

Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide

AIX

6.0

Veritas Storage Foundation and High Availability Solutions Virtualization Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.0

Document version: 6.0.2

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

Symantec 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 Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's 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 Symantec's support offerings, you can visit our Web site at the following URL:

www.symantec.com/business/support/index.jsp

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

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	4	
Chapter 1	Introduction	13
	Purpose	13
	Target audience	13
	About Storage Foundation and High Availability Solutions	14
	Partition Mobility and Workload Migration	15
	About IBM LPARS with dedicated I/O	15
	About Active Memory Sharing	15
	About Active Memory Expansion	16
	Impact of over-provisioning on Storage Foundation and High Availability	16
	Supported Storage Foundation and High Availability Solutions functionality	16
	Veritas Cluster Server	17
	Storage Foundation	17
	Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation High Availability	17
Chapter 2	Veritas Cluster Server support for managing and monitoring logical partitions (LPARs)	19
	VCS support for IBM PowerVM	19
	VCS in the management LPAR	20
	VCS in the managed LPAR	21
	Limitations and unsupported LPAR features	23
	Live partition mobility of management LPARs	23
	Live partition mobility of managed LPARs	24
	VCS requirements for managing LPARs as virtual machines	25
	Setting up management LPAR	25
	Configuring password-less SSH communication between VCS nodes and HMC	26
	Setting up managed LPARs	26
	Creating an LPAR profile	27
	Bundled agents for managing the LPAR	28
	Configuring VCS service groups to manage the LPAR	29

	Failure scenarios handling	30
	Physical server is down	30
	Management LPAR is down but physical server is up	30
Chapter 3	Veritas Cluster Server Solutions for IBM LPARs with Virtual Ethernet	31
	About IBM Virtual Ethernet	31
	Shared Ethernet Adapter (SEA)	31
	VCS configuration in the Virtual Ethernet environment	32
	LLT Private links configuration	32
	VCS Agents	35
	Virtual Ethernet and Cluster Management Software	35
Chapter 4	Storage Foundation and High Availability Virtualization Solutions for IBM LPARs with virtual SCSI Devices	37
	About IBM LPARs with virtual SCSI devices	37
	What is a virtual SCSI (vSCSI) disk?	38
	Using Storage Foundation in the VIO client with virtual SCSI devices	38
	Using Storage Foundation with virtual SCSI devices	39
	Setting up DMP for vSCSI devices in the Virtual I/O Client	39
	About disabling DMP multi-pathing for vSCSI devices in the Virtual IO Client	39
	Preparing to install or upgrade Storage Foundation with DMP disabled for vSCSI devices in the Virtual I/O client	40
	Disabling DMP multi-pathing for vSCSI devices in the Virtual IO Client, after installation	40
	Adding and removing DMP support for vSCSI devices for an array	41
	How DMP handles I/O for vSCSI devices	41
	Using Veritas Cluster Server with virtual SCSI devices	43
Chapter 5	Veritas Dynamic Multi-Pathing for the Virtual I/O Server	45
	Virtual I/O server overview	45
	Virtual I/O Server (VIOS) requirements	46
	DMP administration and management on Virtual I/O Server	46
	Veritas Volume Manager (VxVM) administration and management	47
	Configuring DMP on Virtual I/O Server	48

	Installing Veritas Dynamic Multi-Pathing (DMP) on Virtual I/O Server	48
	Migrating from other multi-pathing solutions to DMP on Virtual I/O Server	48
	Example: migration from MPIO to DMP on Virtual I/O Server for a dual-VIOS configuration	50
	Example: migration from PowerPath to DMP on Virtual I/O Server for a dual-VIOS configuration	55
	Configuring DMP pseudo devices as virtual SCSI devices	59
	Exporting DMP devices as virtual SCSI disks	60
	Exporting a Logical Volume as a virtual SCSI disk	63
	Exporting a file as a virtual SCSI disk	65
	Extended attributes in VIO client for a virtual SCSI disk	67
	Configuration prerequisites for providing extended attributes on VIO client for virtual SCSI disk	67
	Displaying extended attributes of virtual SCSI disks	68
	Virtual IO client adapter settings for Dynamic Multi-Pathing in dual-VIOS configurations	69
Chapter 6	Storage Foundation and High Availability Virtualization Solutions for IBM LPARs with N_Port ID Virtualization	71
	About IBM LPARs with N_Port ID Virtualization (NPIV)	71
	Characteristics of a LUN through NPIV	72
	VIO requirements	73
	Hardware requirements	73
	Support for Storage Foundation in NPIV environment	73
	Storage Foundation	73
	Cluster File System	74
	Veritas Cluster Server	74
	Installation, patching, and configuration requirements	74
Chapter 7	Storage Foundation and High Availability support for Live Partition Mobility	75
	About Live Partition Mobility (LPM)	75
	SFHA supported configuration	76
	SFHA requirements for the Live Partition Mobility	76
	About VCS support for Live Partition Mobility	77
	Overview of partition migration process	77
	Performance considerations	78

