cloud is on premises or with a public cloud service provider. Regardless of which cloud is chosen, this document provides an overview of the migration options available and the technical guidance to complete the installation, configuration, and operational procedures for migrating application workloads into a true hybrid cloud.

Migration Strategy

Business drivers—cost, data security, data locality, compliance, burst capacity, consolidation, mergers and acquisitions, and many others—ultimately impel placement decisions regarding application workloads. There is no “one size fits all” migration approach. Each customer is unique, and VMware gives users the freedom to choose where applications reside. We also understand that where a workload is placed today might not be where it is needed to run tomorrow. VMware aims to provide a highly secure, highly reliable infrastructure that gives users the flexibility to move workloads as needed to best meet business demands. This white paper considers all these factors to help users determine the best tool for a migration strategy.

Migration from Existing vSphere to VCF

VMware Cloud Foundation is a *greenfield* installation only, meaning that Cloud Foundation creates a new vSphere infrastructure and Single Sign-On domain. Existing workloads running on older vSphere versions cannot be managed from Cloud Foundation. For many, the first stage of adopting a hybrid cloud is to migrate current applications running on existing vSphere environments into Cloud Foundation. The second stage in adopting a hybrid cloud then entails enabling workload migration paths between clouds. These clouds can be either private or public clouds. The tools outlined in this whitepaper allow you to migrate your VM workloads to Cloud Foundation and pave the way to true hybrid cloud architecture.

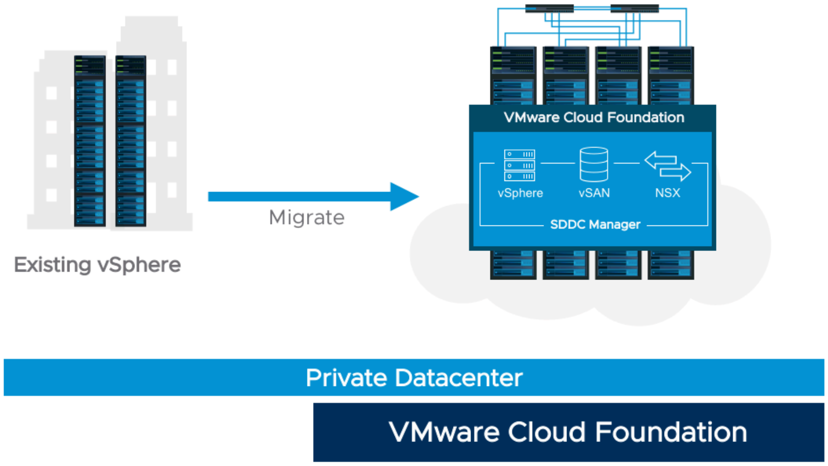

Migration Types

There are three primary methods of migrating application workloads: *live migration*, *warm migration, and cold Migration*. Let's review the meaning and assumptions for each.

Live Migration

Live migration, also referred to as *hot migration* or VMware vSphere® vMotion® migration, is the ability to relocate an application workload with no downtime. This means that virtual machines (VMs) are not powered off and the IP address does not change.

To determine whether hot migration is optimal, first consider network connectivity. The migration target location must be on the same layer 2 (L2) IP subnet. Regarding layer 2 stretch, consider how to best manage the maintenance of these systems as well as all possible downtime scenarios. Review the “Appendix” at the end of this document to learn about many of the possible VM guest settings that can impact all migrations.

Migration Approach	Pros	Cons
Live Migration	• No downtime • No changes to the guest OS or application	• Requires stretching the L2 IP network

Table 1. Pros and Cons of Live Migration

Warm Migration

Warm migration is the most flexible option. With this approach, the VM remains powered on while the data sync's across source and destination. Once the VM data is in sync, there is a brief outage as the VM guest is powered off while it is migrated, then powered back on. Also, because the VM is powered off, it facilitates changing the IP address and subnet during the move. Workloads do not have to be offline for long periods during a warm migration. Very fast warm migrations with little downtime, sometimes less than a few minutes, can be achieved with careful planning and preparation by using the tools discussed in this paper. The following benefits of warm migration

are not available with live migration:

- Warm migration facilitates upgrades of VM compatibility level and VMware Tools™ version.
- Application workloads can take advantage of new capabilities associated with newer virtual hardware.
- Moving to a newer-generation physical server can make new CPU instruction sets available to the operating system (OS) or application.
- NUMA boundaries and configuration can change, based on the new hardware.

See the “Appendix” for details on many of the settings that can be affected by migration.

Migration Approach	Pros	Cons
Warm Migration	• Most flexible • Enables changing IPs as well as upgrading VM compatibility and VMware Tools versions • Enables use of new hardware features	• Requires downtime

Table 2. Pros and Cons of Cold Migration

Cold Migration

Cold migration means the migrated VM is powered off during the entire data migration process. The amount of downtime depends on the size of the VM and how long it takes to migrate with the bandwidth available between sites.

Migration Approach	Pros	Cons
Cold Migration	• Enables changing IPs as well as upgrading VM compatibility and VMware Tools versions • Enables use of new hardware features	• Requires extended downtime

Migration Direction

Migration is not a one-way street in a hybrid-cloud world. Business requirements change. Regulatory requirements change. These changes can require moving workloads migrated to the public cloud back on premises or vice versa. Consider this when designing a migration strategy. This white paper is written with an understanding that migrations occur in both directions. This gives enterprises flexibility and choice with application workloads.

Cloud Foundation Workload Domains

Cloud Foundation has two architecture methodologies, Standard & Consolidated.

The Standard Architecture creates a management workload domain used for infrastructure management appliances. I.E., vCenter, PSCs, NSX manager, etc. Additional Workload Domains can be created for application workloads.

Each Workload Domain has its own vCenter instance connected to the same SSO domain as the Management Domain. Because they share the same SSO domain, they are connected to each with enhanced link mode.

The Consolidated Architecture only creates a single management workload domain. All application workloads are placed within the application resource pool. The consolidated architecture is recommended for small and medium businesses where the management domain is only 4 to 6 hosts. If you are deploying 7 or more hosts, we would recommend using the standard architecture for greater flexibility in future growth.

Cloud Foundation Network Topology

Cloud Foundation leverages the advanced automation capabilities of the VMware SDDC Manager to automate and simplify the deployment of the SDDC. An understanding of the underlying physical and virtual networking architecture implemented by Cloud Foundation is essential to designing a migration strategy and knowing how to route and control network traffic.

Physical network topology for Cloud Foundation since version 3.x is designed and controlled by the customer. Because of this, this whitepaper does not document a specific physical network topology.

Layered on top of the physical network, Cloud Foundation management domain leverages the software-defined networking capabilities provided by VMware vSphere Distributed Switch™ (VDS) together with NSX-V instance. Figure 2 depicts the Cloud Foundation software-defined logical network topology implemented since Cloud Foundation version 3.x.

Using the tools mentioned in this whitepaper, we reference some of the logical networking components in a Cloud Foundation environment. Understanding where to install and connect some migration tools help you better understand how to proceed with migrations in your environment.

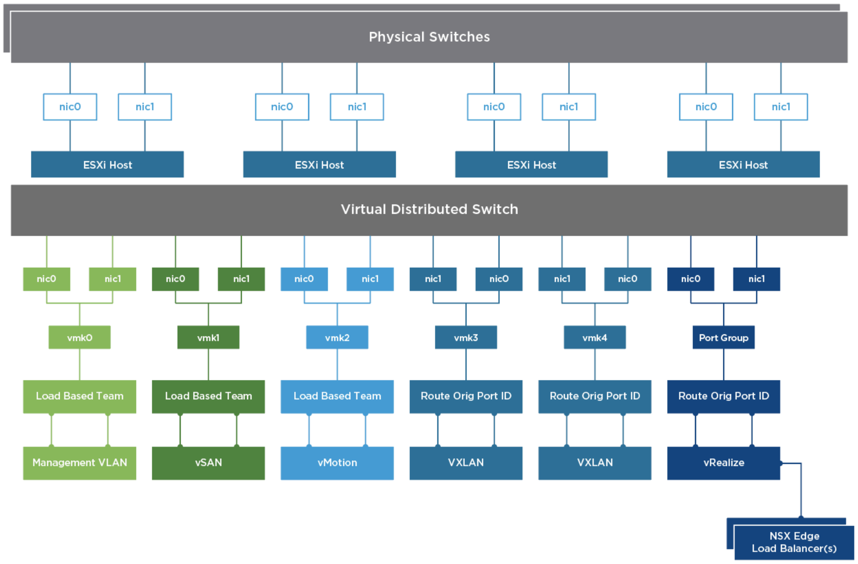

Figure 2. Software-Defined Logical Network Topology Implemented with Cloud Foundation

Migration Tool Options

The migration path can quickly be determined based on the three migration choices—live, warm, and cold—and by the version of any existing vSphere environments.

This white paper reviews two possible tools and methods for migration:

1. VMware HCX
2. VMware vSphere Replication™

Below is a graphical representation of the migration tools in this paper and the vSphere version(s) they support, for live migration. Live migration requires that the L2 VLAN that the VM currently runs on be extended to the Cloud Foundation site.

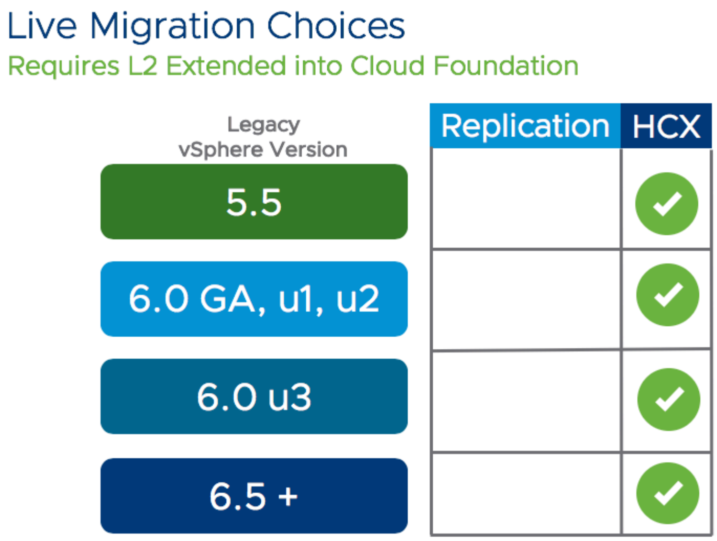

Live Migration Choices

Requires L2 Extended into Cloud Foundation

Legacy vSphere Version	Replication	HCX
5.5		✓
6.0 GA, u1, u2		✓
6.0 u3		✓
6.5 +		✓

Below is a similar chart that again shows the two options. This time, support for cold or low-downtime migration is also shown.

