InfoScale™ 9.0 Virtualization Guide - Linux on ESXi

arctera™

Last updated: 2025-08-14

Legal Notice

Arctera US LLC | www.arctera.io

Technical Support

Technical Support maintains support centers globally. Technical Support’s primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about our support offerings, you can visit our website at the following URL:

www.arctera.io/support

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available here: https://sort.veritas.com/arctera/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to: productdocs@arctera.io

You can also see documentation information or ask a question on the Arctera community site:
https://vox.veritas.com/category/arctera-discussions/discussions/infoscale

Services and Operations Readiness Tools (SORT)

Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:
https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1

Overview

Chapter 1

About InfoScale solutions in a VMware environment

This chapter includes the following topics:

Overview of the InfoScale Virtualization Guide

This document provides information about InfoScale products support for VMware virtualization technology. It contains:

- High-level conceptual information for InfoScale products and how they function in ESXi environments.
- High level implementation information for setting up InfoScale products in ESXi environments.
- Use case chapters with examples of how InfoScale products can improve performance outcomes for common ESXi use cases.

How InfoScale solutions work in a VMware environment

Using InfoScale solutions in a VMware environment means that the InfoScale product runs in the operating system, inside the Virtual Machine (VM).

The InfoScale component, such as InfoScale, does not run inside the VMware ESXi kernel or in the Hypervisor.

Figure 1-1 shows an example of the high-level architecture diagram with InfoScale running in the VM.

Figure 1-1 Architecture overview

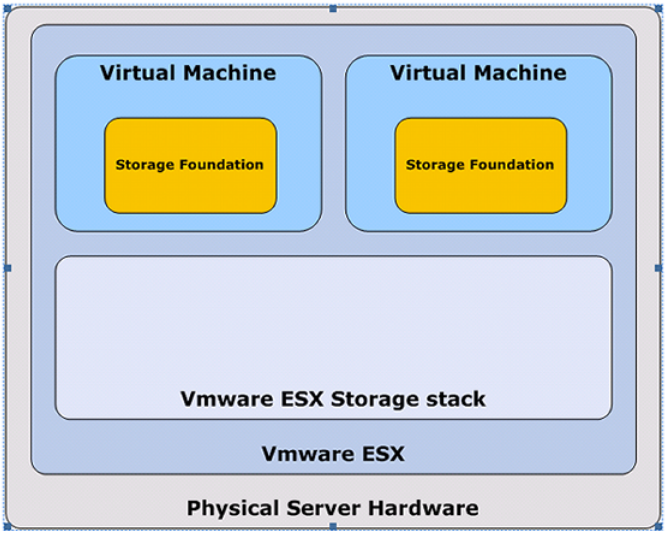

Figure 1-2 shows the I/O path from the Virtual Machine to the storage.

Figure 1-2 I/O path from Virtual Machine to storage

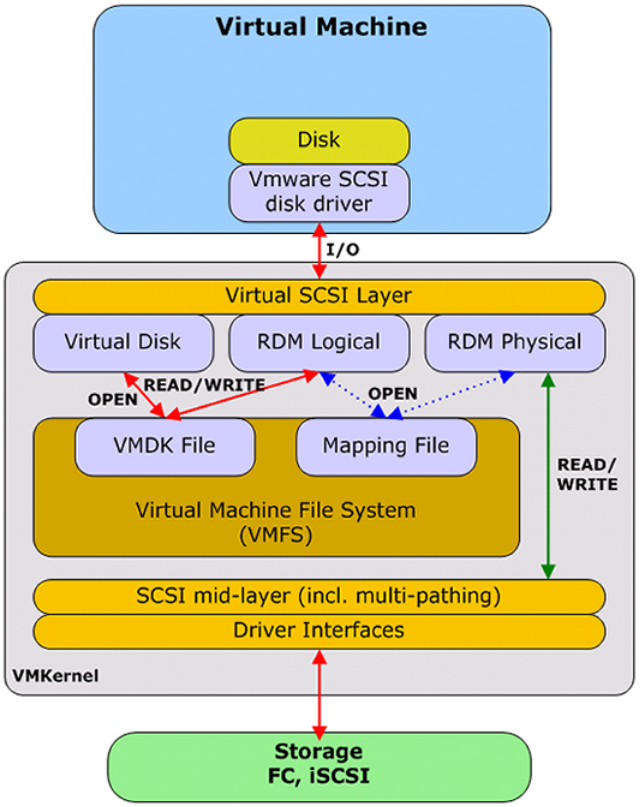

VMware has several different methods to allocate block storage to a virtual machine:

- File-based virtual disks created in VMFS or from NFS - Virtual Disk
- Block storage mapped from local disk, Fibre Channel LUNs or iSCSI – Raw Device Mapping

VMware must be configured to use Raw Device Mapping for certain features of Volume Manager (VxVM) to operate as they do in a physical server environment.

See "When to use Raw Device Mapping and InfoScale" on page 11.

How InfoScale product components enhance VMware capabilities

In VMware vSphere, VMFS does not have the capability to mirror storage. This forces users to use mirrored LUN's to provide this functionality to the virtual machines. With Volume Manager (VxVM) running in the virtual machine, utilizing

raw device mapping, data can be protected with mirroring in the virtual machine, including the ability to mirror between storage arrays.

With InfoScale Enterprise and Flexible Storage Sharing, you can also mirror storage on local disks directly attached to an ESXi host to local disks directly attached to a remote ESXi host, whether these disks are presented as raw device mappings or VMDK files.

InfoScale can make painful migrations of data from physical to virtual environments easier and safer to execute. With InfoScale there is no need to actually copy any data from source to destination; rather, the administrator re-assigns the same storage (or a copy of it for a test migration) to the virtual environment. Once the storage is assigned and configured into the virtual machine, VxVM will scan the device tree and discover the disk group and volume structure.

Figure 1-3 describes an example workflow.

Figure 1-3 Migration workflow

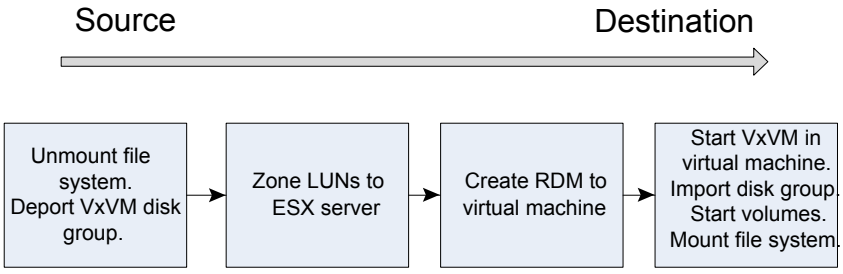

VxVM is agnostic to the actual physical device entry. That is, VxVM does not care if the device is /dev/sdb or /dev/sdaz. This transparency makes it easy to move storage from one node to another, or between physical and virtual machines.

When to use Raw Device Mapping and InfoScale

Raw Device Mapping (RDM) enables a virtual machine to have direct access to the storage rather than going through VMFS. RDM is configured per physical storage device, i.e. a disk or LUN is assigned to one or more virtual machines. It is not possible to assign a part of a physical storage device to a virtual machine. Different types of storage (local SCSI disks, iSCSI disks, Fibre Channel disks) can be used with raw device mapping; Volume Manager (VxVM) supports all three types of disks.

Note: The InfoScale products work well with the iSCSI disks mapped directly to the Virtual Machines.

VMware provides two different modes for raw device mapping:

- Logical mode offers the same functionality and compatibility as a Virtual Disk with respect to VMware ESXi features.
- Physical mode is the most similar method to storage access in a non-virtual environment. Only one SCSI command, REPORT_LUNS, is virtualized as it is required to enable vMotion and a few other features in VMware.
  With InfoScale, physical mode is recommended as it enables maximum functionality of VxVM in a VMware environment.

The different modes affect the functionality and behavior of InfoScale. It is important to use the correct mode for the desired functionality. The benefit of each storage access method is dependent on the workload in the virtual machine. It is easy to get started with one way of deploying storage without considering the long-term implications because of the ease of use of the virtual environment.

For applications with little to no storage need, using raw device mapping is overkill and not recommended. Also, if your environment depends on VMware snapshots, using Raw Device Mapping in physical mode is not possible as it is not supported by VMware.

Raw Device Mapping is a great fit for:

- Applications with large storage needs
- Applications that need predictable and measurable performance
- Multi-node clusters using disk quorums
- Applications with storage that is currently managed by InfoScale but is moving into a virtual environment
- Applications that require direct access to storage, such as storage management applications

Array migration

When a virtual machine (VM) is selected for movement of disk files to a new DataStore (DS) using Storage vMotion, the following takes place:

- The VM home directory (config, log, swap, snapshots) is copied to the destination DS.
  A "shadow" VM is started on the destination DS using the copied files. The “shadow” VM idles waiting for the copying of the VM disk file(s) to complete.
- An initial copy of the VMs disk file(s) is made to the target DS. During the copy changes made to the source are tracked (change block tracking).

- Storage vMotion iteratively repeats this process of copying the changed blocks from the source DS to the destination DS.
- When the amount of outstanding changed blocks is small enough, vMotion invokes a Fast Suspend and Resume (FSR) of the VM (similar to vMotion) to transfer the running VM over to the idling shadow VM. As is the case with regular vMotion, this transfer normally happens so quickly that it will be completely transparent to the VM.
- After the FSR completes the old home directory the VM disk files are deleted from the source DS.

Figure 1-4 describes the high level process for VMware.

Figure 1-4 Virtual machine data migration with ESXi tools

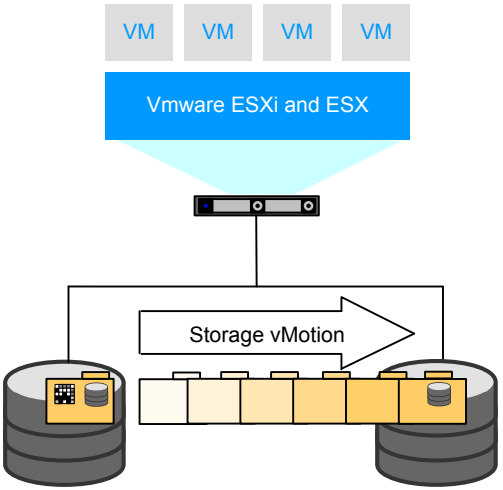

In VMware, you can add disks to the ESXi server and the virtual machine without rebooting. This functionality makes it possible to offer a better process using InfoScale together with raw device mapped storage for online data migration.

Figure 1-5 describes the high level process for InfoScale.