Chapter 8	Storage Foundation and High Availability support for IBM Workload Partitions	79
	About IBM Workload Partitions	79
	System workload partition	80
	Application workload partition	81
	When to use WPARs	81
	Storage Foundation support for WPARs	82
	Using a VxFS file system within a single system WPAR	82
	WPAR with root (/) partition as vxfs	83
	Using VxFS as a shared file system	84
	WPAR mobility	85
	About VCS support for WPARs	86
	Overview of how VCS works with WPARs	86
	The ContainerInfo attribute	86
	The ContainerOpts attribute	87
	About the Mount agent	88
	About the WPAR agent	88
	About configuring VCS in WPARs	88
	Prerequisites for configuring VCS in WPARs	89
	Deciding on the WPAR root location	90
	Creating a WPAR root on local disk	90
	Creating WPAR root on shared storage using NFS	91
	Installing the application	93
	Verifying the WPAR configuration	97
	Maintenance tasks	98
	Troubleshooting information	98
	Configuring AIX WPARs for disaster recovery using VCS	99
Chapter 9	Data migration from Physical to Virtual Clients with NPIV	103
	About migration from Physical to VIO environment	103
	Migrating from Physical to VIO environment	104
	Storage Foundation requirements for migration	104
Chapter 10	Boot device management	105
	Using DMP to provide multi-pathing for the root volume group (rootvg)	105
	Boot device on NPIV presented devices	107
	Hardware and software requirements	107
	Boot Device Management	107
	NPIV for Data volumes	107

Glossary 109

Introduction

This chapter includes the following topics:

- [Purpose](#)
- [Target audience](#)
- [About Storage Foundation and High Availability Solutions](#)
- [Partition Mobility and Workload Migration](#)
- [About IBM LPARS with dedicated I/O](#)
- [About Active Memory Sharing](#)
- [About Active Memory Expansion](#)
- [Supported Storage Foundation and High Availability Solutions functionality](#)

Purpose

This document explains the use of Storage Foundation and High Availability (SFHA) Solutions 6.0 with IBM Power virtualization. The document shows how the SFHA Solutions fit into the configuration for each virtualization solution. Where applicable, this document provides case studies. The scope of the document is limited to SFHA Solutions 6.0. While much of the functionality discussed in this document existed in previously released versions of the product, all case studies and configuration options are based upon this release.

Target audience

The target audience for this document is a Solutions Architect, Solutions Planner, or a senior systems Administrator. Prior knowledge of Storage Foundation and Veritas Cluster Server are a prerequisite.

About Storage Foundation and High Availability Solutions

This topic describes the Storage Foundation and High Availability Solutions:

- *Storage Foundation* - Veritas Storage Foundation provides easy-to-use online storage management, enables high availability of data, optimized I/O performance, and allows freedom of choice in storage hardware investments. Veritas Storage Foundation is the base storage management offering from Symantec. It includes Veritas File System (VxFS) and Veritas Volume Manager (VxVM). Both VxFS and VxVM include advanced features such as journaling file system, storage checkpoints, dynamic multi-pathing, off-host processing, volume snapshots, and tiered storage. Storage Foundation comes in three editions: *Basic*, *Standard* and *Enterprise*.