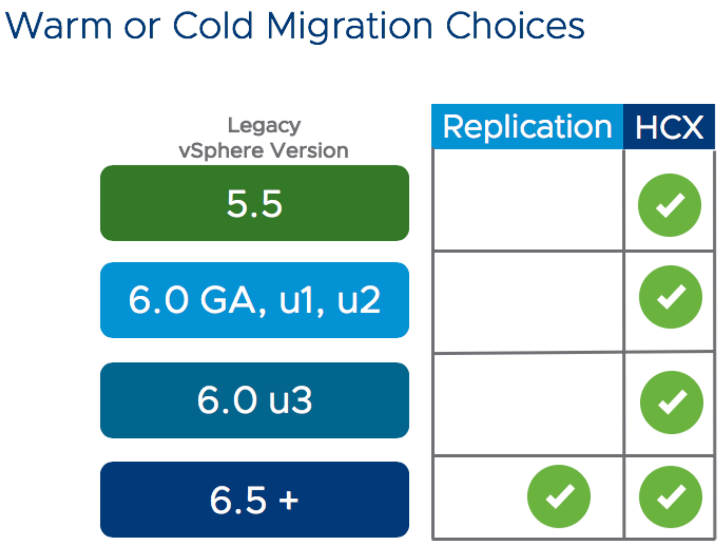

Warm or Cold Migration Choices

Legacy vSphere Version	Replication	HCX
5.5		✓
6.0 GA, u1, u2		✓
6.0 u3		✓
6.5 +	✓	✓

This white paper next discusses the details of the tools, providing an overview, architecture, installation, and migration usage for each.

HCX

Overview

HCX abstracts on-premises and cloud infrastructure resources and presents them to applications as one continuous hybrid cloud. It provides high-performance, secure, and optimized multi-site interconnects. The abstraction and interconnects create a highly secure encrypted hybrid interconnect. With this hybrid interconnect, HCX facilitates secure and seamless application mobility and disaster recovery across multiple on-premises vSphere platforms and VMware public clouds.

HCX is the most feature-rich tool currently available for Cloud Foundation workload mobility. This document explains how to install, configure, and use HCX with a Cloud Foundation site when connecting to a legacy vSphere environment.

Architecture

The HCX migration process begins with two sites, a source, and a target. Each site has unique requirements for software installation. The target site requires the use of NSX; for this document, the target is a Cloud Foundation Workload Domain, as it includes NSX. Each site requires an HCX manager appliance.

The HCX manager appliance is made up of two slightly different manager appliances. One is installed only at the source, and the other is installed only at the target site. Note: although we use the terms *"source"* and "*target,"* migrations with HCX between these sites are bi-directional.

Source site

- Use the HCX *Enterprise* Manager OVA when deploying HCX at the source site

Target Site

- Use the HCX Cloud Manager OVA when deploying HCX at the target site

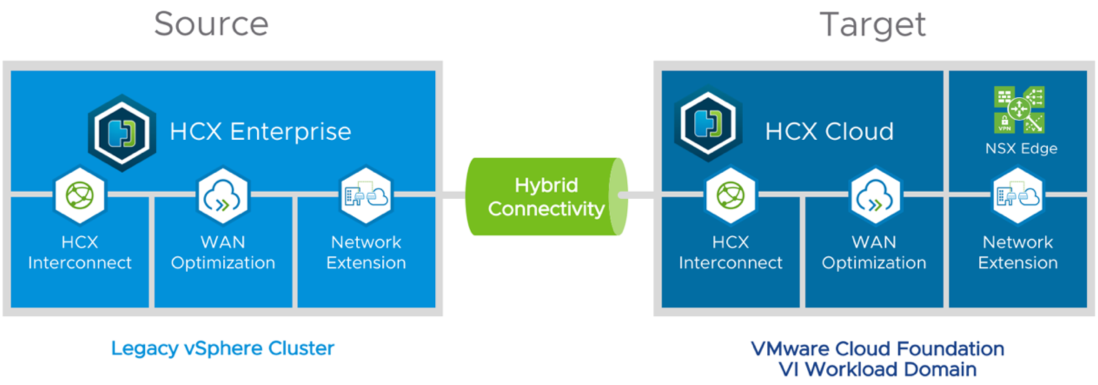

Architecture Overview – HCX

In addition to the HCX manager appliance, there are three service appliances that can be deployed.

HCX Cloud Gateway appliance

This Cloud Gateway appliance is responsible for creating encrypted tunnels across enterprise sites for the vMotion and vSphere-based replication traffic. The appliance makes it easy to connect the source and target sites and reduces the engineering time required to build a hybrid cloud infrastructure that enables workload mobility freedom.

WAN Optimization appliance

This high-performance appliance provides data deduplication and compression. The appliance improves the performance of the tunnel, creating a LAN-like experience. Note: When using HCX to migrate within a datacenter and migrations occur on a LAN, we recommend not deploying the WAN Optimization appliance.

Layer 2 Concentrator appliance

This appliance creates a Layer 2 (L2) network extension between vSphere environments, enabling workloads to maintain the same IP address during migration.

Multi-Site Connectivity

HCX supports multi-site deployment models. Connectivity begins with the HCX Enterprise manager appliance initiating the connection to the HCX Cloud Manager. Each HCX Enterprise manager can create multiple site pair connections to unique HCX Cloud managers. For each site pair, HCX deploys a unique set of service appliances. See the example diagram. Note that each HCX manager can only register with a single vCenter instance.

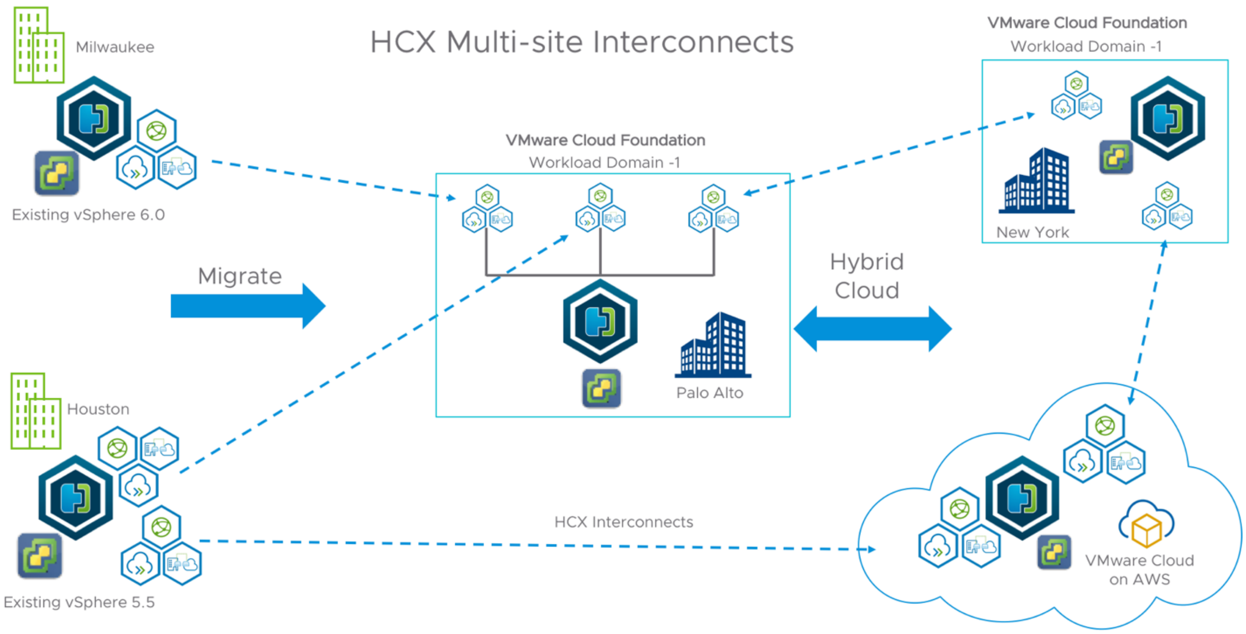

Installation Overview

Install and configure the HCX appliance by following the HCX user guide installation documentation here. *https://docs.vmware.com/en/VMware-HCX/index.html*

Here is a simplified high-level installation for this document.

1. Install and configure the HCX Cloud appliance on Cloud Foundation target site.
2. Create an HCX network profile on destination
3. Retrieve the link for the HCX Enterprise OVA, from the HCX Cloud UI
4. Install and configure the HCX Enterprise OVA on the source site.
5. Create an HCX Site pair.
6. Create an HCX network profile on source site
7. Create an HCX Service Mesh to create a hybridity tunnel between sites.

Appliance Installation Location

When considering where to install the target HCX appliances, install the HCX appliances into the Cloud Foundation Workload Domain where you want to move your application too. You then register the HCX manager with that Workload Domain vCenter. The HCX manager appliance then has access to the vCenter inventory, including the application VMs and networking components.