Figure 1-5 Virtual machine data migration with InfoScale

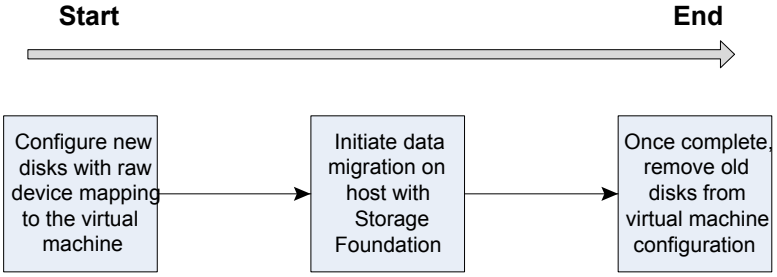

Data migration for InfoScale can be executed either locally in the virtual machine with Volume Manager (VxVM) or in a central location, migrating all storage from an array utilized by InfoScale managed hosts. This powerful, centralized data migration functionality is available in InfoScale Operations Manager.

See the following Web site for information about InfoScale Operations Manager.

https://www.veritas.com/product/storage-management/infoscale-operations-manager

InfoScale component limitations in an ESXi environment

Some limitations apply for using InfoScale Storage and InfoScale Enterprise with VMware ESXi:

- Sharing VMDK files between virtual machines
  When sharing VMDK files between virtual machines, SCSI BUS sharing mode for the corresponding SCSI controllers can be set to either “Physical” or “Virtual” modes. Setting this mode to “Physical” causes SCSI reservation conflict and I/O error on DMP. This issue occurs in LVM and raw disks also.
  Solution:
  Disable simultaneous write protection for the shared disk using the multi-writer flag. The procedure is described in the following VMware Knowledge Base article:
  http://kb.vmware.com/kb/1034165
- Volume Replicator (VVR) option
  VVR is supported inside a Virtual Machine. Keep in mind that VVR can use a significant amount of network bandwidth (depends on the amount of data written to disk) and this can reduce the available network bandwidth for other virtual machines. Since the network card is a shared resource, ensure that enough bandwidth is available to sustain the writes happening to disk.
- VMware snapshots
  The following limitations apply for VMware snapshots:

- VMware snapshots are not supported when raw device mapping is used in physical mode, regardless of whether InfoScale components are installed or not. The REDO-log functionality that is required for VMware Snapshots is not available with raw device mapping used in the physical mode.
  InfoScale components support VMDK file utilization as a backend storage to overcome this limitation.
  InfoScale components also support raw device mapping in the logical mode and VMware snapshots, since RDM-logical mode uses the same level of SCSI virtualization as VMDK files.
  See "When to use Raw Device Mapping and InfoScale" on page 11.
- VMware snapshots are not supported when using VMDK files as backend storage and the multi-writer flag has been set. For example, when using InfoScale Enterprise with VMDK files. This is a VMware limitation.
- VMware snapshots are not supported with any VM using disks managed by the VMwareDisks agent.

- vMotion (Live Migration)
  - For InfoScale Foundation: vMotion is supported.
  - For InfoScale Availability and Cluster Server (VCS):
    VMware vMotion has a limitation that affects all clustering software. vMotion is not supported when a virtual SCSI controller is set to have sharing enabled. Virtual SCSI controller sharing is a virtual machine attribute and is required to be set for the virtual machines that share storage between each other (on the same physical ESXi server or between physical ESXi servers). Essentially all clustering products that rely on SAN storage require this attribute to be set.
    VCS provides the VMwareDisks agent to override this limitation and enable shared storage for InfoScale Availability to operate. InfoScale Availability supports VMDK files and therefore vMotion.
    See "About setting up InfoScale Enterprise on VMware ESXi" on page 57.
    This limitation does not affect the virtual machines that do not have the sharing attribute turned on for their virtual SCSI controllers.
  - For InfoScale Enterprise: vMotion is supported.
- N-Port ID Virtualization (NPIV)
  NPIV used with InfoScale Foundation is fully supported. No additional setup tasks are required for InfoScale Foundation when the storage is NPIV-enabled.
  VMware currently does not support I/O fencing with NPIV for any other 3rd party clustering software other than MSCS. In VMware environments, InfoScale products supports I/O fencing using the Coordination Point server as an arbitration mechanism.

I/O fencing considerations in an ESXi environment

VMware does not support SCSI-3 Persistent Reservations (and hence I/O Fencing) with any other 3rd party clustering software with RDM logical mode or VMDK-based virtual disks. In VMware environments, SFHA and Arctera InfoScale Enterprise support the following methods of fencing:

- disk-based fencing with RDM-P mode.
  Available starting with SFHA and Arctera InfoScale Enterprise version 5.1 Service Pack 1 Rolling Patch 1.
  See the following tech note for details.
  https://www.veritas.com/support/en_US/article.TECH169366
- non-SCSI-3 PR-based fencing using the Coordination Point (CP) server.
  The CP server provides arbitration amongst the multiple nodes.

I/O fencing utilizes HBA World Wide Numbers (WWNs) to create registrations on the storage; this has implications in a virtual environment where the HBA is shared between virtual servers on the same physical ESXi host as the WWN used for I/O fencing ends up being the same for each virtual machine. Therefore, Arctera InfoScale Enterprise virtual machines (in the same Arctera InfoScale Enterprise cluster) cannot share physical servers as the I/O fencing behavior will result in all nodes from that physical ESXi host being fenced out if an event triggers the fencing functionality. In short, if I/O fencing is configured, the Arctera InfoScale Enterprise nodes (in the same Arctera InfoScale Enterprise cluster) have to be running on separate physical ESXi hosts.

About InfoScale™ solutions support for the VMware ESXi environment

InfoScale™components support the VMware ESXi environment as follows:

Table 1-1 InfoScale™solutions support for ESXi virtualization components

InfoScale™ components or solutions	ESXi host	ESXi guest
InfoScale	N	Y
Cluster Server (VCS)	N	Y
Storage Foundation and High Availability (SFHA)	N	Y
InfoScale Enterprise	N	Y

Table 1-1 InfoScale™solutions support for ESXi virtualization components *(continued)*

InfoScale™ components or solutions	ESXi host	ESXi guest
Storage Foundation for Oracle RAC (SF Oracle RAC)	N	Y
Replicator Option	N	Y

Note: * DMP has limited functionality in the guest: multi-pathing works, however it shows only a single path. DMP can encapsulate the disk if mapping the physical disk to the guest.

Virtualization use cases addressed by InfoScale

InfoScale™ supports the following VMware environment use cases:

Table 1-2 Virtualization use cases addressed by InfoScale in an ESXi environment

Virtualization use case	Arctera solution	Implementation details
Application management and availability	InfoScale Availability or InfoScale Enterprise in guest	How to manage application monitoring and failover on virtual machines using VCS. See "About application availability with Cluster Server (VCS) in the guest" on page 38.
High availability for live migration	InfoScale Availability or InfoScale Enterprise in guest	How to use VCS to provide high availability for live migration with vMotion. See "About VCS support for Live Migration" on page 38.
Storage visibility	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in ESXi host	How to use DMP features to improve storage visibility. See DMP documentation.
Storage availability	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in ESXi host	How to use DMP features to improve storage availability. See DMP documentation.

Table 1-2 Virtualization use cases addressed by InfoScale in an ESXi environment *(continued)*

Virtualization use case	Arctera solution	Implementation details
Improved I/O performance	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in ESXi host	How to use DMP features to improve I/O performance.
Simplified path management with DMP	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in guest	How to use DMP features for end-to-end storage path visibility, simplified management, and improved performance. See DMP documentation.
Data protection	InfoScale Storage or InfoScale Enterprise in guest	How to use InfoScale backup and recovery features for data protection. See "Protecting data with InfoScale product components in the VMware guest" on page 44.
Storage optimization	InfoScale Storage or InfoScale Enterprise in guest	How to use InfoScale thin provisioning, FileSnap, SmartTier, and SmartIO features to optimize storage in a VMware environment. See "Optimizing storage with InfoScale product components in the VMware guest" on page 46.
Data migration	InfoScale Storage or InfoScale Enterprise in guest	How to use Storage Foundation Portable Data Containers to migrate data safely and easily in a VMware guest environment. See "Migrating data with InfoScale product components in the VMware guest" on page 51.
Improved database performance	InfoScale Storage or InfoScale Enterprise in guest	How to use InfoScale database accelerators to improve database performance in a VMware guest environment. See "Improving database performance with InfoScale product components in the VMware guest" on page 51.
Simplified storage management	InfoScale Foundation or InfoScale Storage or InfoScale Enterprise in guest	How to use DMP for VMware and InfoScale features for end-to-end storage visibility, performance, optimization, and enhanced ease of management.

Table 1-2 Virtualization use cases addressed by InfoScale in an ESXi environment *(continued)*

Virtualization use case	Arctera solution	Implementation details
Application high availability and fast failover	InfoScale Enterprise in guest	How to manage application high availability and fast failover See "About use cases for InfoScale Enterprise in the VMware guest" on page 54.

Section 2

Deploying Veritas InfoScale products in a VMware environment

Chapter 2

Getting started

This chapter includes the following topics:

Storage configurations and feature compatibility

All block storage topologies that are supported with ESXi are supported when InfoScale is running inside a virtual machine. The storage specific details are hidden for InfoScale by VMware hence FC, iSCSI and locally attached disks are supported.

Table 2-1 vSphere features compatible with shared storage configurations

VM storage configurations	Compatible vSphere features			
	vMotion	DRS	VMware HA	VMware Snapshots
VMDK on VMFS	Y	N	Y	N

Table 2-1 vSphere features compatible with shared storage configurations *(continued)*

VM storage configurations	Compatible vSphere features			
	vMotion	DRS	VMware HA	VMware Snapshots
RDM-P with FC / iSCSI SAN For further information, refer to the InfoScale Support for VMotion in VMware Environments with RDM-P Disks document.	Y	N	N	N
RDM-V with FC / iSCSI SAN All VMware-related vMotion limitations shall continue to apply. For details, refer to the VMware vSphere documentation.	Y	N	N	N
VMDK on NFS	N	N	N	N
iSCSI inside guest	Y	Y	Y	Y *
NFS inside guest (using Mount agent)	Y	Y	Y	Y *

* Taking snapshots is possible if VCS and all applications are fully stopped. Before reverting to the snapshot, shutdown VCS and all applications and then revert to the snapshot. Ensure that no software updates or configuration changes related to VCS have been applied post the snapshot.

Refer to appendix A for more on known limitations before moving to a configuration with VCS in the guest.

About setting up VMware with InfoScale™ products

Before setting up your virtual environment, verify that your planned configuration meets the requirements to install InfoScale™ products.

Refer to the following information for the system requirements, licensing, and other considerations for installation of InfoScale™ product components.