Each targets different environments as described below:

- *Storage Foundation Basic* - Storage Foundation Basic is the freeware version of Storage Foundation. Available as a free download, it is limited to a maximum of 2 CPU and 4 volumes and 4 file systems.
- *Storage Foundation Standard* - Storage Foundation Standard is intended for SAN connected servers with high performance requirements and availability features, such as multiple paths to storage. This product is a minimum requirement for High Availability solutions.
- *Storage Foundation Enterprise* - Storage Foundation Enterprise includes the entire feature set of both File System and Volume Manager. It is designed for servers with large SAN connectivity, where high performance, off-host processing and tiered storage are desired.
- *Storage Foundation Cluster File System* - Veritas Storage Foundation Cluster File System provides an integrated solution for shared file environments. The solution includes Veritas Cluster File System, Cluster Volume Manager and Veritas Cluster Server to help implement robust, manageable, and scalable shared file solutions. Veritas Cluster File System provides linear scalability for parallel applications and is widely used as a fast failover mechanism to ensure that application downtime is minimized in the event of server or software failure. Veritas Storage Foundation Cluster File System enables cluster-wide volume and file system configurations for simplified management. Extending clusters is simplified as new servers adopt cluster-wide configurations.
- *Veritas Cluster Server* - Veritas Cluster Server is the industry's leading cross-platform clustering solution for minimizing application downtime. Through central management tools, automated failover, features to test disaster

recovery plans without disruption, and advanced failover management based on server capacity, Veritas Cluster Server enables IT managers to maximize resources by moving beyond reactive recovery to proactive management of application availability in heterogeneous environments.

- *Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)* - SF Oracle RAC leverages proprietary storage management and high availability technologies to enable robust, manageable, and scalable deployment of Oracle RAC on UNIX platforms. The solution uses Veritas Cluster File System technology that provides the dual advantage of easy file system management as well as the use of familiar operating system tools and utilities in managing databases. The solution stack comprises the Veritas Cluster Server (VCS), Veritas Cluster Volume Manager (CVM), Veritas Oracle Real Application Cluster Support (VRTSdbac), Veritas Oracle Disk Manager (VRTSodm), Veritas Cluster File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

Partition Mobility and Workload Migration

Information on IBM PowerVM configuration can be found in the *IBM Redbook* located at:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf>

About IBM LPARS with dedicated I/O

The baseline configuration is a traditional AIX deployment with dedicated HBAs and NICs. The deployment may include partitions with virtual CPUs or partitions that support DLPAR events.

For more information about specifying dedicated I/O for IBM LPARs, refer to the following IBM Redbook:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf>

About Active Memory Sharing

The Veritas Storage Foundation High Availability stack supports VIO clients that use memory from the Active Memory Sharing (AMS) pool. Active Memory Sharing is a virtualization technology that allows multiple partitions to share a pool of physical memory. AMS increases system memory utilization and reduces the amount of physical memory that the system requires.

Symantec recommends that the ratio of the physical memory in the AMS pool should comply with the AIX guidelines.

See the IBM Redpaper PowerVM Virtualization Active Memory Sharing document for the AIX guidelines.

About Active Memory Expansion

Active Memory Expansion relies on compression of in-memory data to increase the amount of data that can be placed into memory. This feature expands the effective memory capacity of a POWER7 system. The operating system manages the in-memory data compression, which is transparent to applications and users.

Active Memory Expansion is configurable per logical partition (LPAR) . Active Memory Expansion can be selectively enabled for one or more LPARs on a system. When Active Memory Expansion is enabled for a LPAR, the operating system compresses a portion of the LPAR's memory and leave the remaining portion of memory uncompressed. The memory is effectively broken up into two pools – a compressed pool and an uncompressed pool. The operating system dynamically varies the amount of memory that is compressed, based on the workload and the configuration of the LPAR.

Impact of over-provisioning on Storage Foundation and High Availability

Over-provisioning occurs when you select a high compression ratio. Select the compression ratio carefully, keeping in mind the minimum application requirement and the maximum load threshold.

Note that operating system commands present a different view of memory if Active Memory Expansion (AME) is enabled. The `prtconf -m` command displays actual physical memory. The `topas` command displays the relay memory considering AME.

Supported Storage Foundation and High Availability Solutions functionality

In this release, Storage Foundation and High Availability Solutions products are fully functional in the IBM virtualization environment. The following functionality is supported:

- Veritas Cluster Server
See [“Veritas Cluster Server”](#) on page 17.
- Veritas Storage Foundation

See [“Storage Foundation”](#) on page 17.

- Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation High Availability
See [“Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation High Availability”](#) on page 17.