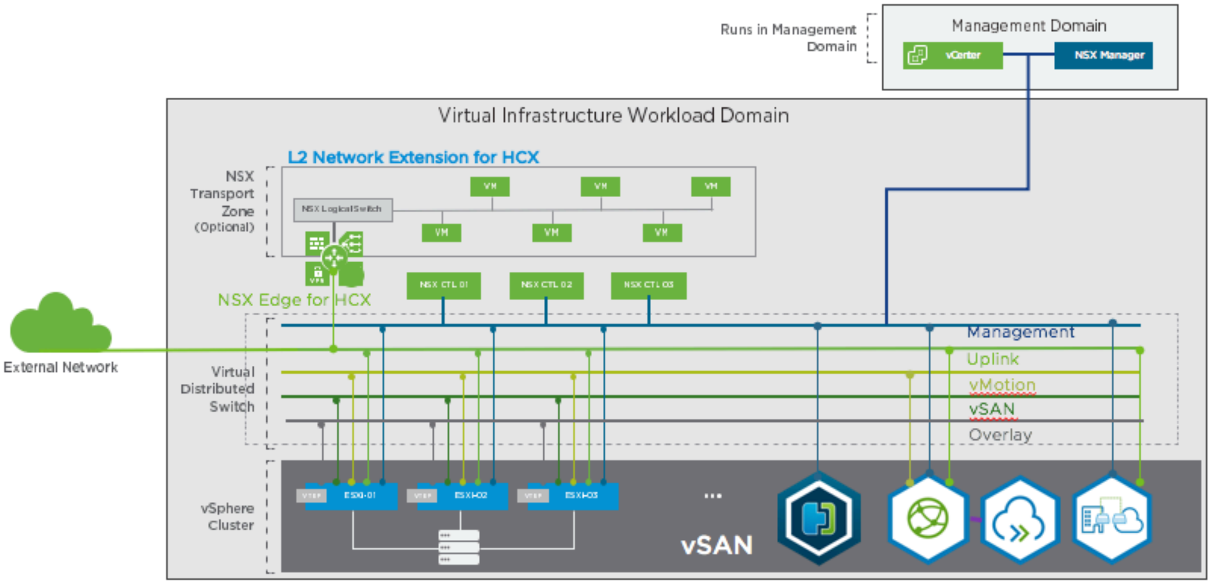

Image of VCF Standard Architecture installation with HCX appliances

If you are deploying a Cloud Foundation consolidated architecture, you would then install the HCX appliances into the Management Workload Domain.

Keep in mind, once you migrate a VM into a Cloud foundation environment, all Workload Domains are connected with enhanced link mode. Enhanced link mode allows you to freely vMotion VMs between all workload domains in a Cloud Foundation environment.

Network Connectivity

When migrating from a legacy vSphere environment that is in the same datacenter as your target Cloud Foundation environment, the HCX Interconnect appliance needs to connect to an uplink network. This uplink network needs to be a subnet that is reachable between source and target sites. For instance, If the management subnet is shared between the two sites, the uplink network can be the management subnet. Doing this keeps the installation very simple for migrating source and target in the same data center. For more advanced network topologies, see the HCX documentation.

Migrations

HCX Advanced has three options for migrations, vMotion (Live), Bulk Migration (warm), and Cold Migration. Each of these migration options has pros/cons, as described above. Let's dive into the technical details on each of these as they pertain to HCX and Cloud Foundation.

HCX vMotion

Using HCX vMotion leverages the VMware vMotion protocol to move a live VM to the target site. During the migration, HCX copies the active state of the VM, including virtual disks, CPU cycles, RAM, IP address, and MAC address. Listed in the HCX documentation are several prerequisites for HCX vMotion.

Some essential perquisites include;

- The VM virtual hardware is version 9 or above.
- The Layer 2 network subnet of the VM you are migrating is extended from source to target.

HCX provides a simple user interface for extending a Layer 2 network. Once you have extended the network, you can proceed to live-migrate one VM at a time. You also have the option to select upgrading the VM tools or VM hardware during HCX vMotion. Note; selecting these options does not automatically update either of these components during the vMotion process. Instead, it creates a pending operation in the VM. These updates are then completed after the next reboot of the VM, as these changes require the VM to be powered off.

Bulk Migration

HCX Bulk Migration uses the host-based replication protocol to copy the disks of the VM to the target site. The VM can remain powered on during this disk replication period. Once the disk replication is complete, a delta sync occurs on the disk(s). After the delta sync is complete, HCX proceeds with the switchover process. During the switchover phase, the source VM is powered off, and the replica is powered on in the target site. The switch over process can run immediately or wait until a scheduled downtime window. The downtime period of a bulk migration is concise. The workflow for this process is; power-off source, then, power-on target. Bulk migration, as the name implies, can occur on multiple VMs at once.

Bulk Migrations do not require Layer 2 extension but are recommended for a faster switchover. Personalization scripts can be run during the switchover phase allow you to change the IP address or make other VM modifications. Additional HCX automated options include; updating the VM tools driver(s), and VM Hardware level. These settings are all selectable from the HCX UI. They are then executed immediately to the target side replica VM during the switchover process.

Cold Migration

HCX has a third method for VM migration, this option is not listed in the UI for migration, but is instead dependent on the power status of the VM you have chosen to migrate. If the VM is powered off, a cold migration process occurs by default. Cold migration uses the NFC protocol to migrate the VM. During a cold migration, the Virtual Machine IP address and MAC address remain preserved. Cold migrations must satisfy the same requirements as vMotion.

HCX Enterprise

Additional migration services are available with HCX Enterprise edition.

Replication Assisted vMotion (RAV)

RAV combines the power of vSphere Replication and vMotion. This advance migration service allows you to live migrate multiple VMs all at once. Select all the VMs you would like to migrate, and HCX begins the process of replicating the VM state to the target site. When it is time to migrate, HCX then switches from Replication to vMotion and moves the active state of the VM.

OS Assisted Migration (OSAM)

OSAM enables you to migrate application VMs that are not running in a vSphere based environment. OSAM is an agent-based migration service that copies the state of a VM and recreates it in a vSphere Environment. The migration process is a warm migration, once the VM copies to the target site. HCX powers off the source, and powers on the target VM.

Migration time

The time required to migrate VMs depends on many factors, including VM size and available bandwidth between sites, disk read/write speeds at source and destination.

Best Practices

- When installing HCX appliances, it is recommended to connect the appliances directly to the Management and vMotion networks for improved performance and reliability
- Keep the network topology simple. Installing the HCX appliances on dedicated subnets is possible, but increases network downtime with firewall and router interdependencies.
- For best migration performance, set resource reservations on the HCX appliances.
- When using HCX to migrate within a datacenter and migrations occur on a LAN, we recommend not deploying the WAN optimization appliance.

HCX Summary

HCX is extremely powerful and easy to use. It enables live migrations, warm, and cold migrations to VMware Cloud Foundation sites without the need to re-engineer the physical network infrastructure. This tool quickly creates a hybrid cloud.

vSphere Replication with SRM

vSphere Replication is a hypervisor-based, asynchronous replication solution for vSphere VMs. It is included at no extra cost with most vSphere versions. Originally designed for data protection and disaster recovery, it can also be used as a low-cost, easy-to-use vSphere migration tool.

Site Recovery Manager is a business continuity and disaster recovery solution that helps you to plan, test, and run the recovery of virtual machines between a protected vCenter Server site and a recovery vCenter Server site.

Architecture

In this white paper, we install the vSphere Replication appliance on a Cloud Foundation instance with a private replication subnet so bandwidth usage can be controlled with VMware vSphere Network I/O Control. Using Network I/O Control, or CoS/DSCP tagging ensures that migrations do not negatively impact production workloads. Cloud Foundation supports configuration of additional VMkernel adapters that are dedicated to vSphere Replication for network traffic isolation. Do not modify the VMkernel adapters created by the SDDC Manager.

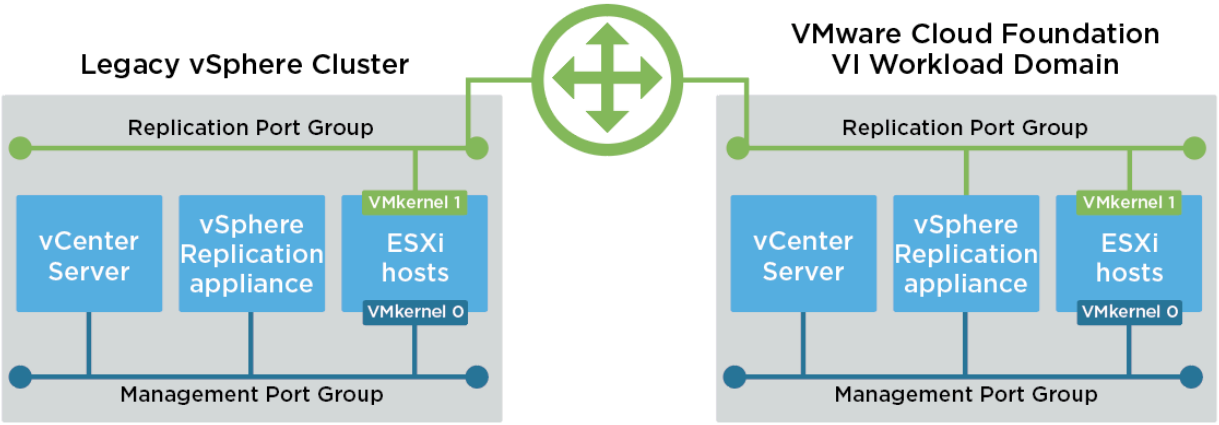

Figure 9. vSphere Replication Architecture Overview

Installation

Cloud Foundation uses an architecture based on workload domains. By default, a single management workload domain is created for all Cloud Foundation management components. In addition, virtual infrastructure (VI) workload domains can be created for application workloads to maintain separation. When installing the vSphere Replication appliance, know in advance where application workloads will be run.

For a Cloud Foundation single-rack consolidated workload domain, the vSphere Replication appliance is installed in the management domain. This white paper explains how to configure it. The same process can be followed for configuring a VI workload domain.

Prior to installing SRM install vSphere Replication appliances at both sites. For installation and configuration of SRM appliances see Deploying the Site Recovery Manager Appliance

With the appliances installed and configured in source and target sites, we now proceed to connect the two.