- Licensing: Customers running InfoScale or InfoScale Enterprise in a VMware environment are entitled to use an unlimited number of guests on each licensed server or CPU.
- InfoScale™ product requirements: Each InfoScale™ product has system requirements and supported software. Refer to the *Release Notes* for each product.
  In addition, InfoScale™ may require specific configuration in the VMware environment.
  See "Storage configurations and feature compatibility" on page 21.
- Release Notes: For each InfoScale™ product, the release notes contains last-minute news and important details, including updates to system requirements and supported software. Review the *Release Notes* for the latest information before you start installing the product.
  The product documentation is available on the web at the following location:
  https://sort.veritas.com/documents

InfoScale™ products support for VMware environments

InfoScale™ products are a set of components that provide storage administration and management in a heterogeneous storage environment. Arctera supports the InfoScale™ product components running within the virtual machine on VMware ESXi.

For information about the versions of VMware ESXi and guest operating systems on which Arctera tests and supports the InfoScale™ product components, refer to the software compatibility list (SCL) at:

For information about the supported operating systems, see the Linux Release Notes for the InfoScale™ product version you are using.

See "About InfoScale™ solutions support for the VMware ESXi environment" on page 16.

Installing and configuring storage solutions in the VMware virtual environment

To set up a guest in VMware environment with InfoScale™ products, install and configure the InfoScale™ product component on the required virtual machines.

Refer to the installation guide for your InfoScale™ product for installation and configuration information.

See "InfoScale documentation" on page 71. for the installation guide

Recommendations for improved resiliency of InfoScale clusters in virtualized environments

Arctera recommends that you configure the following settings to improve the resiliency of InfoScale cluster configurations in virtualized environments:

- **Peerinact**: Set the default LLT tunable parameter peerinact to 32 seconds instead of 16 seconds. Doing so helps improve the stability of the cluster in virtualized environments, where multiple external factors as described further in this list, can affect the stability of the cluster.
- **Provisioning ratio**: The CPU and memory provisioning ratio affects the stability of the InfoScale cluster. To ensure maximum stability, set the ratio to the lowest value possible. For critical solutions that require maximum resiliency, the ratio must be set to **1:1**.
- **CPU load on host operating systems**: Although the provisioning ratio is low, the CPU load on the host operating systems still plays a part in cluster stability. If the load on the host operating system is very high, it can affect how vCPUs on the guest VMs are scheduled, because vCPUs are processes from the perspective of the host servers.
- **CPU requirement of the actual workload on guests**: When the total CPU requirement for workloads exceeds the available physical CPU capacity, it causes node evictions due to heartbeat timeouts.
- **External events**: External events like live migration of the guest VMs, virtualized disk backups, and so on, are known to add CPU load on the host servers. To reduce this additional load on the CPU, watch the stun duration in your environment caused by these events, and increase the peerinact value, if required. Increase the peerinact value only in these conditions and not in any other circumstances.
- **Hypervisor**: Always follow the best practices for the hypervisor.
  For details, refer to the VMware article at:
  https://kb.vmware.com/s/article/2001003

Chapter 3

Understanding Storage Configuration

This chapter includes the following topics:

Configuring storage

There are two options to provide storage to the Virtual Machines (VMs) that will host the Cluster File System:

- The first option, Raw Device Mapping Protocol (RDMP), uses direct access to external storage and supports parallel access to the LUN, but does not allow vMotion or DRS. For RDMP configuration, you must map the raw device to each VM and make sure you select the Physical (RDM-P) configuration, so SCSI-3 PGR commands are passed along to the disk.

- The second option, VMFS virtual disk (VMDK), provides a file that can only be accessed in parallel when the VMFS multi-writer option is enabled. This option supports server vMotion and DRS, but does not currently support SCSI-3 PR IO fencing. The main advantage of this architecture is the ability to move VMs around different ESXi servers without service interruption, using vMotion.

Let us consider an example where VMDK is configured with the multi-writer option enabled. In the following section, steps required to configure the virtual machines on ESXi server with shared VMDK files is covered. Also steps to configure Arctera InfoScale Enterprise in the same environment to consume the storage and create file system is also mentioned.

By following the steps in the VMware article, simultaneous write protection provided by VMFS is disabled using the multi-writer flag. When choosing this configuration, users should be aware of the following limitations and advantages

To understand the support for VMDK files with the multi-writer option, refer to the VMware article at: http://kb.vmware.com/kb/1034165

By following the steps in the VMware article, simultaneous write protection provided by VMFS is disabled using the multi-writer flag. When choosing this configuration, users should be aware of the following limitations and advantages.

Limitations:

- Virtual disks must be thick provisioned eager zeroed.
- VMDK sharing is limited to eight ESXi servers
- Linked clones and snapshots are not supported. Be aware that other vSphere activities utilize cloning and that backup solutions leverage snapshots via the vAPIs, so backups may be adversely impacted.
- SCSI-3 PR IO fencing is not supported by VMDK files.
- Special care needs to be taken when assigning VMDKs to VMs. Inadvertently assigning a VMDK file already in use to the wrong VM will likely result in data corruption.
- Storage vMotion is not supported

Advantages

Server vMotion is supported.

The lack of SCSI-3 PR IO fencing support requires the usage of at least three Coordination Point servers, to provide non-SCSI-3 fencing protection. In case of a split brain situation, CP servers will be used to determine what part of the sub-cluster will continue providing service. Once the multi-writer flag is enabled

on a VMDK file, any VM will be able to mount it and write, so special care in the provisioning phase needs to be taken.

Note: Note that if the number of Arctera InfoScale Enterprise nodes is greater than eight, several nodes will have to run in the same ESXi server, based on the limitation that a maximum of eight ESXi servers can share the same VMDK file. For example, if you are running at the Arctera InfoScale Enterprise maximum of 64 nodes, those 64 VMs would share the same VMDK file, but you could only use eight ESXi servers to host the cluster.

These are the steps that need to be taken when configuring VMDKs as shared backed storage and that will be presented in the next sections:

Table 3-1 Steps to configure VMDK

Storage deployment task	Deployment steps
Enabling Disk UUID on virtual machines (VMs)	See "Enabling disk UUID on virtual machines" on page 27.
Installing Array Support Library (ASL) for VMDK on cluster nodes	See "Installing Array Support Library (ASL) for VMDK on cluster nodes" on page 28.
Excluding the boot disk from the Volume Manager configuration	See "Excluding the boot disk from the Volume Manager configuration" on page 29.
Creating the VMDK files	See "Creating the VMDK files" on page 29.
Mapping VMDKs to each virtual machine	See "Mapping the VMDKs to each virtual machine (VM)" on page 31.
Enabling the multi-write flag	See "Enabling the multi-write flag" on page 32.
Getting consistent names across nodes	See "Getting consistent names across nodes" on page 34.
Creating a Cluster File System	See "Creating a file system" on page 35.

Enabling disk UUID on virtual machines

You must set the **disk.EnableUUID** parameter for each VM to "TRUE", if you plan to setup a high availability configuration. This step is necessary so that the VMDK always presents a consistent UUID to the VM, thus allowing the disk to be mounted properly. For each of the virtual machine nodes (VMs) that will be participating in the cluster, follow the steps below from the vSphere client:

To enable disk UUID on a virtual machine

1. Power off the guest.
2. Select the guest and select **Edit Settings**.
3. Select the **VM Options** tab on top.
4. Check to see if the parameter **disk.EnableUUID** is set, if it is there then make sure it is set to **TRUE**.

   If the parameter is not there, select **Add Configuration Params** and add it.
5. Power on the guest.

Installing Array Support Library (ASL) for VMDK on cluster nodes

In order for the cluster file system to work properly with the VMDK files, an ASL must be installed in the virtual server. The ASL package (VRTSaslapm) version that contains the VMDK ASL is 8.0.0.2501.

Note: Any future updates to the VMDK ASL will be published on http://sort.veritas.com and will have a higher revision than 6.0.100.100.

To download the ASL package

1. Go to http://sort.veritas.com.
2. Under **Downloads**, click **ASL/APM/DDI/DDL**.
3. Select the ASL package by appropriately filtering the array information.
4. Click **Download** in the **Array Support Library (ASL)/Array Policy Module (APM) Details** page.

On each clustered file system node, perform the following steps. The steps are illustrated with the example installation and your details may vary.

To install the ASL package

1. To install the package, follow the instructions outlined in the Readme file (VRTSaslap_readme.txt) which is displayed towards the end of the **Array Support Library (ASL)/Array Policy Module (APM) Details** page.

   You can also save this Readme file by selecting the **Save As...** option.
2. Follow the steps outlined in the *Installation Procedure* section of the Readme.

In case of the mentioned installation exmaple, RHEL 7 is present as a guest OS.

After installing the ASL, you will notice that the disk has been renamed from disk_0 to vmdk0_0. Before ASL:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
disk_0          auto:none     -             - online       invalid
```

After ASL has been deployed:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
vmdk0_0         auto:none     -             - online       invalid
```

vmdk0_0 is the boot disk that is to be excluded from Volume Manger configuration.

Excluding the boot disk from the Volume Manager configuration

It is a best practice to exclude the boot disk from Volume Manager. This allows the shared VMDK files to be configured to use the same name. In order to exclude the disk, run the command `vxdmpadm` with the name of the boot disk. In the example installation:

```
# vxdmpadm exclude dmpnodename=vmdk0_0
```

Verify that the boot disk is no longer reported under the VxVM configuration:

```
# vxdisk list
DEVICE        TYPE          DISK          GROUP          STATUS
```

Creating the VMDK files

The VMDKs that will be used by InfoScale Enterprise can be created either by the vSphere GUI or using the command line. Using the GUI, there is no control for the name of the file used, and they will be stored under the folder belonging to the VM that is creating the files. We would prefer in this case to control those file names, so we will use the command line to create the following configuration:

Table 3-2 Virtual disk configuration

Data Store	Virtual Disk on ESXi	VMDK NAME	Virtual device	SCSI Driver	Virtual size (GB)
DS1	Hard disk 2	cfs0/shared1.vmdk	SCSI 1:0	Paravirtual	90
DS2	Hard disk 3	cfs0/shared2.vmdk	SCSI 1:1	Paravirtual	90
DS3	Hard disk 4	cfs0/shared3.vmdk	SCSI 1:2	Paravirtual	90
DS4	Hard disk 5	cfs0/shared4.vmdk	SCSI 1:3	Paravirtual	90
DS5	Hard disk 6	cfs0/shared5.vmdk	SCSI 1:4	Paravirtual	90

To create the infrastructure

1 Connect to one of the ESXi virtual machines.

2 Create a folder called cfs0 (the name of the cluster) in each of the datastores:

```
mkdir /vmfs/volumes/DS1/cfs0
mkdir /vmfs/volumes/DS2/cfs0
mkdir /vmfs/volumes/DS3/cfs0
mkdir /vmfs/volumes/DS4/cfs0
mkdir /vmfs/volumes/DS5/cfs0
```

3 Create each of the VMDKs that will be used:

```
vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS1/cfs0/shared1.vmdk

vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS2/cfs0/shared2.vmdk

vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS3/cfs0/shared3.vmdk

vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS4/cfs0/shared4.vmdk

vmkfstools -c 90G -d eagerzeroedthick
/vmfs/volumes/DS5/cfs0/shared5.vmdk
```

To addd VMDK through vsphere GUI

1 Shut down the VM.

2 Select the VM and select Edit Settings

3 Select **Add New Device** > **Hard disk**. A new hard disk will be added.

4 Change the default size to 90 GB and disk provisioning to Thick Provision Eager zeroed. Select SCSI Controller 1. If not present, add a new SCSI Controller.