Veritas Cluster Server

Veritas Cluster Server is supported in the IBM virtualization environment. LLT is supported with virtual NICs. VxFEN is supported in NPIV environment but not with the virtual SCSI.

Storage Foundation

Storage Foundation stack is fully functional, including support for Portable Data Containers (Cross-platform Data Sharing format).

In the IBM virtualization environment, Storage Foundation runs on physical HBAs, and traditional LUN presentation is in place. Full DMP functionality is supported; the ability to query device information allows for LUN attribute discovery. This includes ALUA functionality, as well as thin reclamation and automated SSD discovery.

Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation High Availability

Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation High Availability are fully function in the IBM virtualization environment.

In this release, I/O fencing is supported. DMP is fully functional and supported.

Veritas Cluster Server support for managing and monitoring logical partitions (LPARs)

This chapter includes the following topics:

- [VCS support for IBM PowerVM](#)
- [Limitations and unsupported LPAR features](#)
- [VCS requirements for managing LPARs as virtual machines](#)
- [Setting up management LPAR](#)
- [Setting up managed LPARs](#)
- [Failure scenarios handling](#)

VCS support for IBM PowerVM

IBM PowerVM is virtualization solution for AIX environments on IBM POWER technology. In IBM PowerVM environment, multiple logical partitions (LPARs) can be carved in a physical server, also called the managed system. The LPARs can be assigned physical or virtual resources. The Virtual I/O Server is a dedicated partition and is a software appliance with which you can associate physical resources and that allows you to share these resources among multiple client logical partitions. The Virtual I/O Server can use both virtualized storage and network adapters. The managed system, LPARs and the resources are managed

using the Hardware Management Console (HMC) appliance sitting outside the physical frame.

You can use Veritas Cluster Server (VCS) by Symantec in a IBM PowerVM virtualization environment to provide mission-critical clustering and failover capabilities.

The following configurations are supported:

- LPAR to LPAR Clustering and Fail-over for LPAR availability (virtual machine) availability

VCS runs in one LPAR on each physical server, known as the management LPAR. The management LPAR behaves as a control point, and provides high-availability to the other LPARs running on the same physical server. The management LPAR views the LPARs that it manages as virtual machines but does not have visibility into the applications on the managed LPARs.

You can create a cluster of the management LPARs on more than one physical server to provide failover between the managed LPARs on the different physical servers.

See [“VCS in the management LPAR”](#) on page 20.

- VCS in the managed LPAR

VCS can also run inside the individual managed LPARs to provide high-availability for applications running inside the LPAR.

See [“VCS in the managed LPAR”](#) on page 21.

VCS in the management LPAR

VCS provides high availability for the AIX LPARs within a physical server. VCS is run in the control point which is an LPAR that is designated for management of other LPARs. The management LPARs on different physical servers form a VCS cluster.

VCS runs in one management LPAR on each physical server. The management LPAR provides high availability to the other LPARs on the same physical server, known as managed LPARs. Each managed LPAR is simply a resource that is managed and monitored by VCS running on the management LPAR, with the help of LPAR agent. This capability allows VCS to monitor the individual LPAR as an individual resource. VCS can restart the service group that has the LPAR resource on the same physical server or fail-over to another physical server.

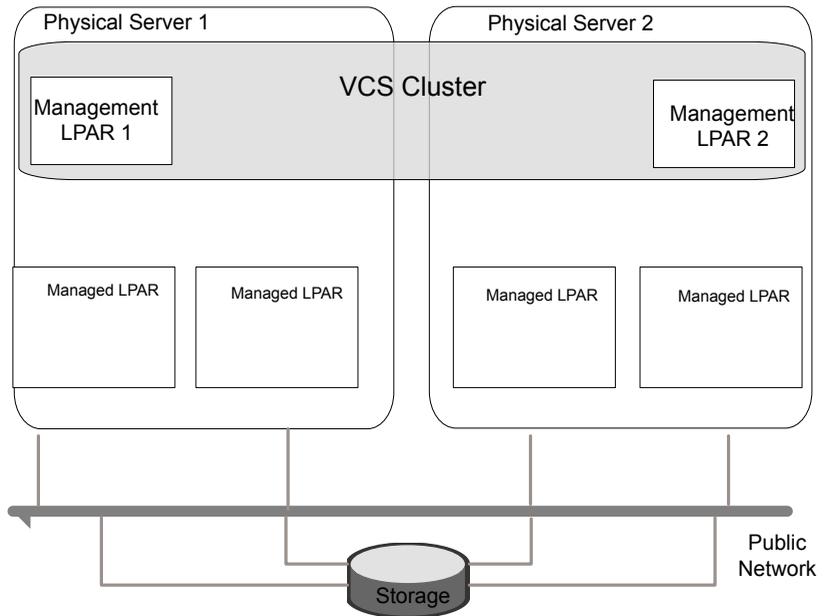
The management LPAR views the LPARs that it manages as virtual machines but does not have visibility into the applications on the managed LPARs. The management LPAR cluster does not monitor resources inside the managed LPARs.