From the legacy source vCenter Server instance, navigate to ***Menu > Site Recovery > Open Site Recovery***

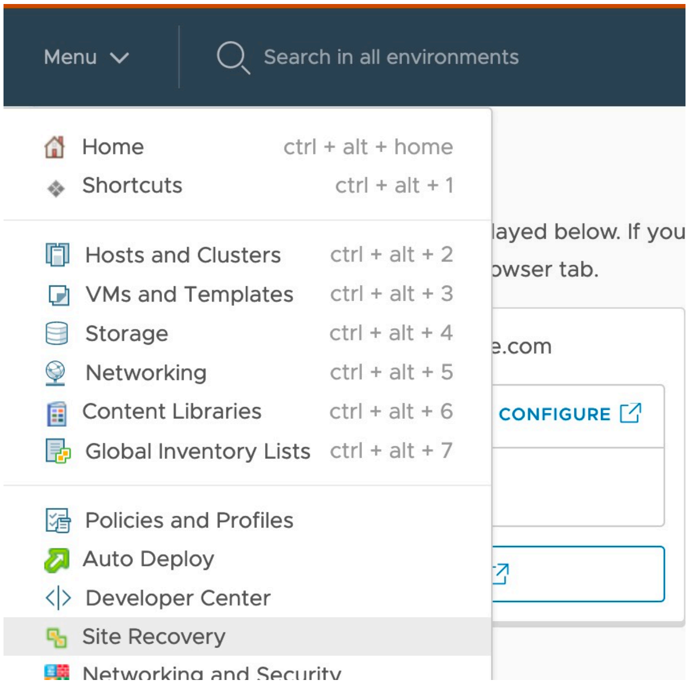

Select Open Site Recovery

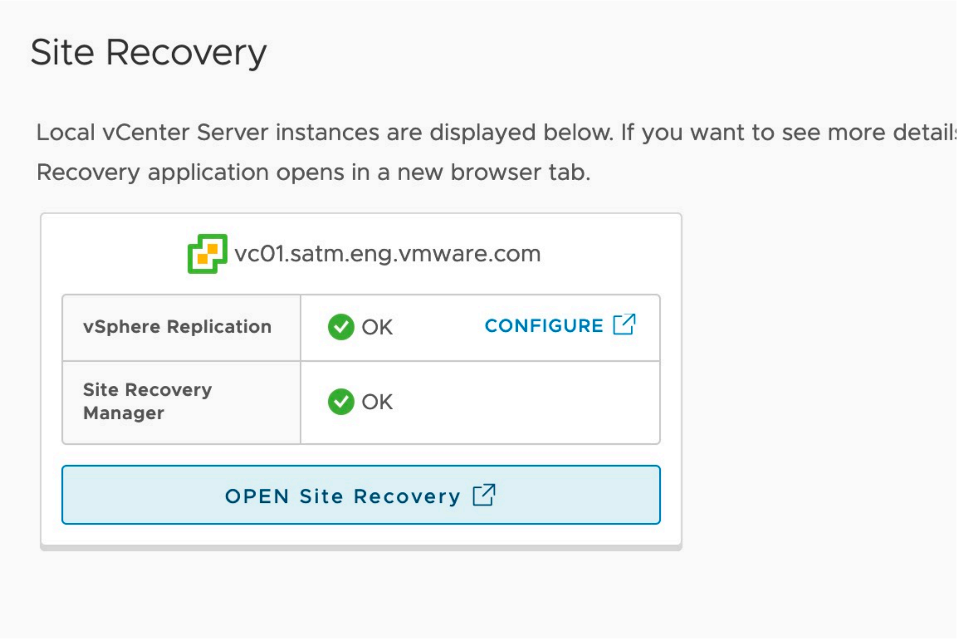

Site Recovery

Local vCenter Server instances are displayed below. If you want to see more detail
Recovery application opens in a new browser tab.

vc01.satm.eng.vmware.com

vSphere Replication	✓ OK	CONFIGURE
Site Recovery Manager	✓ OK	

NOTE: Ensure that the DNS name for each site is resolvable from each site.

How to Use

After the vSphere Replication appliance has been installed, it is quite easy to use for migration.

Right-click the VM to be migrated. Navigate to All *Site Recovery Actions > Configure Replication*.

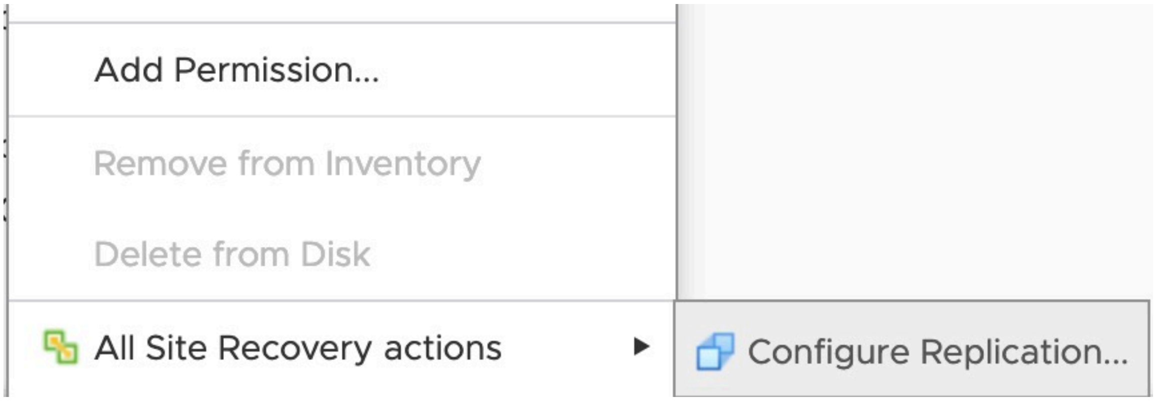

This will redirect you into Site Recovery Manager. Navigate to Replication > Forward Replication > New

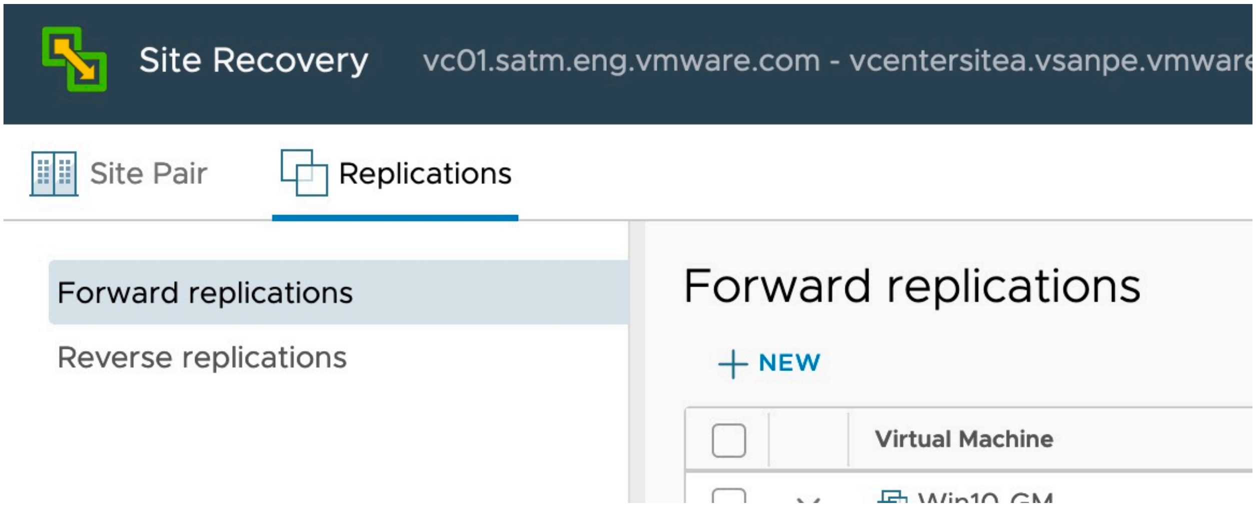

Select the virtual machines to replicate.

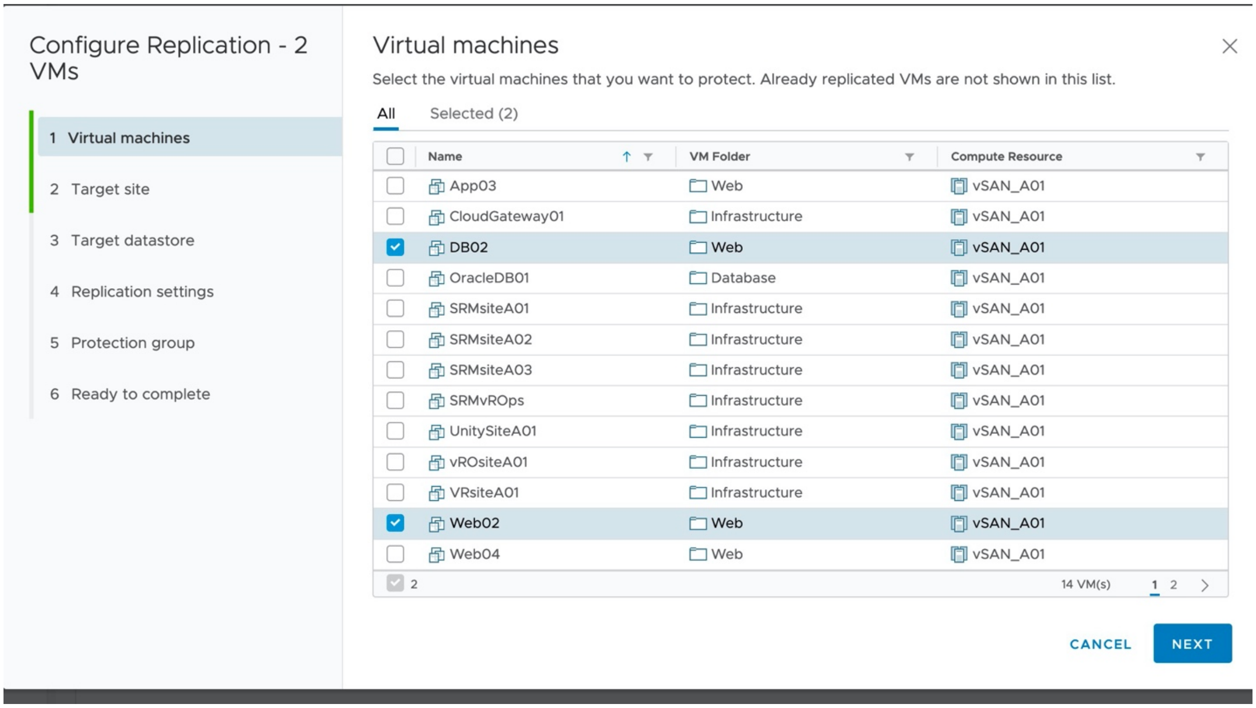

Configure Replication - 2 VMs

1 Virtual machines

2 Target site

3 Target datastore

4 Replication settings

5 Protection group

6 Ready to complete

Virtual machines

Select the virtual machines that you want to protect. Already replicated VMs are not shown in this list.