5 Repeat steps 3 and 4 for the rest of the disks.

Mapping the VMDKs to each virtual machine (VM)

Map each of the created VMDK files to each VM. The example procedure illustrates mapping the VMDKs to the cfs01 node: all steps should be followed for each of the other nodes.

To map the VMDKs to each VM

1 Shut down the VM.

2 Select the VM and select **Edit Settings**.

3 Ensure there is a new SCSI controller present (SCSI 1:0), if not add a new scsi controller and modify type to paravirtual if not default, and check that SCSI Bus Sharing is set to None, as this is key to allow vMotion for the VMs.

4 Select **Add** , select **Hard disk** and click **Next**.

5 Select **Use an existing virtual disk** and click **Next**.

6 Select **Browse** and choose DS1 data store.

7 Select the folder **cfs0** and select **shared1.vmdk file** and click **Next**.

8 On **Virtual Device Node**, select **SCSI (1:0)**

and click **Next**.

9 Review the details to verify they are correct and click **Finish**.

10 Since this is the first disk added under SCSI controller 1, a new SCSI controller is added.

Modify the type to **Paravirtual**, if that is not the default, and check that **SCSI Bus Sharing** is set to **None**, as this is key to allow vMotion for the VMs.

11 Follow steps 3 to 8 for the rest of the disks that will be added to each of the VMs.

For the example configuration, the parameters for steps 5-7 are given in the table below:

Data Store	VMDK Name	Virtual Device
DS1	cfs0/shared1.vmdk	SCSI 1:0
DS2	cfs0/shared2.vmdk	SCSI 1:1
DS3	cfs0/shared3.vmdk	SCSI 1:2
DS4	cfs0/shared4.vmdk	SCSI 1:3
DS5	cfs0/shared5.vmdk	SCSI 1:4

The final configuration for the first node of the example cluster (cfs01):

Now follow the same steps for each node of the cluster and map each VMDK file to the VM following the instructions above. Once all the steps are completed, all the VMs should have access to the same VMDK files. Note that at this point, all the VMs are still powered off and that multi-writer flag has not been enabled yet (it will be done in the next step). Any attempt to power on the VMs in this state will prevent a second VM start because it will violate the restrictions to access a VMDK by only a host at a time.

Enabling the multi-write flag

Enable the multi-write flag if you plan to setup an InfoScale Enterprise configuration. For detailed instructions on how to enable the multi-writer flag, see the steps in the following VMware article:

http://kb.vmware.com/kb/1034165

The steps given below illustrate the example where five VMDK files are configured and are shared by four virtual machines (VMs). These VMs constitute the four nodes of the cluster and they are powered off. Now it is time to enable the multi-writer flag for each of the VMs.

To enable the multi-write flag for a virtual machine

1 On the vSphere Client, right-click on the cfs01 virtual machine. Go to **Edit Settings** > **Advanced Parameters**.

2 .Enter **scsi1:0.sharing** in the **Attribute** column.

3 Enter **multi-writer** in the **Value** column.

4 Repeat steps 2 through 4 and enter the multi-writer value for the rest of the SCSI controllers and targets. For the example configuration:

```
scsi1:1.sharing multi-writer
scsi1:2.sharing multi-writer
scsi1:3.sharing multi-writer
scsi1:4.sharing multi-writer
```

Once those steps are done, the VM configuration will resemble:

disk.EnableUUID	true
scsi1:0.sharing	multi-writer
scsi1:1.sharing	multi-writer
scsi1:2.sharing	multi-writer
scsi1:3.sharing	multi-writer
scsi1:4.sharing	multi-writer

5 Press **OK** to confirm.

6 Repeat steps 1 to 6 for the other virtual machines (cfs02, cfs03 and cfs04 in the example configuration).

7 Once all the virtual machines have been configured correctly, power them on and verify that there are no issues. Note that the disks have been added to each of the hosts.

Example configuration for cfs01:

```
# vxdisk list
DEVICE          TYPE          DISK         GROUP        STATUS
vmdk0_0         auto:none     -            - online     invalid
vmdk0_1         auto:none     -            - online     invalid
vmdk0_2         auto:none     -            - online     invalid
vmdk0_3         auto:none     -            - online     invalid
vmdk0_4         auto:none     -            - online     invalid
vmdk0_5         auto:none     -            - online     invalid
[root@cfs01 ~]#
```

Getting consistent names across nodes

It is likely that the VMDK files are presented in a different order on each system and that the names given by Volume Manager may vary. The recommended best practice for a consistent deployment is to rename the disk so the configuration is clear.

As an example of the initial discrepancies between cfs01 and cfs03, cfs01 the disk name associated to device ending on serial number 226 is vmdk0_5:

```
[root@cfs01 ~]# /etc/vx/bin/vxgetdmpnames
enclosure vendor=VMware product=Virtual disk serial=OTHER_DISKS name=other_di
        dmpnode serial= name=sda
enclosure vendor=VMware product=Virtual disk serial=vmdk name=vmdk0
        dmpnode serial=6000C290CF1309E963CBBAD021F264F6 name=vmdk0_1
        dmpnode serial=6000C29186006C35507FEC737F9F8FDA name=vmdk0_2
        dmpnode serial=6000C291C4629CE5694FC723353138C3 name=vmdk0_3
        dmpnode serial=6000C2935F2B20B10FD8FC211FCB0936 name=vmdk0_4
        dmpnode serial=6000C29FAD20B38E6595D8BC167D27F2 name=vmdk0_5
```

Observe how cfs03 named the same device vmdk_0_0:

```
[root@cfs03 ~]#  /etc/vx/bin/vxgetdmpnames
enclosure vendor=VMware product=Virtual disk serial=OTHER_DISKS name=other_di
        dmpnode serial= name=sda
enclosure vendor=VMware product=Virtual disk serial=vmdk name=vmdk0
        dmpnode serial=6000C290CF1309E963CBBAD021F264F6 name=vmdk0_1
        dmpnode serial=6000C291C4629CE5694FC723353138C3 name=vmdk0_2
        dmpnode serial=6000C29FAD20B38E6595D8BC167D27F2 name=vmdk0_3
        dmpnode serial=6000C29186006C35507FEC737F9F8FDA name=vmdk0_4
        dmpnode serial=6000C29F83E347A4B9BC350750579889 name=vmdk0_5
```

```
[root@cfs01 ~]#  vxddladm assign names
[root@cfs01 ~]#  /etc/vx/bin/vxgetdmpnames
enclosure vendor=VMware product=Virtual disk serial=OTHER_DISKS name=other_di
        dmpnode serial= name=sda
enclosure vendor=VMware product=Virtual disk serial=vmdk name=vmdk0
        dmpnode serial=6000C291595D0A4755FBF37872E105FD name=vmdk0_1
        dmpnode serial=6000C29186006C35507FEC737F9F8FDA name=vmdk0_2
        dmpnode serial=6000C291C4629CE5694FC723353138C3 name=vmdk0_3
        dmpnode serial=6000C2935F2B20B10FD8FC211FCB0936 name=vmdk0_4
        dmpnode serial=6000C29F83E347A4B9BC350750579889 name=vmdk0_5
```

In order to get the same names across all the cluster nodes the command `vxddladm` is used. For each node of the cluster, run the command:

```
# vxddladm assign names
```

Observe now how cfs03 got the right name for device ending at 226 serial number:

```
[root@cfs03 ~]# vxddladm assign names
[root@cfs03 ~]#  /etc/vx/bin/vxgetdmpnames
enclosure vendor=VMware product=Virtual disk serial=OTHER_DISKS name=other_di
        dmpnode serial= name=sda
enclosure vendor=VMware product=Virtual disk serial=vmdk name=vmdk0
        dmpnode serial=6000C291595D0A4755FBF37872E105FD name=vmdk0_1
        dmpnode serial=6000C29186006C35507FEC737F9F8FDA name=vmdk0_2
        dmpnode serial=6000C291C4629CE5694FC723353138C3 name=vmdk0_3
        dmpnode serial=6000C2935F2B20B10FD8FC211FCB0936 name=vmdk0_4
        dmpnode serial=6000C29F83E347A4B9BC350750579889 name=vmdk0_5
```

Creating a file system

The next step will be to configure a common mount point across all the nodes, mounted on the same storage. In order to simplify the examples given here, a single disk group containing all the disks and a single volume will be created. Depending on the application requirements the number of disk groups and volumes may vary.

The boot disk has been excluded from Volume Manger configuration, so the 5 available disks (vmdk0_1, vmdk0_2, vmdk0_3, vmdk0_4 and vmdk0_5) will be the ones added to the disk group. These are the steps:

To create a clustered file system

1. Initialize the disks:
2. Create a new disk group and add the disks.
3. Verify the configuration. Note the DISK and GROUP information.
4. Create a striped volume with the 5 disks available.
5. Create a File System.
6. If you plan to configure a clustered file system environment, add the newly created file system to the cluster configuration. Given that this will be mounted by all the nodes at the same time, we will add it as a cluster resource, and commands cfsmntadm and cfsmount will be used.

7 In case of a clustered file system, verify that the new directory is available in all the nodes by running the cfscluster status command or by verifying with df in each of the nodes.

Section 3

Use cases for InfoScale product components in a VMware environment

Chapter 4

Application availability using Cluster Server

This chapter includes the following topics:

About application availability with Cluster Server (VCS) in the guest

Using VCS virtual-to-virtual or in-guest clustering in a VMware environments provides high availability of applications inside the guest by providing protection from host failures, hardware failures, OS crashes and also application failures at software layer. For example, in cases of application hang, file-level corruption at the OS level cannot be resolved with a reboot.

Since there is a cost involved in maintaining standby virtual machines (VMs), you may choose to protect only specific applications using VCS in-guest and protect the remaining applications using VMware HA. By using VMware HA in conjunction with VCS in the guest, when a host fails, standby VCS nodes running on that host are automatically restarted by VMware HA on a new host without the need for user-intervention, potentially eliminating the need to maintain multiple standbys.