A VCS cluster is formed among the management LPARs in this configuration. The VCS cluster provides failover for the managed LPARs between the management

LPARs. The managed LPARs can be either on the same physical server or on different physical servers. Each physical server should have a management server. Symantec recommends that the management LPARs are on separate physical servers. Each managed LPAR resource can have only one VCS system on one physical server in the system list.

Figure 2-1 provides an example of VCS in the management LPAR.

Figure 2-1 VCS in the management LPAR



This configuration also provides high availability for applications running on the management LPAR. The VCS cluster manages and controls the applications and services that run inside the management LPARs. Any faulted application or service is failed over to other management LPARs in the cluster.

Note: The managed LPARs cannot be in a cluster configuration.

VCS in the managed LPAR

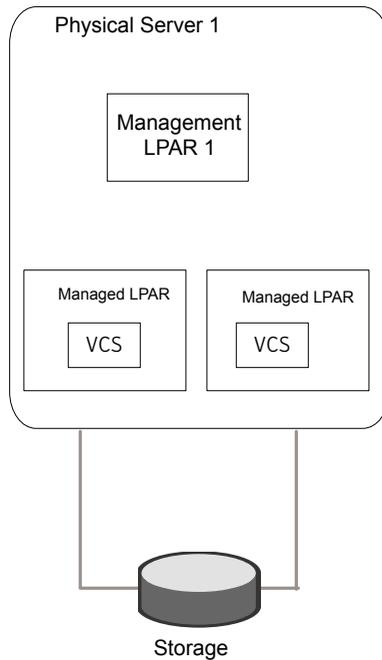
Veritas Cluster Server (VCS) can run within each managed LPAR to provide application monitoring and fault handling of applications running within the LPAR. VCS manages and controls the applications and services that run inside

the LPARs. This configuration provides restart of the application within the LPAR, but not failover between managed LPARs or physical servers.

In this configuration, the LPARs do not form a cluster. VCS runs as single-node VCS.

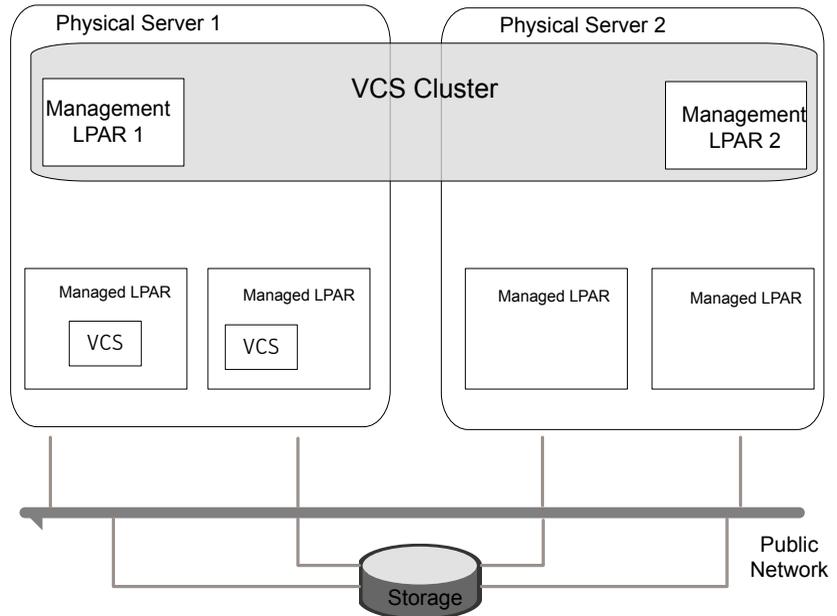
[Figure 2-2](#) provides an example of VCS in the managed LPAR.