All | Selected (2)

□	Name	VM Folder	Compute Resource
□	App03	Web	vSAN_A01
□	CloudGateway01	Infrastructure	vSAN_A01
☑	DB02	Web	vSAN_A01
□	OracleDB01	Database	vSAN_A01
□	SRMsiteA01	Infrastructure	vSAN_A01
□	SRMsiteA02	Infrastructure	vSAN_A01
□	SRMsiteA03	Infrastructure	vSAN_A01
□	SRMvROps	Infrastructure	vSAN_A01
□	UnitySiteA01	Infrastructure	vSAN_A01
□	vROsiteA01	Infrastructure	vSAN_A01
□	VRsiteA01	Infrastructure	vSAN_A01
☑	Web02	Web	vSAN_A01
□	Web04	Web	vSAN_A01

☑ 2 — 14 VM(s) 1 2 >

Select the **Target site** to replicate to. Note you can manually select a replication server or have Site Recovery automatically assign a replication server.

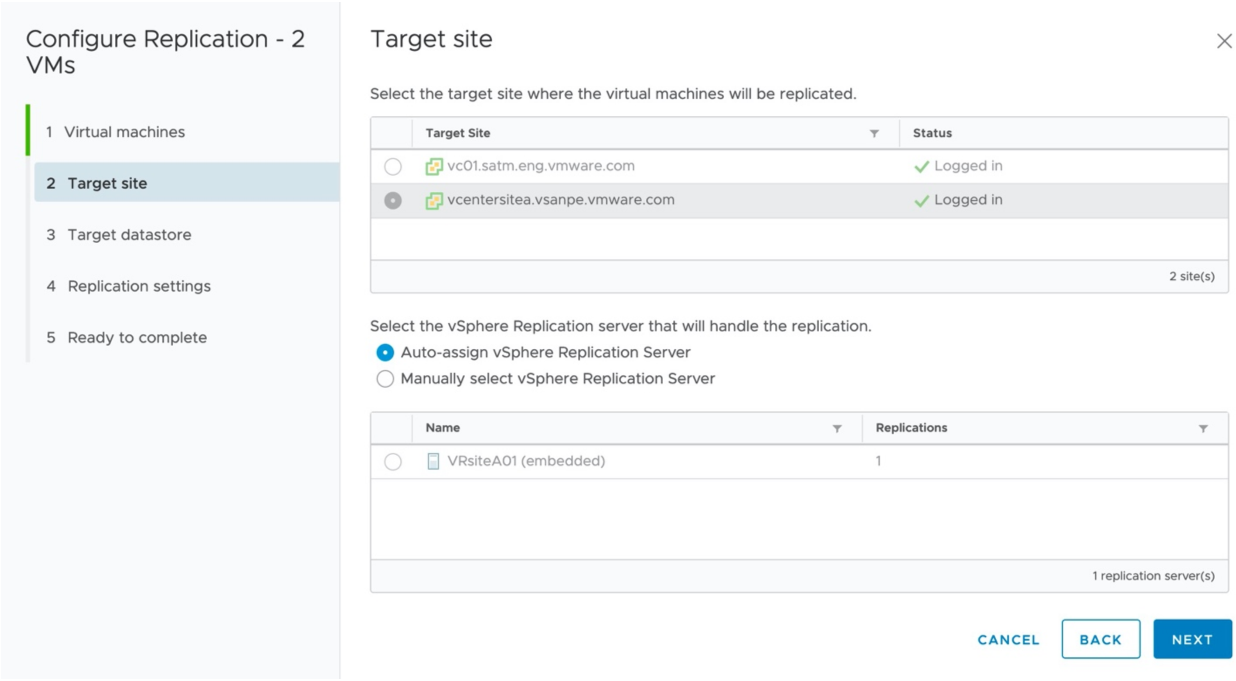

Configure Replication - 2 VMs

1 Virtual machines

2 Target site

3 Target datastore

4 Replication settings

5 Ready to complete

Target site

Select the target site where the virtual machines will be replicated.

	Target Site	Status
○	vc01.satm.eng.vmware.com	✓ Logged in
●	vcentersitea.vsanpe.vmware.com	✓ Logged in

2 site(s)

Select the vSphere Replication server that will handle the replication.

● Auto-assign vSphere Replication Server

○ Manually select vSphere Replication Server

	Name	Replications
○	VRsiteA01 (embedded)	1

1 replication server(s)

Select the **Target location**. For the Cloud Foundation site. In this example we have selected a “Oracle_Database” policy that has highlighted a supported vSAN datastore.

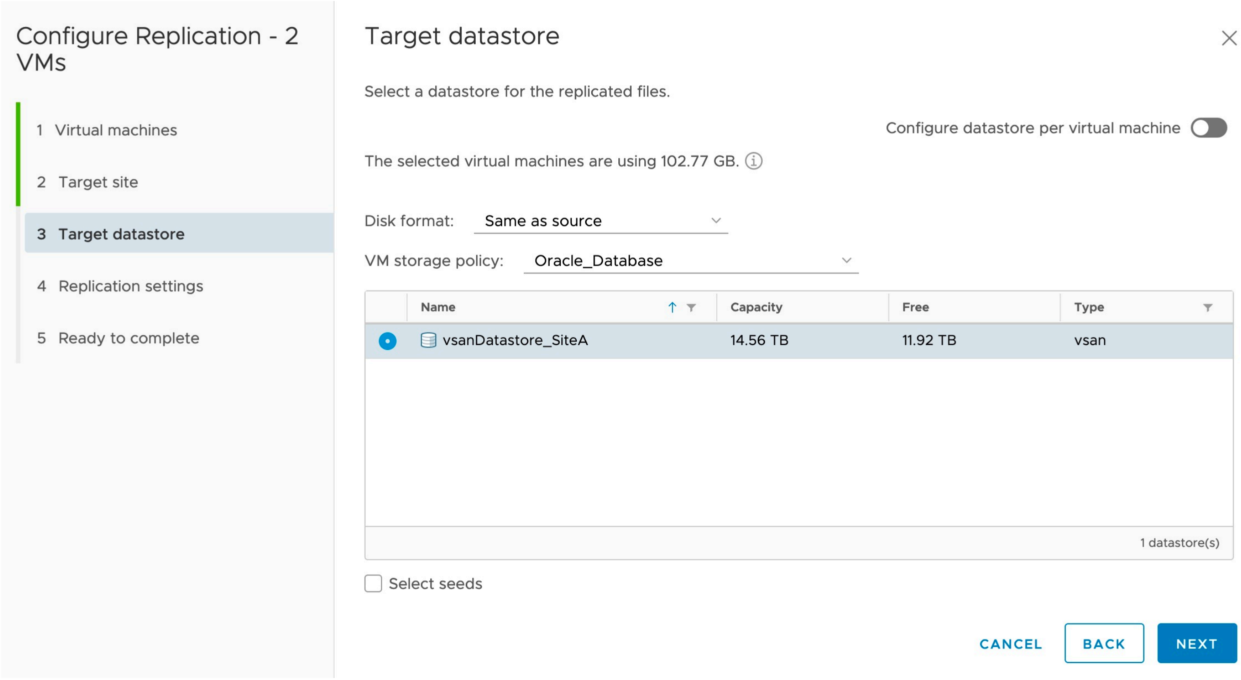

Configure Replication - 2 VMs

1 Virtual machines

2 Target site

3 Target datastore

4 Replication settings

5 Ready to complete

Target datastore

Select a datastore for the replicated files.

Configure datastore per virtual machine

The selected virtual machines are using 102.77 GB.

Disk format: Same as source

VM storage policy: Oracle_Database

	Name	Capacity	Free	Type
●	vsanDatastore_SiteA	14.56 TB	11.92 TB	vsan

1 datastore(s)

□ Select seeds

Select **Replication options**.

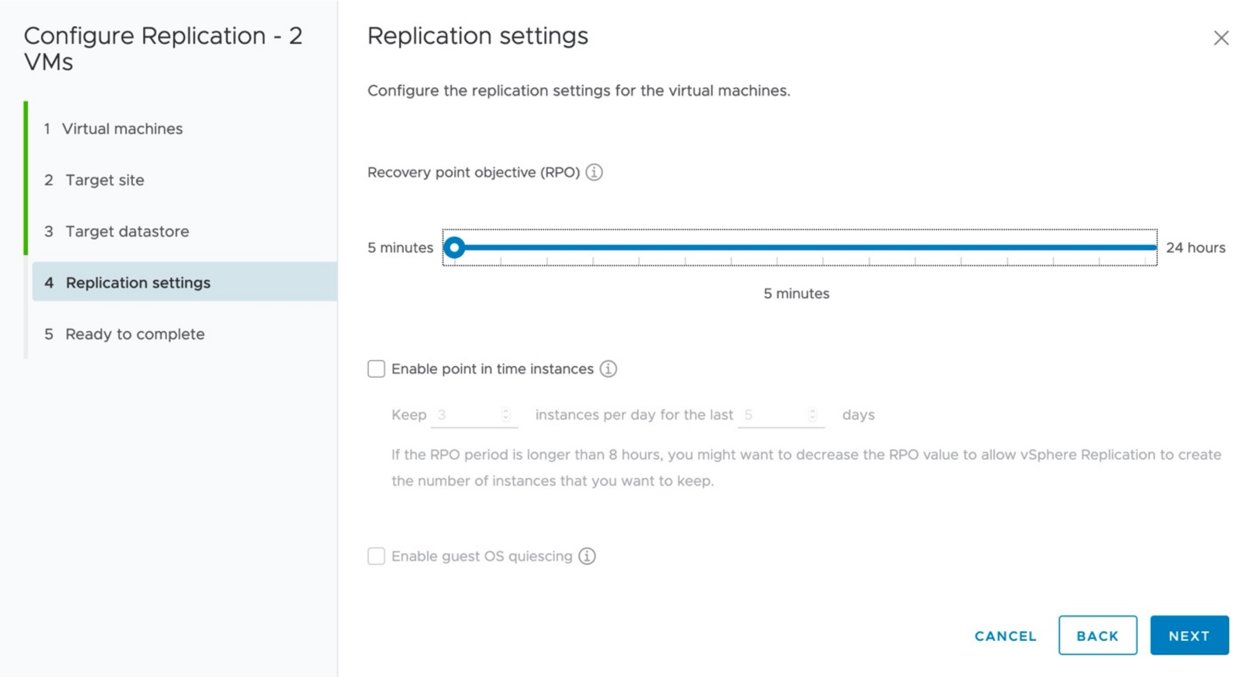

Configure Replication - 2 VMs

1 Virtual machines

2 Target site

3 Target datastore

4 Replication settings

5 Ready to complete

Replication settings

Configure the replication settings for the virtual machines.