About VCS support for Live Migration

VCS in-guest clustering continues to provide high availability of applications on virtual machines, in live migration scenarios initiated by the virtualization technology. You can use Live migration to perform a stateful migration of an

virtual machine in a VCS environment. During this period, you may see notifications if the migrating node is unable to heartbeat with its peers within LLT's default peer inactive timeout. To avoid false failovers, determine how long the migrating node is unresponsive in your environment. If that time is less than the default LLT peer inactive timeout of 16 seconds, VCS operates normally. If not, increase the peer inactive timeout to an appropriate value on all the nodes in the cluster before beginning the migration. Reset the value back to the default after the migration is complete.

Chapter 5

Multi-tier business service support

This chapter includes the following topics:

About Virtual Business Services

The Virtual Business Services feature provides visualization, orchestration, and reduced frequency and duration of service disruptions for multi-tier business applications running on heterogeneous operating systems and virtualization technologies. A virtual business service represents the multi-tier application as a consolidated entity that helps you manage operations for a business service. It builds on the high availability and disaster recovery provided for the individual tiers by InfoScale™ products such as Cluster Server.

Application components that are managed by Cluster Server or Microsoft Failover Clustering can be actively managed through a virtual business service.

You can use the InfoScale Operations Manager Management Server console to create, configure, and manage virtual business services.

Sample virtual business service configuration

This section provides a sample virtual business service configuration comprising a multi-tier application. Figure 5-1 shows a Finance application that is dependent on components that run on three different operating systems and on three different clusters.

- Databases such as Oracle running on Solaris operating systems form the database tier.
- Middleware applications such as WebSphere running on AIX operating systems form the middle tier.
- Web applications such as Apache and IIS running on Windows and Linux virtual machines form the Web tier.
  Each tier can have its own high availability mechanism. For example, you can use Cluster Server for the databases and middleware applications for the Web servers.

Figure 5-1 Sample virtual business service configuration

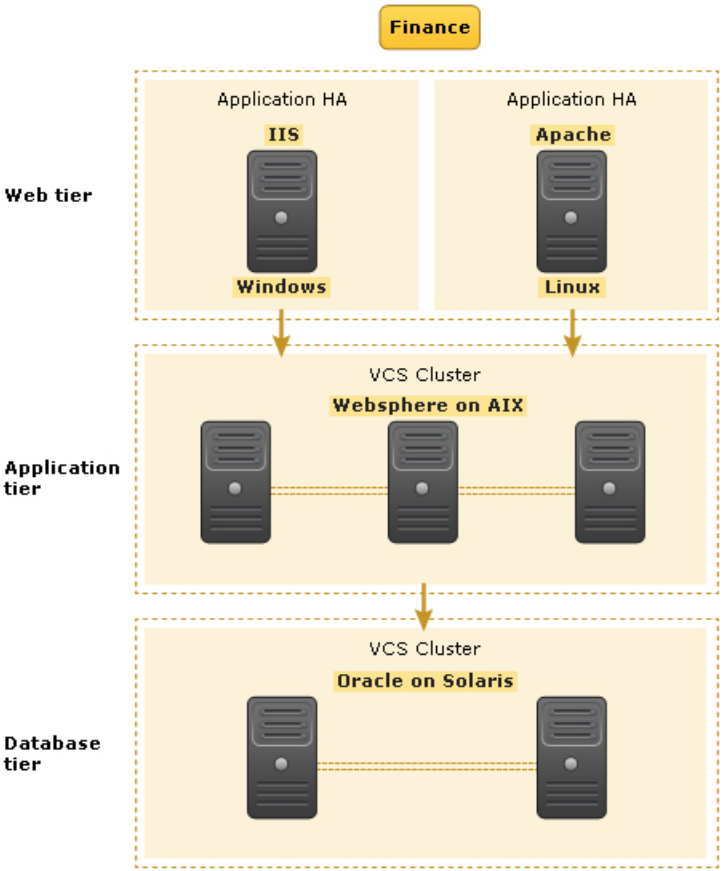

Each time you start the Finance business application, typically you need to bring the components online in the following order – Oracle database, WebSphere, Apache and IIS. In addition, you must bring the virtual machines online before you start the Web tier. To stop the Finance application, you must take the components offline in the reverse order. From the business perspective, the Finance service is unavailable if any of the tiers becomes unavailable.

When you configure the Finance application as a virtual business service, you can specify that the Oracle database must start first, followed by WebSphere and the Web servers. The reverse order automatically applies when you stop the virtual business service. When you start or stop the virtual business service, the components of the service are started or stopped in the defined order.

For more information about Virtual Business Services, refer to the *Virtual Business Service–Availability User's Guide*.

Chapter 6

Improving data protection, storage optimization, data migration, and database performance

This chapter includes the following topics:

Use cases for InfoScale product components in a VMware guest

InfoScale product components provide many features for enhancing storage in a VMware environment. You can use the InfoScale products for data protection, storage optimization, ease of data migration, and optimized performance.

Using the InfoScale product in VMware guest enables the following use cases:

Table 6-1 InfoScale support for VMware environment use cases

Virtualization use case	InfoScale™ product component	Implementation details
Data protection	InfoScale or SFHA in the guest	How to use InfoScale point-in-time copy features for data protection. See "Protecting data with InfoScale product components in the VMware guest" on page 44.
Storage optimization	InfoScale or SFHA in the guest	How to use InfoScale thin provisioning, SmartTier, and SmartIO features to optimize storage in a VMware environment. See "Optimizing storage with InfoScale product components in the VMware guest" on page 46.
Data migration	InfoScale or SFHA in the guest	How to use InfoScale Portable Data Containers to migrate data safely and easily in a VMware guest environment. See "Migrating data with InfoScale product components in the VMware guest" on page 51.
Improved database performance	InfoScale or SFHA in the guest	How to use InfoScale database accelerators to improve database performance in a VMware guest environment. See "Improving database performance with InfoScale product components in the VMware guest" on page 51.

Protecting data with InfoScale product components in the VMware guest

When you install InfoScale in the VMware guest, you can use the following InfoScale product point-in-time copy technologies to protect your data:

Table 6-2 Data protection options in InfoScale product components

Data protection feature	Description
FlashSnap	Volume-based point-in-time snapshot and recovery method.

Table 6-2 Data protection options in InfoScale product components *(continued)*

Data protection feature	Description
Database FlashSnap	Not supported
Storage Checkpoints	File system-based point-in-time copy and recovery method.
Database Storage Checkpoints	Not supported
FileSnap	File-level point-in-time snapshot and recovery method.
Volume replication	Ongoing method of volume replication to a remote location.
File replication	Ongoing method of file system replication to a remote location.
DST cloud tearing	Copies the files from tier1-tier2 (InfoScale cluster to s3 and vise versa).

About point-in-time copies

Storage Foundation offers a flexible and efficient means of managing business-critical data. Storage Foundation lets you capture an online image of an actively changing database at a given instant, called a point-in-time copy.

More and more, the expectation is that the data must be continuously available (24x7) for transaction processing, decision making, intellectual property creation, and so forth. Protecting the data from loss or destruction is also increasingly important. Formerly, data was taken out of service so that the data did not change while data backups occured; however, this option does not meet the need for minimal down time.

A point-in-time copy enables you to maximize the online availability of the data. You can perform system backup, upgrade, or perform other maintenance tasks on the point-in-time copies. The point-in-time copies can be processed on the same host as the active data, or a different host. If required, you can offload processing of the point-in-time copies onto another host to avoid contention for system resources on your production server. This method is called off-host processing. If implemented correctly, off-host processing solutions have almost no impact on the performance of the primary production system.

Point-in-time snapshots for InfoScale products in the VMware environment

The point-in-time snapshots are fully supported for InfoScale products in the VMware environment. In addition to using point-in-time snapshots to create

back-ups or perform off-host processing, you can use them to provision additional virtual machines.

For more information about point-in-time snapshots, see *InfoScale Solutions Guide*.

Optimizing storage with InfoScale product components in the VMware guest

When you install InfoScale or Storage Foundation High Availability in the VMware guest, you can use the following InfoScale storage optimization technologies to maximize your storage utilization:

Table 6-3 Storage optimization options in InfoScale product components

Storage optimization feature	Description
Compression	Maximize storage utilization by reducing file sizes.
Thin Reclamation	Optimize your thin array usage by setting set up thin storage.
SmartMove	Optimize your thin array usage with thin reclamation solutions to maintain thin storage.
SmartTier	Maximize storage efficiency by moving data to storage tiers based on age, priority, and access rate criteria.
SmartTier for Oracle	Oracle-optimized method to maximize storage efficiency by moving data to storage tiers based on age, priority, and access rate criteria.

About Flexible Storage Sharing

Flexible Storage Sharing (FSS) enables network sharing of local storage, cluster wide. The local storage can be in the form of Direct Attached Storage (DAS) or internal disk drives. Network shared storage is enabled by using a network interconnect between the nodes of a cluster.

FSS allows network shared storage to co-exist with physically shared storage, and logical volumes can be created using both types of storage creating a common storage namespace. Logical volumes using network shared storage provide data redundancy, high availability, and disaster recovery capabilities, without requiring physically shared storage, transparently to file systems and applications.

FSS allows growing the existing storage by growing the LUN, which is termed as Dynamic LUN Expansion (DLE). To use the DLE feature, you invoke the `vxdisk resize` command with the `length` option from the master node in an FSS

configuration. Alternatively, you can use the Management Server console of InfoScale Operations Manager to resize a disk, which internally invokes this command. The command intelligently detects the remote disks and gets the required protocol executed to complete the LUN expansion; you do not need to specify any additional options. No explicit master switching is required.

Note: Do not resize an LVM disk or any other disk on which the OS is installed.

For details on growing the existing storage by growing the LUN, refer to the platform-specific *Storage Foundation Administrator's Guide*.

FSS can be used with SmartIO technology for remote caching to service nodes that may not have local SSDs.

FSS is supported on clusters containing up to 64 nodes with CVM protocol versions 140 and above.

The following figure depicts a FSS environment.