Figure 2-2 VCS in the managed LPAR



[Figure 2-3](#) provides an example of VCS in both the managed LPAR and the management LPAR. VCS running in each managed LPAR provides high availability to applications running within the managed LPAR. VCS running in the management LPARs provide high availability to managed LPARs. The VCS instances work independently.

Figure 2-3 VCS in both the managed LPARs and the management LPARs



Limitations and unsupported LPAR features

The following limitations apply to VCS support for LPARs:

- Live partition mobility of management LPARs is not supported.
 If LPARs are managed by VCS running in the management LPAR, then the live partition mobility of the management LPAR is blocked by VCS.
 If you need to migrate the management LPAR, follow the recommended steps.
 See “ [Live partition mobility of management LPARs](#)” on page 23.
- If the LPAR agent crashes, the migration of the management LPAR will remain blocked even if it is not managing any LPARs. To unblock, you can perform the following:

```
# /usr/sbin/drmgr -u vcs_blockmigrate.sh
```

Live partition mobility of management LPARs

Live partition mobility is not supported if the VCS is running in the management LPAR on a physical system.

If LPARs are managed by VCS running in the management LPAR, then the live partition mobility of the management LPAR is blocked by VCS. If you need to migrate the management LPAR, use the following procedure.

To migrate the management LPAR

- 1 Migrate or fail over the managed LPARs to another physical server before migrating the management LPAR.
- 2 Stop the LPAR agent.
- 3 Migrate the management LPAR.
- 4 When management LPAR is back on the source physical server (which matches with the value of PhysicalServer in the VCS system), start the LPAR agent.

Live partition mobility of managed LPARs

When LPAR is managed by VCS, set the virtual SCSI adapter with any partition and any slot on the VIO. Map the virtual SCSI adapter to the correct SCSI adapter on the managed LPAR. This step needs to be part of the initial configuration. Otherwise, reboot the VIO so that the configuration takes effect before you perform the migration.

To migrate the managed LPAR

- 1 From the source managed system, back up the LPAR profile. After migration completes, the LPAR and its profiles are automatically deleted from the source.

For VCS to manage the LPAR, the profile is required on the managed physical system of the management VCS that is part of the system list of the LPAR resource.
- 2 On the destination system, rename the LPAR profile that was created during initial configuration of LPAR as a resource on all systems. LPM validation fails if it finds the profile with same LPAR name on the destination managed physical system
- 3 Migrate the managed LPAR.
- 4 Perform one of the following:
 - If migration succeeds, the profile on source is removed. Restore and rename the LPAR profile from the backup that was taken in step 1. Remove the renamed LPAR profile on the destination.
 - If migration fails, remove the backup profile on the source. On the destination, rename the renamed LPAR profile to original LPAR profile.

VCS requirements for managing LPARs as virtual machines

To use VCS to manage LPARs as virtual machines, the following requirements must be met.

Table 2-1 System requirements for the LPAR-supported configurations

VCS version	6.0
Supported OS version in LPARs	AIX 6.1 TL 5 and above. AIX 7.1 and above.
Supported VIOS version	2.1.3.10-FP-23 and above
Supported HMC version	7.2.0.0 Note: All the physical servers that are part of a cluster should be managed by the same HMC.
Supported hardware	Power5,6,7

Setting up management LPAR

Following is a high-level overview of the steps required for setting up the management LPARs.

Setting up management LPAR

- 1 Install VCS on all nodes of the management LPARs cluster. For information about installing VCS, refer to the Veritas Cluster Server Installation Guide.
- 2 Make sure that the HMC is at the supported level.
- 3 Make sure that the VIOS are at the supported level.
- 4 Configure password-less SSH communication from the management LPAR to the HMC. This step is required for all nodes of the cluster even if a node has no LPAR resource.
See “[Configuring password-less SSH communication between VCS nodes and HMC](#)” on page 26.
- 5 The managed LPARs are managed via HMC. Ensure that the network connection between the physical servers and the HMC has redundancy.
- 6 Set auto restart of management LPAR to “on”.