Recovery point objective (RPO) ⓘ

5 minutes — 24 hours

5 minutes

□ Enable point in time instances ⓘ

Keep 3 instances per day for the last 5 days

If the RPO period is longer than 8 hours, you might want to decrease the RPO value to allow vSphere Replication to create the number of instances that you want to keep.

□ Enable guest OS quiescing ⓘ

Set the **Recovery Point Objective (RPO)** for replication. If **Guest OS quiescing** was not previously selected, RPO can be set for as little as 5 minutes. In this example we will not keep point in time recovery or add guest OS quiescing as we will be doing a planned migration and do not need the additional performance or capacity overhead that these features require.

Go to Protection Groups. Select "New" and create a protection group for the first group of virtual machines you wish to migrate as a group. In this example we will move all 5 virtual machines associated with a Payroll application as we need low latency between the virtual machines, and a partial migration would impact performance.

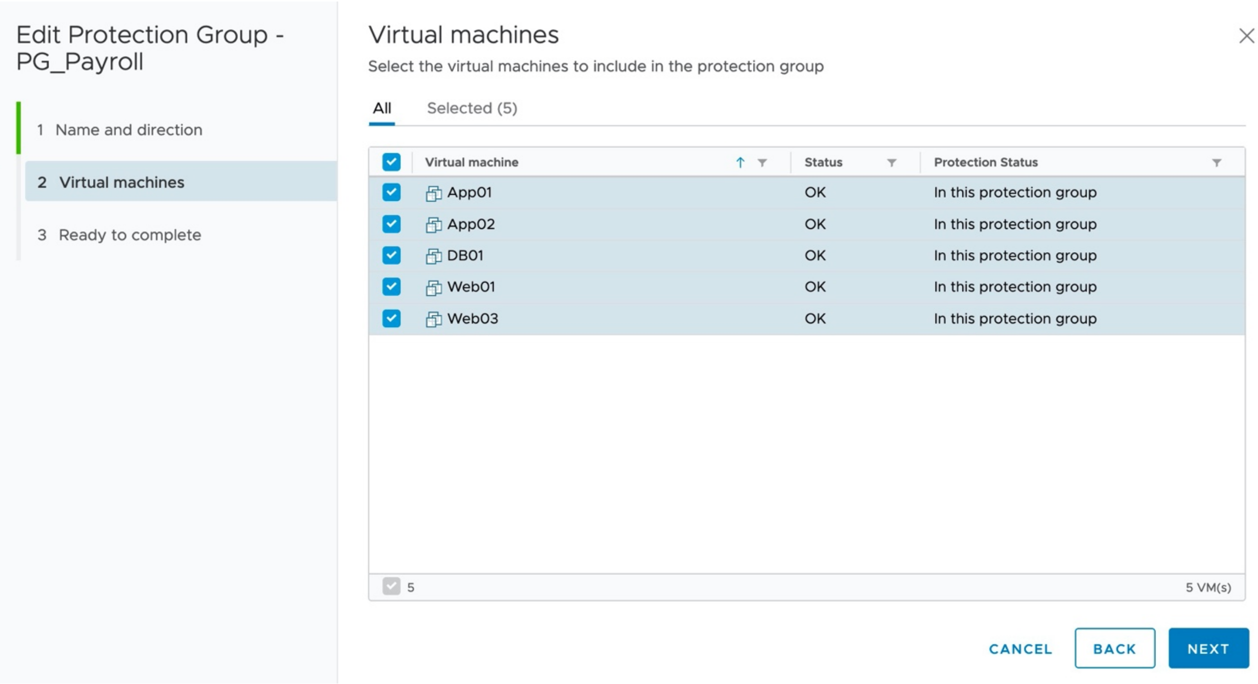

Edit Protection Group - PG_Payroll

1 Name and direction

2 Virtual machines

3 Ready to complete

Virtual machines

Select the virtual machines to include in the protection group

All | Selected (5)

☑	Virtual machine	Status	Protection Status
☑	App01	OK	In this protection group
☑	App02	OK	In this protection group
☑	DB01	OK	In this protection group
☑	Web01	OK	In this protection group
☑	Web03	OK	In this protection group

☑ 5 — 5 VM(s)

Next, we will build a Recovery Plan, and Associate the previously created protection group with it.

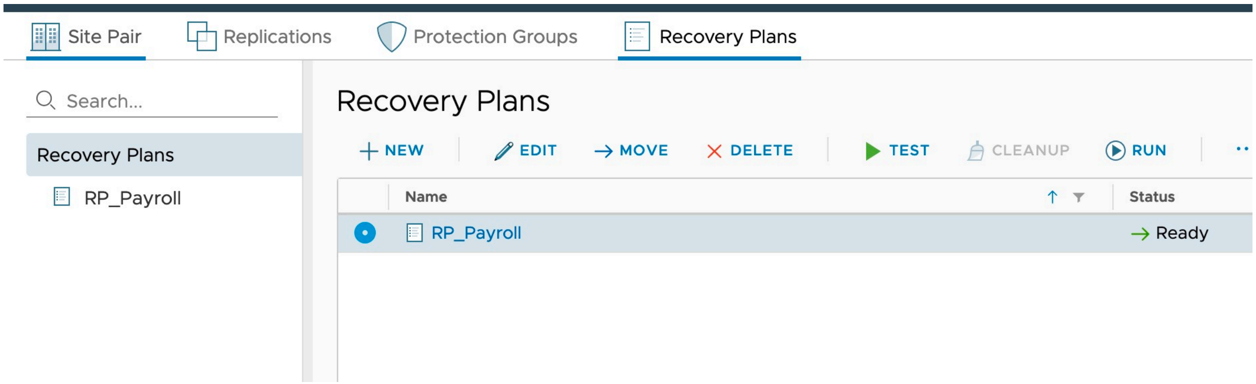

Recovery Plans

	Name	Status
●	RP_Payroll	→ Ready

Adding the entire Payroll Protection group automatically adds the 5 virtual machines to this group.

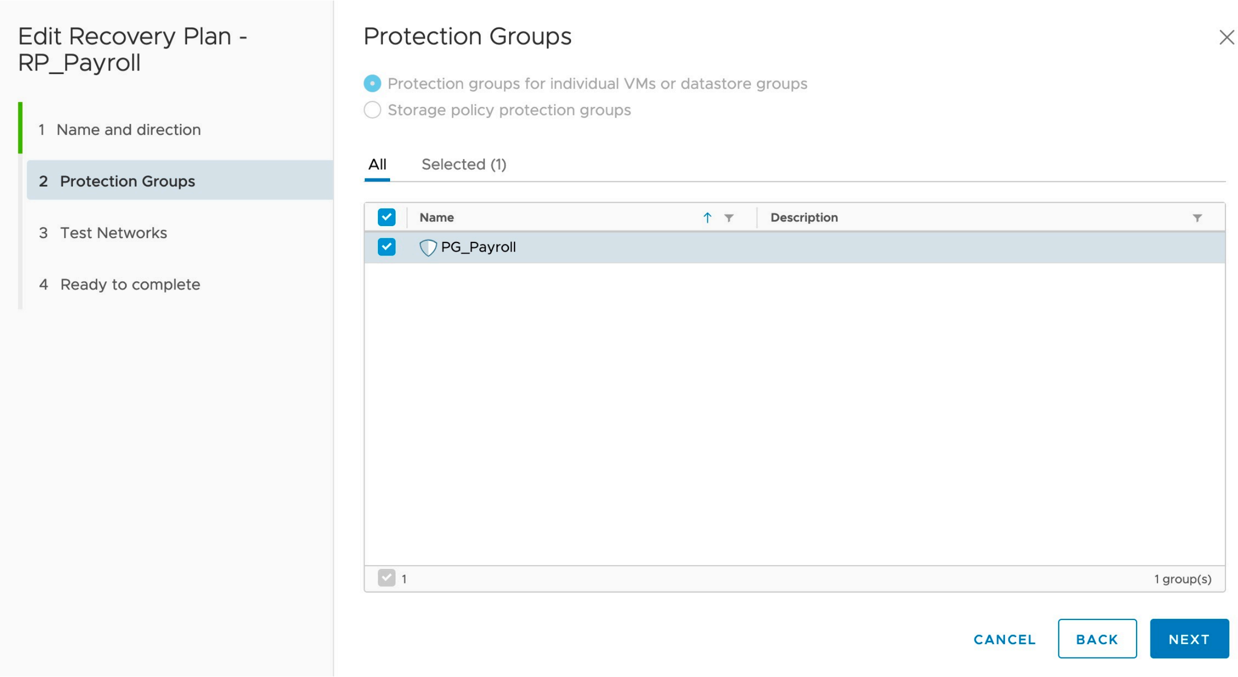

● Protection groups for individual VMs or datastore groups
○ Storage policy protection groups

☑	Name	Description
☑	PG_Payroll	

1 — 1 group(s)

We can now test our recovery Plan using the “Test” button from the recovery plan view.