Figure 6-1 Flexible Storage Sharing environment

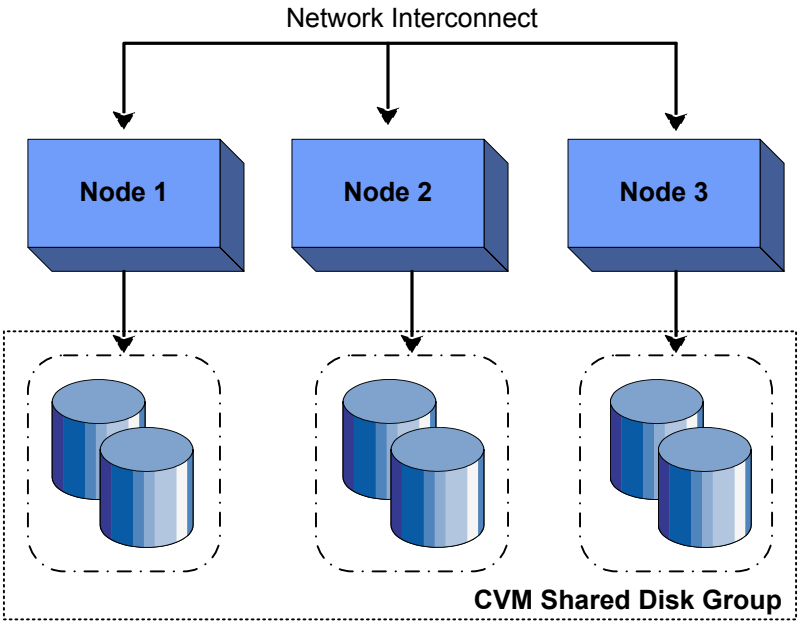

Limitations of Flexible Storage Sharing

Note the following limitations for using Flexible Storage Sharing (FSS):

- Disk initialization operations should be performed only on nodes with local connectivity to the disk.
- FSS does not support the use of boot disks, opaque disks, and non-VxVM disks for network sharing.
- Hot-relocation is disabled on FSS disk groups.
- FSS does not support non-SCSI3 disks connected to multiple hosts.
- FSS supports only those instant data change objects (DCOs) that are either created using the `vxsnap` operation or by specifying the **`logtype=dco`** and **`dcoversion=20`** attributes during volume creation.
- By default, creating a mirror between SSD and HDD is not supported through `vxassist`, as the underlying media types are different. To workaround this issue, you can create a volume with one media type, for instance the HDD, which is the default media type, and then later add a mirror on the SSD. For example:

```
# vxassist -g diskgroup make volume size init=none
# vxassist -g diskgroup mirror volume mediatype:ssd
# vxvol -g diskgroup init active volume
```

About SmartTier in the VMware environment

SmartTier is a VxFS feature that enables you to allocate file storage space from different storage tiers according to rules you create. SmartTier provides a more flexible alternative compared to current approaches for tiered storage. Static storage tiering involves a manual one-time assignment of application files to a storage class, which is inflexible over a long term. Hierarchical Storage Management solutions typically require files to be migrated back into a file system name space before an application access request can be fulfilled, leading to latency and run-time overhead. In contrast, SmartTier allows organizations to:

- Optimize storage assets by dynamically moving a file to its optimal storage tier as the value of the file changes over time
- Automate the movement of data between storage tiers without changing the way users or applications access the files
- Migrate data automatically based on policies set up by administrators, eliminating operational requirements for tiered storage and downtime commonly associated with data movement

---

Note: SmartTier is the expanded and renamed feature previously known as Dynamic Storage Tiering (DST).

You can use the SmartTier functionality of InfoScale in a VMware environment similar to how you use it in the physical environment. For example, you may have two different types of storage arrays, some SSD and some HDD. You can allocate one VMFS on each type of disk. You can use the SSD devices as one storage tier and the HDD devices as the second storage tier.

For more information about SmartTier, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About compression with InfoScale product components in the VMware guest

Compressing files reduces the space used by files, while retaining the accessibility of the files and being transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files. Reads cause data to be uncompressed in memory, only; the on-disk copy of the file remains compressed. In contrast, after a write, the new data is uncompressed on disk.

The compression functionality is the same in the VMware environment as in the physical devices.

For more information about file compression, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About thin reclamation with InfoScale product components in the VMware guest

InfoScale product components support reclamation of the unused storage on thin-reclamation capable arrays and LUNs. InfoScale Solutions automatically discover LUNs that support thin reclamation.

If you configure InfoScale product components in Raw Device Mapping mode, InfoScale product components can control the storage in the VMware guests directly. With this configuration, you can use the thin reclamation functionality of InfoScale product components on the VMware guest. This configuration may limit some VMware functionality such as vMotion. You cannot use the thin reclamation functionality of InfoScale product components if the storage is configured in VMDK mode.InfoScale product components do not have the visibility to the physical storage under the VMware guest to reclaim the freed storage.

For more information about thin reclamation, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About SmartMove with InfoScale product components in the VMware guest

The SmartMove utility optimizes move and copy operations. The SmartMove utility leverages the knowledge that File System (VxFS) has of the Volume Manager (VxVM) storage. VxFS lets VxVM know which blocks have data. When VxVM performs an operation that copies or moves data, SmartMove enables the operation to only copy or move the blocks used by the file system.

This capability improves performance for synchronization, mirroring, and copying operations because it reduces the number of blocks that are copied. SmartMove only works with VxFS file systems that are mounted on VxVM volumes. If a file system is not mounted, the utility has no visibility into the usage on the file system.

The SmartMove functionality is the same in the VMware environment as in the physical devices.

For more information about SmartMove, see *Storage Foundation Cluster File System High Availability Administrator's Guide* or *Storage Foundation Administrator's Guide*.

About SmartTier for Oracle with InfoScale product components in the VMware guest

SmartTier is a VxFS feature that enables you to allocate file storage space from different storage tiers according to rules you create. SmartTier provides a more flexible alternative compared to current approaches for tiered storage.

In an Oracle database environment, the access age rule can be applied to archivelog files and Flashback files. Oracle updates the header of each datafile at every database checkpoint and hence access age rules cannot be used for datafiles. For a partitioned table, we can use the name base rule to relocate files belonging to a given partition, for instance last year, to the secondary storage tier. However if a database does not have partitioned tables, current methods for relocation do not fit to the Oracle database environment.

To adopt a successful SmartTier policy and save storage costs, a method is needed for relocating Oracle objects which are accessed infrequently to the secondary tier. Relocating the entire file is not useful in all cases.

SmartTier for Oracle in the VMware environment works the same as in the physical environment.

For more information about SmartTier for Oracle, see *InfoScale Storage and Availability Management for Oracle Databases*.

Migrating data with InfoScale product components in the VMware guest

When you install the InfoScale products, InfoScale Storage or InfoScale Enterprise in the VMware guest, you can use the following InfoScale data migration technologies to migrate your data safely and easily:

- Portable Data Containers
- Volume mirroring

Types of data migration

This section describes the following types of data migration:

- Migrating data between platforms using Cross-platform Data Sharing (CDS)
  Storage Foundation lets you create disks and volumes so that the data can be read by systems running different operating systems. CDS disks and volumes cannot be mounted and accessed from different operating systems at the same time. The CDS functionality provides an easy way to migrate data between one system and another system running a different operating system.
- Migrating data between arrays
  Storage Foundation supports arrays from various vendors. If your storage needs change, you can move your data between arrays.

Note: The procedures are different if you plan to migrate to a thin array from a thick array.

Improving database performance with InfoScale product components in the VMware guest

When you install the InfoScale products in the VMware guest, you can use the following InfoScale database accelerator technologies to enhance your database performance:

- Veritas Oracle Disk Manager (VxODM)
- Veritas Cached Oracle Disk Manager
- Veritas Concurrent I/O

Database accelerators enable your database to achieve the speed of raw disk while retaining the management features and convenience of a file system.

About InfoScale product components database accelerators

The major concern in any environment is maintaining respectable performance or meeting performance service level agreements (SLAs). InfoScale product components improve the overall performance of database environments in a variety of ways.

Table 6-4 InfoScale product components database accelerators

InfoScale database accelerator	Supported databases	Use cases and considerations
Oracle Disk Manager (ODM)	Oracle	■ To improve Oracle performance and manage system bandwidth through an improved Application Programming Interface (API) that contains advanced kernel support for file I/O. ■ To use Oracle Resilvering and turn off Volume Manager (VxVM) Dirty Region Logging (DRL) to increase performance, use ODM. ■ To reduce the time required to restore consistency, freeing more I/O bandwidth for business-critical applications, use SmartSync recovery accelerator.
Cached Oracle Disk Manager (Cached ODM)	Oracle	To enable selected I/O to use caching to improve ODM I/O performance, use Cached ODM.
Concurrent I/O	DB2 Sybase	Concurrent I/O (CIO) is optimized for DB2 and Sybase environments To achieve improved performance for databases run on VxFS file systems without restrictions on increasing file size, use VxFS Concurrent I/O.

These database accelerator technologies enable database performance equal to raw disk partitions, but with the manageability benefits of a file system.

InfoScale database accelerators enable you to manage performance for your database with more precision.

Chapter 7

Setting up virtual machines for fast failover using InfoScale Enterprise on VMware disks

This chapter includes the following topics:

About use cases for InfoScale Enterprise in the VMware guest

In addition to the application availability and storage capabilities, InfoScale Enterprise adds the fast failover capability of a parallel cluster.

InfoScale Enterprise operation in VMware virtualized environments

InfoScale Enterprise has two supported modes of operation when used inside a guest operating system operating system (OS) that is running on a VMware hypervisor:

- Mode 1: InfoScale Enterprise is connected to external storage via RDM-P
  Use InfoScale Enterprise running in a guest OS and connected to external storage via RDM-P when you require highly reliable SCSI-3 PGR keys for split-brain protection and data fencing
- Mode 2: InfoScale Enterprise connected to VMFS with the VMFS multi-writer flag enabled
  Use InfoScale Enterprise running in a guest OS and connected to VMFS with the VMFS multi-writer flag enabled when you require guest OS mobility via VMware vMotion.

Requirements for Mode 1:

- InfoScale Enterprise must be connected to a physical LUN has been mapped to a virtual machine disk (VMDK) using the VMware raw disk mapping physical (RDMP) compatibility mode.
- This option provides full InfoScale Enterprise high availability support for fast failover, split brain protection, and data fencing, but does not allow use of VMware snapshots, vMotion, or other VMware HA features.
- InfoScale Enterprise snapshots, clones, and other storage optimization features are fully supported in this configuration.
- vmotion should work with RDM-P where the virtual SCSI controller is set to have SCSI Bus sharing enabled.
  https://www.veritas.com/content/support/en_US/doc/infoscale_support_vmware_vmotion_rdmp

Requirements for Mode 2:

- The guest OS must use a VMFS virtual machine disk (VMDK).
- VMFS normally prohibits multiple guest Operating Systems from connecting to the same virtual disk, thus precluding use of that VMDK with a parallel file system such as. However, by enabling a new VMFS option, multi-writer, an administrator can create a VMDK that can be read/written to by multiple guest Operating Systems simultaneously. This is described in the following VMware article - https://kb.vmware.com/s/article/1034165
- Since VMFS does not allow InfoScale Enterprise to see the SCSI-3 persistent group reservation (PGR) bit on the LUN, SCSI-based I/O fencing will not work.

To use InfoScale Enterprise with the VMFS multi-writer option, Coordination Point Servers (CPS) must be used for split-brain protection.

