



Configuring Virtualized Workload Mobility

This chapter contains the following sections:

- [Information About Virtualized Workload Mobility \(DC to DC vMotion\), on page 1](#)
- [Prerequisites for Virtualized Workload Mobility \(DC to DC vMotion\), on page 2](#)
- [Guidelines and Limitations, on page 2](#)
- [Migrating a VSM, on page 3](#)
- [Verifying and Monitoring the Virtualized Workload Mobility \(DC to DC vMotion\) Configuration, on page 4](#)

Information About Virtualized Workload Mobility (DC to DC vMotion)

This section describes the Virtualized Workload Mobility (DC to DC vMotion) configurations and includes the following topics:

- Stretched Cluster
- Split Cluster

Stretched Cluster



Note A stretched cluster is a cluster with ESX/ESXi hosts in different physical locations.

In an environment where the same Cisco Nexus 1000 instance spans two data centers, this configuration allows you to have Virtual Service Engine (VSEs) in different data centers be part of the same vCenter Server cluster.

By choosing this configuration, you ensure that the VSEs in either data center (in a two data center environment) are a part of the same Dynamic Resource Scheduling (DRS) / VMware High Availability (VMW HA) / Fault Tolerance (FT) domain that allows for multiple parallel virtual machine (VM) migration events.

Split Cluster

The Split Cluster configuration is an alternate to the Stretched Cluster deployment. With this configuration, the deployment consists of one or more clusters on either physical site with no cluster that contains VSEs in multiple data centers. While this configuration allows for VM migration between physical data centers, these events are not automatically scheduled by DRS.

Prerequisites for Virtualized Workload Mobility (DC to DC vMotion)

Virtualized Workload Mobility (DC to DC vMotion) has the following prerequisite:

- Layer 2 extension between the two physical data centers over the DCI link.

Guidelines and Limitations

Virtualized Workload Mobility (DC to DC vMotion) has the following guidelines and limitations:

- The VSM HA pair must be located in the same site as their storage and the active vCenter Server.
- Quality of Service bandwidth guarantees for control traffic over the DCI link.
- Limit the number of physical data centers to two.
- A maximum latency of 10 ms is supported for VSM-VSM control traffic when deployed across datacenters.
- A maximum latency of 100 ms is supported for VSM-VSE control traffic for both L2 and L3 mode of deployments.

Physical Site Considerations

When you are designing a physical site, follow these guidelines:

- Check the average and maximum latency between a Virtual Supervisor Module (VSM) and VSE.
- Follow the procedures to perform actions you would intend to do in normal operation. For example, VSM migration.
- Design the system to handle the high probability of VSM-VSE communication failures where a VSE must function in headless mode due to data center interconnect (DCI) link failures.

Handling Inter-Site Link Failures

If the DCI link or Layer 2 extension mechanism fails, a set of VSE modules might run with their last known configuration for a period of time.

Headless Mode of Operation

For the period of time that the VSM and VSE cannot communicate, the VSE continues to operate with its last known configuration. Once the DCI link connectivity is restored and the VSM-VSE communication is reestablished, the system should come back to its previous operational state. This mode type is no different than the headless mode of operation within a data center and has the following limitations for the headless VSE:

- No new VM ports will be brought up on Cisco N1KVE during headless mode as VSM-to-VSE communication is essential to bring up ports.
- Queries on BRIDGE and IF-MIB processed at the VSM give the last known status for the hosts in headless mode.

Handling Additional Distance/Latency Between the VSM and VSE

In a network where there is a considerable distance between the VSM and VSE, latency becomes a critical factor.

Because the control traffic between the VSM and VSE faces a sub-millisecond latency within a data center, latency can increase to a few milliseconds depending on the distance.

With an increased round-trip time, communication between the VSM and VSE takes longer. As you add VSEs and vEthernet interfaces, the time it takes to perform actions such as configuration commands, module insertions, port bring-up, and **showshow** commands increase because that many tasks are serialized.

Migrating a VSM

This section describes how migrate a VSM from one physical site to another.



Note If you are migrating a VSM on a Cisco Nexus 1010, see the Cisco Nexus 1010 Software Configuration Guide, Release 4.2(1)SP1(3).

Migrating a VSM Hosted on an ESX

Use the following procedure to migrate a VSM that is hosted on an ESX or ESXi host from the local data center to the remote data center:



Note For information on vMotion or storage vMotion, see the VMware documentation.

Before you begin

Before beginning this procedure, you must know or do the following:

- Reduce the amount of time where the VSM runs with remote storage in another data center.

- Do not bring up any new VMs or vMotion VMs that are hosted on any VSEs corresponding to the VSM that is being migrated.

Procedure

- Step 1** Migrate the standby VSM to the backup site.
- Step 2** Perform a storage vMotion for the standby VSM storage.
- Step 3** `switch#system switchover`
Initiates a system switchover.
- Step 4** Migrate the original active VSM to the backup site.
- Step 5** Perform a storage vMotion for the original active VSM storage.
-

Verifying and Monitoring the Virtualized Workload Mobility (DC to DC vMotion) Configuration

Refer to the following section for verifying and monitoring the Virtualized Workload Mobility (DC to DC vMotion) configuration:

Procedure

`switch#show module`

Displays the virtualized workload mobility (DC to DC vMotion) configuration.