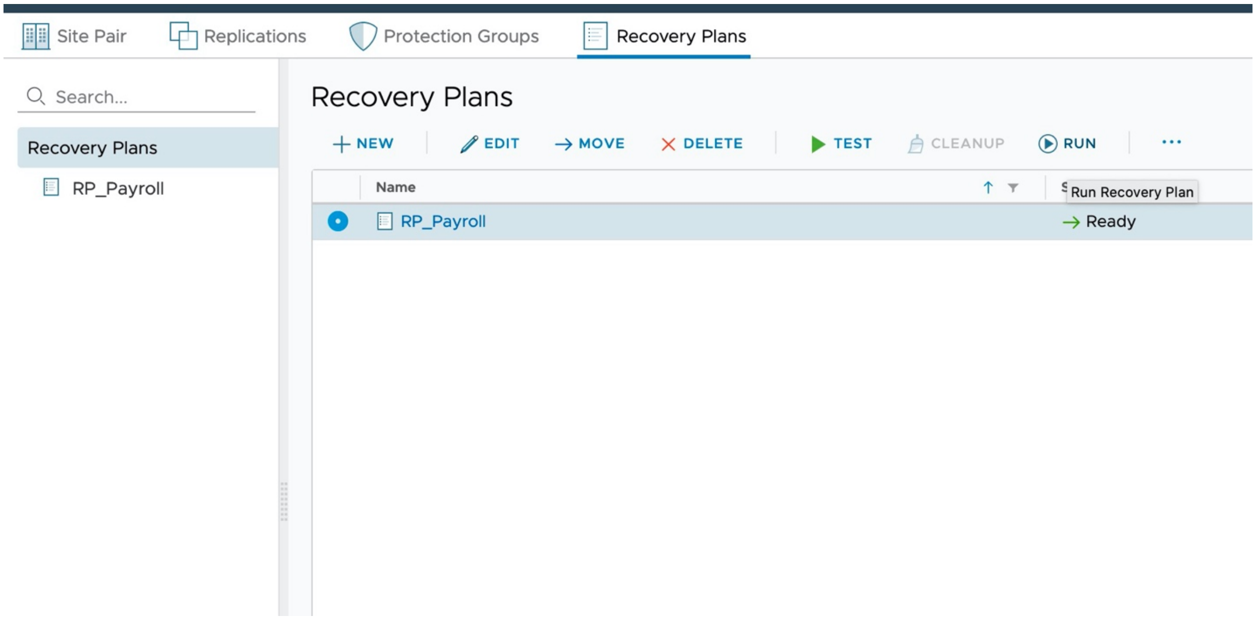

	Name	S
●	RP_Payroll	→ Ready

,

When running the test, Replicate Recent changes will force a resync to the existing RPO that has been set for the Virtual Machine.

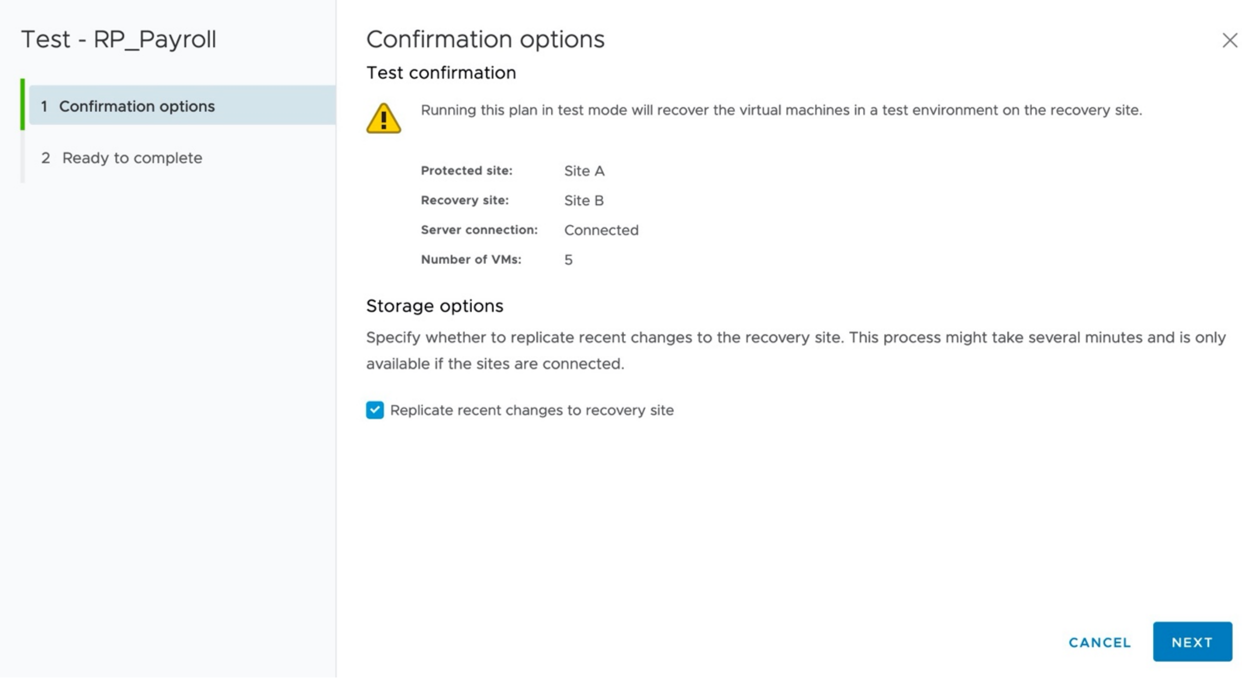

Confirmation options

Test confirmation

Running this plan in test mode will recover the virtual machines in a test environment on the recovery site.

Protected site: Site A
Recovery site: Site B
Server connection: Connected
Number of VMs: 5

Storage options

Specify whether to replicate recent changes to the recovery site. This process might take several minutes and is only available if the sites are connected.

☑ Replicate recent changes to recovery site

Once satisfied with testing, “Run” can be selected for the recovery plan. Selecting “Planned migration” will ensure a final synchronization before failover to prevent loss of data. This is the ideal selection for a migration of stateful applications.

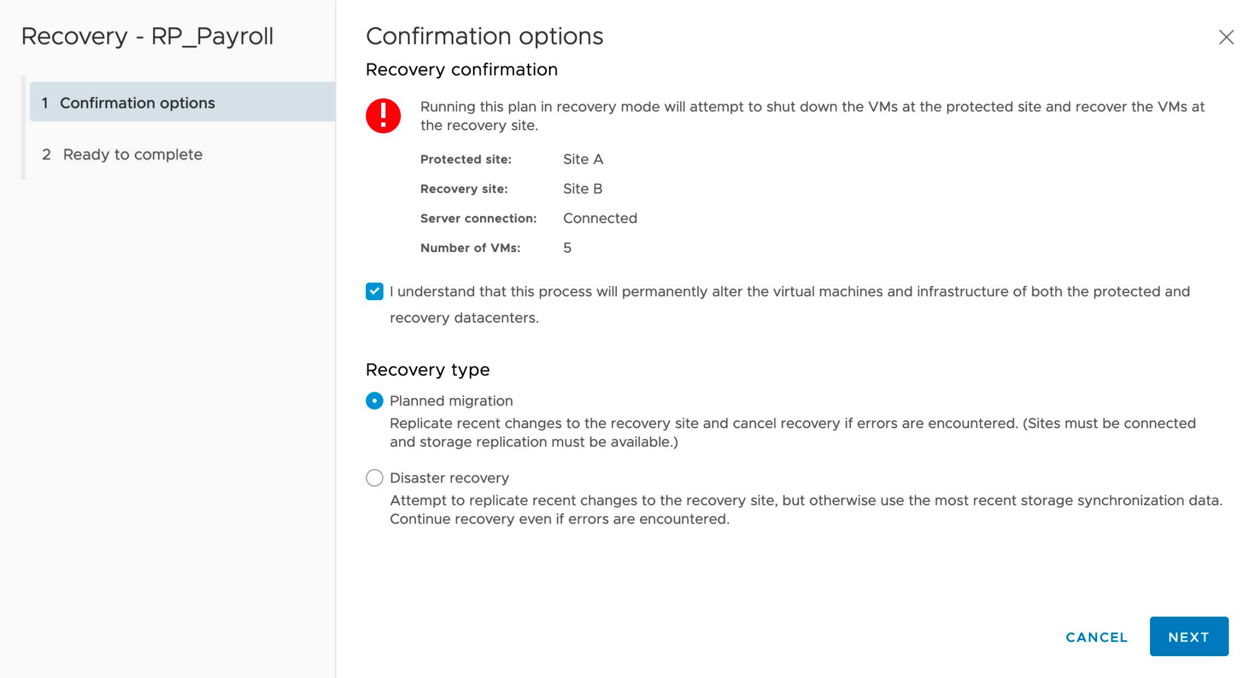

Recovery - RP_Payroll

1 Confirmation options

2 Ready to complete

Confirmation options

Recovery confirmation

Running this plan in recovery mode will attempt to shut down the VMs at the protected site and recover the VMs at the recovery site.

Protected site:	Site A
Recovery site:	Site B
Server connection:	Connected
Number of VMs:	5

☑ I understand that this process will permanently alter the virtual machines and infrastructure of both the protected and recovery datacenters.

Recovery type

● Planned migration
Replicate recent changes to the recovery site and cancel recovery if errors are encountered. (Sites must be connected and storage replication must be available.)

○ Disaster recovery
Attempt to replicate recent changes to the recovery site, but otherwise use the most recent storage synchronization data. Continue recovery even if errors are encountered.

After the final sync has completed, the VM will be registered with the target vCenter Server instance.

Additional Network Configuration

Replication traffic can be significant and can impact other traffic classes that share the same physical uplinks, or limited WAN bandwidth. Prioritization of traffic can be accomplished by two methods.

Network I/O Control (NIOC) – a mechanism to reserve bandwidth for system traffic based on the capacity of the physical adapters on a host. It enables fine-grained resource control at the VM network adapter level similar to the model that you use for allocating CPU and memory resources. This is deployed and configured automatically by VMware Cloud Foundation. By default, replication is give a “low” share weight.