For information on configuring CPS fencing, see the Storage Foundation Cluster File System High Availability Administrator's Guide.

- CPS does not support data protection via fencing, so care must be taken to prevent inadvertent data corruption caused by a non-Arctera InfoScale Enterprise guest OS writing to a VMDK.

  InfoScale Enterprise snapshots, clones, and other storage optimization features are fully supported in this configuration.

Support for InfoScale Enterprise in a guest OS:

- Arctera will attempt to resolve any issues as if they were occurring in a physical OS environment. Once a potential problem has been identified, Arctera support personnel will recommend an appropriate solution that would be applicable on the native OS on a physical server.
- If that solution does not work in the VMware virtualized environment, Arctera reserves the right to ask the customer to replicate the problem in a physical environment. In some cases the customer may be referred to VMware for support.

InfoScale functionality and compatibility matrix

Table 7-1 shows the InfoScale functionality and compatibility with VMware ESXi disk modes.

Table 7-1 InfoScale functionality and compatibility matrix with VMware disk modes

InfoScale	VMware ESXi disk mode: Virtual Disk (VMDK)	VMware ESXi disk mode: Raw Device Mapping Logical mode	VMware ESXi disk mode: Raw Device Mapping Physical mode
VxVM Disk format: simple, sliced	Yes	Yes	Yes
VxVM Disk format: cdsdisk	Yes	Yes	Yes
I/O fencing	Yes (with non-SCSI3-PR based fencing)[1]	Yes (with non-SCSI3-PR based fencing) [1]	Yes (with disks configured in RDM-P mode) [1]

Table 7-1 InfoScale functionality and compatibility matrix with VMware disk modes *(continued)*

InfoScale	VMware ESXi disk mode: Virtual Disk (VMDK)	VMware ESXi disk mode: Raw Device Mapping Logical mode	VMware ESXi disk mode: Raw Device Mapping Physical mode
Portable Data Containers	No	No	Yes
Volume Replicator	Yes	Yes	Yes
CVM/VVR	Yes	Yes	Yes
Bunker node (non-CVM environment)	Yes	Yes	Yes
DDL extended attributes	No	No	Yes
Thin reclamation	No	No	Yes

See "I/O fencing considerations in an ESXi environment" on page 16.

About setting up InfoScale Enterprise on VMware ESXi

This sample deployment illustrates how to install and configure InfoScale Enterprise in a VMware virtual server using VMware filesystem (VMFS) virtual disks (VMDKs) as the storage subsystem.

The information provided here is not a replacement or substitute for InfoScale Enterprise documentation nor for the VMware documentation. This is a deployment illustration which complements the information found in other documents.

The following product versions and architecture are used in this example deployment:

- RedHat Enterprise Linux (RHEL) Server 8.6
- InfoScale Enterprise 9.0
- ESXi 8.0

A four node virtual machine cluster will be configured on two VMware ESXi servers. Shared storage between the two ESXi servers using Fibre Channel has

been setup. The Cluster File System will exist across four virtual machines: cfs01, cfs02, cfs03, and cfs04. Three Coordination Point (CP) servers will be used: cps1, cps2, and cps3 (this one placed in a different ESXi server). For storage, five data stores will be used and one shared VMDK file will be placed in each data store.

Figure 7-1 InfoScale Enterprise virtual cluster on VMware ESXi

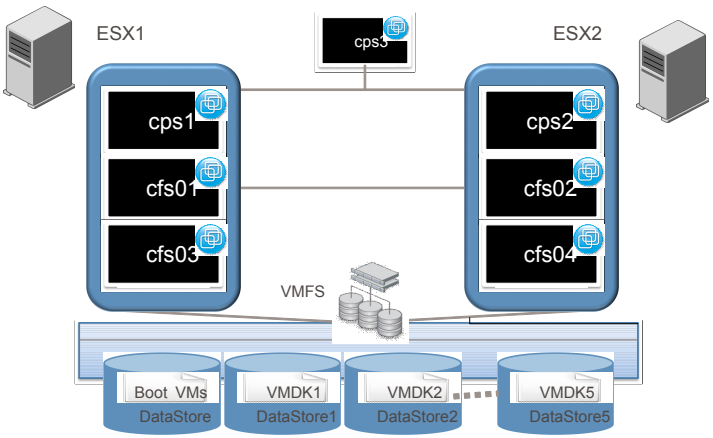

Two private networks will be used for cluster heartbeat. They are called PRIV1 and PRIV2. Virtual switch vSwitch2 also has the VMkernel Port for vMotion enabled. vSwitch0 is used for management traffic and the public IP network.

Some blades have a two network limit. If this is the case, configure one network for heartbeats and the other one as a heartbeat backup (low priority setting).

Planning a InfoScale Enterprise configuration

The deployment example illustrates the recommended steps for installing and configuring InfoScale Enterprise 9.0 using CP servers. To achieve a successful InfoScale Enterprise implementation on VMware, follow the recommended steps. Different versions of RHEL or ESXi may require different steps.

See the InfoScale Enterprise documentation for additional information.

Table 7-2 Tasks for deploying InfoScale Enterprise on VMware

InfoScale Enterprise deployment task	Deployment steps
Deploy virtual machines (VMs)	See RHEL and ESXi documentation.
Enable password-less SSH	See "Enable Password-less SSH" on page 59.

Table 7-2 Tasks for deploying InfoScale Enterprise on VMware *(continued)*

InfoScale Enterprise deployment task	Deployment steps
Enable traffic to CP server and management ports	See "Enabling TCP traffic to coordination point (CP) Server and management ports" on page 60.
Configure CP servers	See "Enabling TCP traffic to coordination point (CP) Server and management ports" on page 60.
Deploy InfoScale Enterprise software	See "Deploying InfoScale Enterprise software" on page 64.
Configure InfoScale Enterprise	See "Configuring InfoScale Enterprise" on page 65.
Configure non-SCSI3 fencing	See "Configuring non-SCSI3 fencing" on page 67.

Enable Password-less SSH

The installer will configure password-less SSH/RSH among the cluster nodes, but it will not be able to enable this required functionality between the cluster nodes and the CP servers. This is necessary if there are any future changes in the IO fencing configurations, etc. In this configuration we are going to configure password-less SSH between the node that we will be using to configure the cluster and the rest of the nodes.

The following are the instances where Password-less SSH will be enabled:

Table 7-3 Instances to enable password-less SSH

Source	Target	Reason
cfs01	cfs02	Cluster configuration
cfs01	cfs03	Cluster configuration
cfs01	cfs04	Cluster configuration
cfs01	cps01	Non-SCSI-3 fencing configuration
cfs01	cps2	Non-SCSI-3 fencing configuration
cfs01	cps3	Non-SCSI-3 fencing configuration

Table 7-3 Instances to enable password-less SSH *(continued)*

Source	Target	Reason
cps01	cps01	Secure cluster configuration
cps2	cps01	Secure cluster configuration
cps3	cps01	Secure cluster configuration

If you do not enable Password-less SSH, then follow the manual configuration instructions given at the end of this document.

Enabling TCP traffic to coordination point (CP) Server and management ports

For successful intra-cluster communication, make sure that cluster nodes and CP servers can be reached on port 14250 (or any other if you changed the default). If RedHat Firewall has been enabled, make sure there is a rule to allow the connection to ports 14250 and 14149.

To enable TCP traffic to CP server and management ports

1 Stop iptables service:

```
[root@cps3 sysconfig]# service iptables stop
iptables: Flushing firewall rules:                         [ OK ]
iptables: Setting chains to policy ACCEPT: filter          [ OK ]
iptables: Unloading modules:                               [ OK ]
```

2 Enter the following lines at the /etc/sysconfig/iptables file:

```
-A INPUT -p tcp -m tcp --dport 14250 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 14149 -j ACCEPT
```

3 Start the service:

```
[root@cps2 ~]# service iptables restart
iptables: Flushing firewall rules:                         [ OK ]
iptables: Setting chains to policy ACCEPT: filter          [ OK ]
iptables: Unloading modules:                               [ OK ]
iptables: Applying firewall rules:                         [ OK ]
[root@cps2 ~]#
```

4 Verify the new rule is in place:

```
[root@cps3 sysconfig]# iptables --list

Chain INPUT (policy ACCEPT)
target     prot opt source               destination
ACCEPT     tcp  -- anywhere             anywhere            tcp dpt:cps
ACCEPT     tcp  -- anywhere             anywhere            tcp dpt:vrts-tdd
ACCEPT     tcp  -- anywhere             anywhere            tcp dpt:xprtld
```

For the InfoScale Enterprise deployment example, these rules must be enabled on cfs01, cfs02, cfs03, cfs04, cps1, cps2 and cps3

Configuring coordination point (CP) servers

In order to protect the cluster from data corruption caused by a cluster failure that results in "split-brain," coordination point (CP) servers must be deployed. If you do not have CP servers configured already, you must configure the CP servers that are used to provide non-SCSI-3 fencing protection at this point in the deployment process.

CP servers can be used for multiple clusters, so this step will only be necessary the first time a cluster is configured and no CP servers are available.

Tasks for configuring a new CP server:

Table 7-4

Step	Description
Configure a new Coordination Point server for InfoScale Enterprise	See "Configuring a Coordination Point server for InfoScale Enterprise" on page 62.
Configure a Cluster Server (VCS) single node cluster	See "Configuring a Cluster Server (VCS) single node cluster" on page 62.
Configure a Coordination Point server service group	See "Configuring a Coordination Point server service group" on page 63.

Configuring a Coordination Point server for InfoScale Enterprise

In the example InfoScale Enterprise deployment, three CP servers are used. Each CP server will be housed in a virtual machine, with each one forming a single node CP server cluster. Each of the two physical ESX Servers will contain one CP server, and a third CP server will be located in another location. To illustrate the process, steps are given to deploy the cps1 node.

The installer script will be used to perform the deployment of Cluster Server (option 2). If required packages are not present on the nodes, they will be displayed during the Pre-Installation check. Please make sure to install the rpms on the nodes mentioned.

Configuring a Cluster Server (VCS) single node cluster

Configure a VCS single node cluster for the Coordination Point server. The following precedure illustrates the process for the example configuration.

To configure a single-node cluster

1. Run the installer script and perform a Pre-Installation Check to verify that the server is ready to have VCS installed.
2. Select the product installation option, and when prompted to make a selection choose VCS.
3. Select the option to install all RPMs. The VRTScps package is included.
4. Enter the name of the host where you are running the configuration, cps1 in this example. After reviewing all the packages that will be installed, the package installation begins. In the example environment, the keyless licensing is required because a InfoScale Operation Manager host is deployed. Global Cluster Option will not be enabled.
5. When prompted to configure VCS, enter **y**.
6. When prompted to configure fencing, enter **n**. Fencing is not needed for a single-node cluster.
7. Although configuring a single node, you must enable LLT and GAB for Coordination Point Server (CP server) clients to connect to the CP server. When prompted to start GAB and LLT, choose **y**.
8. Enter the cluster name (same as node). For the configuration example, it is cps1.
9. Select **Configure heartbeat links using LLT over Ethernet**.

10 In the configuration example, an additional NIC must be configured for the private network. This step is necessary even with a single node cluster.

Enter a unique cluster ID that is not already in use. The installer verifies that the cluster ID is not already in use.

11 You do not need to enter a Virtual IP, as the same one used for the host will be fine to use.

12 You must configure secure mode or not according to your own configuration requirements. For the example configuration, InfoScale Enterprise is configured to use secure mode, so the Coordination Point Servers must also use secure mode.

13 The FIPS option has not been qualified yet with CP servers. When prompted to choose, select **Configure the cluster in secure node without fips**.

This also matches the configuration done for the cluster nodes.

14 For the example configuration, SMTP and SNMP are not used. For your own configuration you must choose based on your requirements.

15 The VCS configuration starts after the selections are completed. Once finished, you can verify if VCS is running.

```
# hastatus -sum

-- SYSTEM STATE
-- System               State                Frozen

A  cps1                 Running              0
```

Configuring a Coordination Point server service group

Even in a single node cluster, a virtual IP address (VIP) is used. This enables the creation of a VCS resource to control the VIP availability. For the example configuration, a VIP for each CP server is assigned to illustrate the process.

To configure CP server service group

1 Verify that you have a VIP available for each of your CP servers.

2 Run the command:

```
# /opt/VRTS/install/installer <version> -configcps
```

Where <*version*> is the specific release version

3 When the installer prompts if you want to configure a CP server, select **Configure Coordination Point Server on single node VCS system**.

4 The name of the CP server is the same as the host name plus "v" at the end. For the example configuration CP server it is cps1v.

5 Enter the Virtual IP address that is associated to the single node cluster. In the example of node cps1, it is 192.168.1.1. Accept the default port suggested.

6 As discussed before, security is enabled for the example and is recommended as a best practice.

7 When prompted, enter the location of the database. In the example, the database will be installed locally, so you can accept the default location.

8 After reviewing the configuration parameters, continue with the configuration of the CP server Service Group. The NIC used at cps1 is eth4. The example does not use NetworkHosts. Enter the netmask, and the configuration is complete.

The CPSSG Service Group is now online.

```
# hastatus -sum

-- SYSTEM STATE
-- System            State                Frozen

A  cps1              Running              0

-- GROUP STATE
-- Group          System      Probed      AutoDisabled   State


B  CPSSG          cps1        Y           N              ONLINE
```

Deploying InfoScale Enterprise software

Deploy InfoScale Enterprise in each of the four virtual machines that make up the cluster. To do that, you could have either selected **yes** in the previous CPI step, so the install could continue, or you could have run the installer script again. If you ran the installer script again, use the following procedure to deploy InfoScale Enterprise.

To deploy InfoScale Enterprise after running the installer script again

1 Make the following selections as you run the installer script:

- I) Install a Product
- 5) Arctera InfoScale Enterprise

- Agree the terms of the EULA.

2 When you are prompted to select the RPMs to install on all systems, select option 3, Install all RPMs.

3 Enter the name of the nodes of the cluster (cfs01 cfs02 cfs03 cfs04) on which to install the RPMs.

The installer verifies the pre-requisites again, and a list of all the RPMs that will be installed will be printed. The RPMs are installed.

4 Once the RPMs have been installed, the installer prompts you for the license. Select option 2 to enable keyless licensing so this cluster can be managed by InfoScale Operation Manager):

5 You do not need licenses to enable the following options for this InfoScale Enterprise deployment:

- Replication
- Global Cluster Option

Once the licenses have been registered, the installer finishes with the deployment.

You can now start the configuration. If you want a Virtual Machine template with the InfoScale Enterprise software already deployed, stop here and take a snapshot or other copy of this image. The next step will be to run "installer –configure" to continue with the configuration.

Configuring InfoScale Enterprise

To configure InfoScale Enterprise cluster settings

1 Run **installer –configure** or just continue from where we left in the previous step entering **y**.

2 Fencing would normally be the next step in configuring InfoScale Enterprise. However, the I/O fencing configuration depends on other factors which are not yet determined:

- VMDKs or RDMP storage devices are used
- How I/O and network paths are configured
- Configuration of coordination point (CP) server (or, in some cases, Coordination Disks)

For now you can enter **n** when prompted to configure IO fencing in enabled mode and come back to it later in the configuration process.

3 Configure the cluster name when prompted.

The cluster name for the example deployment is cfs0

4 Configure the NICs used for heartbeat when prompted.

LLT (Low Latency Protocol) can be configured over Ethernet or UDP. UDP is needed only when routing between the nodes is necessary. If UDP is not needed, then Ethernet is the clear recommendation.

In the example deployment, eth4 and eth5 are the private links. Eth3 is the public link, and it will be only used as low priority heartbeat path (so it only will be used if the other two paths fail).

All media speed checking should succeed. If not, please review your node interconnections.

5 Configure the cluster ID when prompted. A unique cluster ID is needed: it is vital to choose a number that is not used in any other cluster. This is especially true when using the same network interconnection (both private and public). The CPI generates a random number, and checks the network to make sure that packets with that ID do not exist. However the CPI cannot guarantee that the ID is not being used in a cluster that is currently powered off. The best practice is to maintain a register of the cluster IDs used across the data center to avoid use of duplicate IDs. In the example configuration, no other clusters with that ID have been found.

6 At this point a summary of the configuration to be deployed is presented. Examine the summary and enter **y** if everything is correct. If not enter **n** and go through the steps again.

7 The installer prompts for a Virtual IP to manage the cluster. This is not mandatory, and the cluster can be configured without that IP. Depending on your implementation, it might be a best practice.

8 Decide whether or not to use secure mode.

In the past, the difficulty in configuring Cluster Server secure mode deterred many users from using it. For InfoScale Enterprise:

- Secure mode configuration is much easier
- The installer takes care of the entire configuration
- A validated user and password from the OS is used instead of the traditional admin/password login

For demonstration purposes, secure mode is used in the example deployment, but feel free to choose the option that best suits your needs.

FIPS is not used for the example configuration as it is not certified for deployment with CP servers. Option 1, **secure mode without FIPS** is used.

9 SMTP is not needed for the example.

10 SNMP notifications are not needed for the example.

At this point the cluster configuration will be initiated.

Configuring non-SCSI3 fencing

VMDK files do not currently support SCSI-3 Persistent Reservation and therefore non-SCSI-3 PR fencing must be used. Coordination point (CP) servers provide the required level of server based fencing. At this point in the configuration process, the three CP servers that are to be used with this cluster should be available and the CP service should be up and running.

To configure non-SCSI-3 fencing

1 If you started at the beginning of the installer process and selected the enable fencing option, you are prompted to configure fencing.

If you chose not to enable fencing at that point, then the cluster configuration is finished. You should now run `/opt/VRTS/install/installer` to enable fencing in the cluster.

2 Regardless of how you navigated to the fencing configuration of the installer, select option 1 for Coordination Point client-based fencing.

3 When prompted if your storage environment supports SCSI-3 PR, select **n**, since VMDK files do not support SCSI-3 PR.

4 When prompted if you want to configure Non-SCSI-3 fencing, select **y**.

5 For production environments, three CP servers are recommended. Enter **3** when prompted for the number of coordination points.

6 Specify how many interfaces the CP servers will be listening on and the IP address of each interface. If a CP server is reachable via several networks, the best practice is to configure every interface. This allows the InfoScale Enterprise nodes maximum communication flexibility, should a race condition occur.

Enter the host names and VIPs for the other CP servers and review the fencing configuration.

7 When prompted, select secure mode. All the trusted relationships between cluster nodes and CP servers are automatically set up.

8 Verify that the cluster information is correct. Each node is registered with each of the CP servers. Once this is done, the installer will restart VCS to apply the fencing configuration. At this point we don´t have any file system configured yet.

9 When prompted, it is a recommended best practice to configure the Coordination Point Agent on the client, so CP servers are proactively monitored from the cluster. This step completes the fencing configuration.

Once fencing configuration is complete, you can verify if it is correct.

To verify the fencing configuration

1 Query each of the CP servers to verify each node has been registered.

```
# CCPS_USERNAME=CPSADM@VCS_SERVICES
# CPS_DOMAINTYPE=vx
[root@cfs01 install]# cpsadm -s cps1v -a list_nodes
ClusterName UUID                                 Hostname(Node ID) Registered
=========== ======================================== ================  ==========
cfs0        {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs01(0)             1
cfs0        {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs02(1)             1
cfs0        {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs03(2)             1
cfs0        {38910d38-1dd2-11b2-a898-f1c7b967fd89} cfs04(3)             1
[root@cfs01 install]# cpsadm -s cps1v -a list_membership -c cfs0
List of registered nodes: 0 1 2 3
```

2 Run the same command against the each CP server.

3 Using the VCS Cluster Explorer screen, we can see that the vxfen service group has been created to monitor CP servers and that it is healthy.

Section 4

Reference

Appendix A

Known issues and limitations

This appendix includes the following topics:

Prevention of Storage vMotion

In a configuration where VMDK files are used with the multi-writer flag, any attempt of migrating the VMDK file to another data store will be prevented with an error.

The operation is unable to succeed.

In order to migrate VMDK to different storage, Arctera InfoScale Enterprise functionalities can be used to transparently migrate data between different disks.

Appendix B

Where to find more information

This appendix includes the following topics:

InfoScale documentation

The latest documentation is available on the Services and Operations Readiness Tools (SORT) website in the Adobe Portable Document Format (PDF).

See the release notes for information on documentation changes in this release.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The documents are updated periodically for errors or corrections.

https://sort.veritas.com/documents

You need to specify the product and the platform and apply other filters for finding the appropriate document.

Service and support

To access the self-service knowledge base, go to the following URL:

https://www.veritas.com/support/en_US.html

About Services and Operations Readiness Tools (SORT)

Services and Operations Readiness Tools (SORT) is a website that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Arctera products.

SORT can help you do the following:

Prepare for your next installation or upgrade	■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture. ■ Analyze systems to determine if they are ready to install or upgrade Arctera products. ■ Download the latest patches, documentation, and high availability agents from a central repository. ■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems.
Manage risks	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository. ■ Identify and mitigate system and environmental risks. ■ Display descriptions and solutions for hundreds of Arctera error codes.
Improve efficiency	■ Find and download patches based on product version and platform. ■ List installed Arctera products and license keys. ■ Tune and optimize your environment.

Note: Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

https://sort.veritas.com