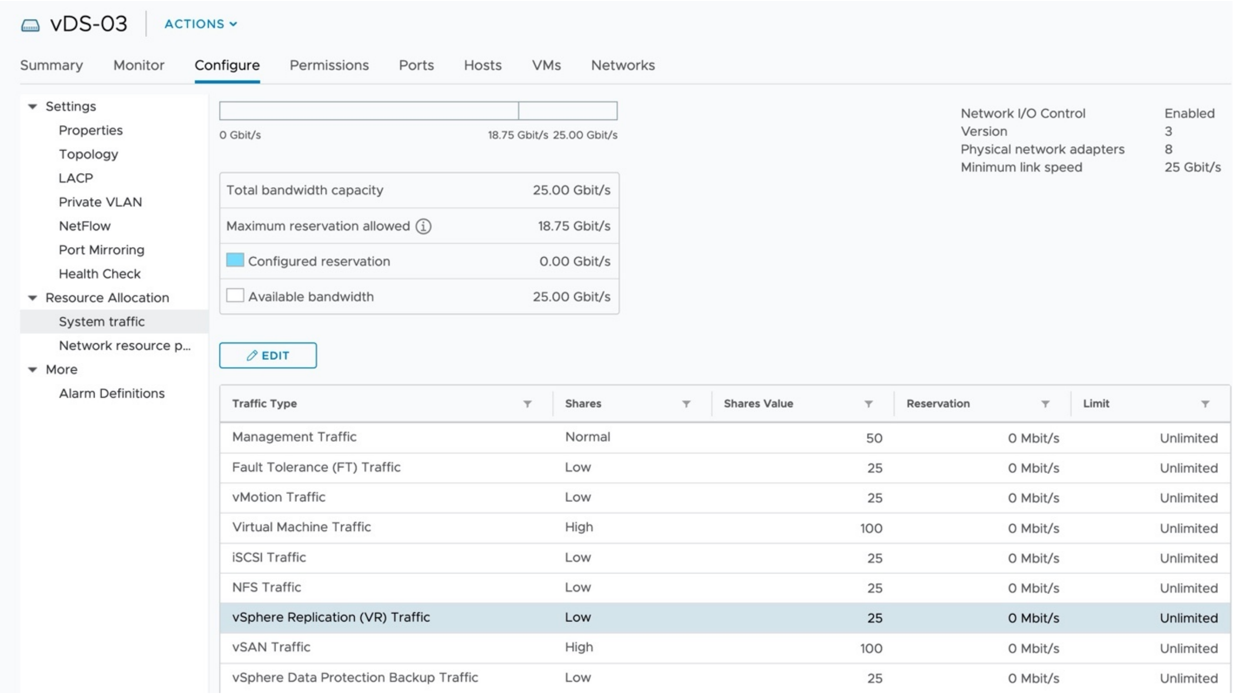

Network I/O Control	Enabled
Version	3
Physical network adapters	8
Minimum link speed	25 Gbit/s

Total bandwidth capacity	25.00 Gbit/s
Maximum reservation allowed ⓘ	18.75 Gbit/s
🟦 Configured reservation	0.00 Gbit/s
□ Available bandwidth	25.00 Gbit/s

Traffic Type	Shares	Shares Value	Reservation	Limit
Management Traffic	Normal	50	0 Mbit/s	Unlimited
Fault Tolerance (FT) Traffic	Low	25	0 Mbit/s	Unlimited
vMotion Traffic	Low	25	0 Mbit/s	Unlimited
Virtual Machine Traffic	High	100	0 Mbit/s	Unlimited
iSCSI Traffic	Low	25	0 Mbit/s	Unlimited
NFS Traffic	Low	25	0 Mbit/s	Unlimited
vSphere Replication (VR) Traffic	Low	25	0 Mbit/s	Unlimited
vSAN Traffic	High	100	0 Mbit/s	Unlimited
vSphere Data Protection Backup Traffic	Low	25	0 Mbit/s	Unlimited

Traffic filtering and marking – Found within the port group settings menu, this option allows for the applying of Quality of Services tags in the form of Class of Service (CoS) or Differentiated Services Code Point (DSCP). Note for this to work it will require all switches and devices within the network path be configured to trust or appropriately translate these tags. Please consult with your networking vendor, as well as circuit provider when using this option.

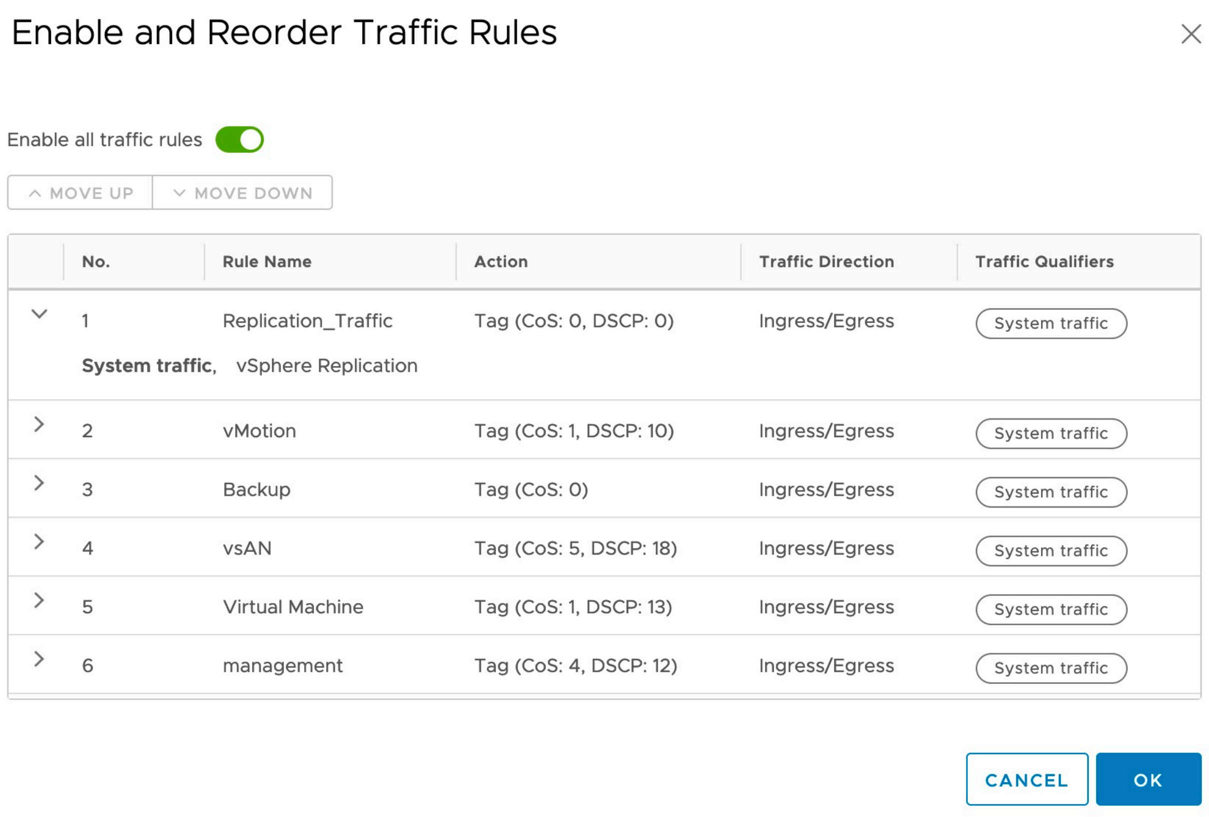

Enable and Reorder Traffic Rules

Enable all traffic rules ●

	No.	Rule Name	Action	Traffic Direction	Traffic Qualifiers
⌄	1	Replication_Traffic	Tag (CoS: 0, DSCP: 0)	Ingress/Egress	System traffic
	System traffic,	vSphere Replication			
›	2	vMotion	Tag (CoS: 1, DSCP: 10)	Ingress/Egress	System traffic
›	3	Backup	Tag (CoS: 0)	Ingress/Egress	System traffic
›	4	vsAN	Tag (CoS: 5, DSCP: 18)	Ingress/Egress	System traffic
›	5	Virtual Machine	Tag (CoS: 1, DSCP: 13)	Ingress/Egress	System traffic
›	6	management	Tag (CoS: 4, DSCP: 12)	Ingress/Egress	System traffic

Other Migration Methods

For topics such as RDM Migrations, Physical to Virtual (P2V), Windows Server Failover Clusters, and vMotion ,see the *Migrating to vSAN guide*.

Using Third-Party Migration Software

Overview

Cloud Foundation supports the use of 3rd party software that migrates data using approved supported API's such as VAIO and VADP. For a list of certified VAIO solutions see the *VAIO VCG page.*

Recommendation

Of the migration tools mentioned in this whitepaper, we recommend the use of HCX as the tool of choice for workload migrations. Many customers still have legacy vSphere 5.x sites that need live migration for mission-critical applications. HCX is the only migration tool that can meet this requirement. The remaining applications can then be migrated in bulk using a warm migration process. HCX allows customers to migrate business applications to new modernized infrastructure quickly. After legacy datacenter migrations are complete, HCX can then be used to create a hybrid cloud architecture. HCX allows customers to move applications between on-premises clouds and public cloud infrastructure freely. Only HCX provides this much flexibility for application migrations. Continue reading, and you will see why we recommend HCX.

Appendix: VM Guest–Setting Considerations

Migrating a guest VM can impact many things. Research the following issues before beginning migration activities.

When migrating a VM to different hardware or different sites, the underlying type of physical hardware will probably change. Newer physical servers can have different CPU types, instruction sets, and core counts. Some customers have tuned their database VMs to the underlying physical CPU NUMA node boundaries. For applications that are NUMA aware, research some of the latest performance-tuning settings for VMs. *This article* dives deep into best-practice performance settings for VMs.

Enhanced vMotion Compatibility (EVC) mode is another thing to consider with the physical CPU. If EVC mode is turned on in the legacy site but not in the Cloud Foundation site, it is possible to live-migrate to the Cloud Foundation site, but it might not be possible to migrate back because the VM might have acquired a new instruction set from new CPUs. Consider enabling EVC mode in Cloud Foundation to enable migration to and from Cloud Foundation instances.

When migrating to another site, consider any other local application dependencies that the VM might communicate with. Cloud Foundation Enterprise includes vRealize Network Insight, which can enable development of application dependency maps to itemize entities a VM might be communicating with.

There are many things to consider before migrating a VM to a new site or infrastructure. The following list provides some guest settings to research that might be affected by their location when migrating VMs. This list might not be all inclusive, but it provides preliminary suggestions for migration planning.

Guest Settings

- NUMA awareness
- EVC mode
- Custom CPU instruction set masking
- DNS settings

- Gateway
- Microsoft Active Directory sites and services for Windows
- Backups
- Agents
- Governance
- Tags
- Firewall rules
- Affinity rules
- Media access control (MAC) address (might be tied to software licensing)
- vSphere Network I/O Control
- VMware vSphere DirectPath I/O™
- ISO or CD-ROM attachments
- VMware vSphere Virtual Machine Encryption (VM Encryption)
- Fibre Channel NPIV
- Custom swap file location

About the Authors

Heath Johnson is a Sr. Technical Marketing Manager based in Wisconsin. Heath has been with VMware since 2015 is a VMware vExpert and has been working with VMware products since 2004. When he's not working, he is spending time with his family, flying airplanes, flying drones, cycling, or enjoying the outdoors. You can follow Heath on his personal blog at FlyingVirtually.com or on Twitter @heathbarj

John Nicholson is a Sr. Technical Marketing Manager based in Texas. John has been with VMware since 2015. He hosts a podcast vspeakingpodcast.com and You can follow John on his personal blog at www.Thenicholson.com or on Twitter @Lost_Signal

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.