- 7 Ensure that PhysicalServer system level attribute has the physical server name.

Use the following command to retrieve the physical server name.

```
# lssyscfg -r sys -F name
```

- 8 Set up the managed LPARs.

See “[Setting up managed LPARs](#)” on page 26.

- 9 Configure the LPARs that need to be managed and monitored as VCS LPAR resources.

See “[Bundled agents for managing the LPAR](#)” on page 28.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

- 10 Set the Service Group level attribute SysDownPolicy = {"AutoDisableNoOffline"} for groups that have LPAR resources.

See “[Failure scenarios handling](#)” on page 30.

For more information on the Service Group level attribute SysDownPolicy, see the *Veritas Cluster Server User's Guide*.

Configuring password-less SSH communication between VCS nodes and HMC

To use remote command operations on the HMC, you must have SSH installed on the LPARs in the VCS cluster. You must configure the HMC to allow password-less SSH access from these LPARs. Refer to the appropriate IBM AIX documentation for information.

To verify that you have password-less SSH

- ◆ From each LPAR in the cluster, execute the following command to test if the password-less access works:

```
> ssh -l hscroot hmc2.veritas.com
Last login:Thur Jun 16 22:46:51 2011 from 10.182.9.34
hscroot@hmc2:~>
```

Setting up managed LPARs

The following procedure provides a high-level overview of how to set up LPARs that VCS manages.

For detailed instructions on creating and configuring a LPAR, refer to the IBM PowerVM Guide.

To set up managed LPARs

- 1 Ensure CPU and memory resources are available to create managed LPARs on all physical servers in the cluster, where the managed LPAR can start.
- 2 Install VCS on all the management LPARs, to manage the LPAR.
For information about installing VCS, see the *Veritas Cluster Server Installation Guide*.
- 3 Create the LPAR profile on all the physical servers whose management LPAR is in the systemlist for the LPAR resource.
See [“Creating an LPAR profile”](#) on page 27.
- 4 Set auto restart of the managed LPAR to “off” via HMC when VCS is managing the LPAR.
- 5 The boot disk should be shared and should be accessible from all the physical servers where the LPAR can fail over.
- 6 Verify if the LPAR can fail over to the other physical servers.
- 7 Configure the LPAR as a resource in VCS.
See [“Configuring VCS service groups to manage the LPAR”](#) on page 29.

Creating an LPAR profile

The following steps describe how to create LPAR profile:

To create an LPAR profile

- 1 Identify the disk on which the AIX OS is to be installed for the LPAR. This disk should be on shared storage in order for LPAR to be capable of failover across physical servers. Change the following attributes of the disk on all VIOS of the physical servers that the LPAR will be configured to boot on.

```
vio1#chdev -a hcheck_cmd=inquiry -l hdisk7
vio1#chdev -a hcheck_interval=60 -l hdisk7 -P
vio1#chdev -a pv=yes -l hdisk7
vio1#chdev -a reserve_policy=no_reserve
```

- 2 Create the Virtual SCSI Host adapter on all VIOS on which the LPAR will be configured to boot on. Reboot the VIO, and then map the OS disk to this host adapter.

- 3 Log in to HMC and create the LPAR profile. The following example shows creating an LPAR profile.

```
hscadmin1@hmc2.veritas.com:~> lssyscfg -r sys -F name
PServer1-SN100129A
PServer2-SN100130A

hscadmin1@hmc2.veritas.com:~> lssyscfg -m PServer1-SN100129A -r lpar \
-F name
Pserver1_VIO1

hscadmin1@hmc2.veritas.com:~> mksyscfg -m PServer1-SN100129A -r lpar \
-i name=lpar_test,lpar_env=aixlinux,profile_name=lpar_test,min_mem=512,\
desired_mem=512,max_mem=512,proc_mode=shared,sharing_mode=uncap,\
uncap_weight=128,min_proc_units=0.1,desired_proc_units=0.4,\
max_proc_units=2.0,min_procs=1,desired_procs=2,max_procs=4,\
lpar_io_pool_ids=none,max_virtual_slots=10,auto_start=1,\
boot_mode=norm,power_ctrl_lpar_ids=none,conn_monitoring=0,\
virtual_eth_adapters=2/1/1//0/1,virtual_scsi_adapters=3/client/1//10/1"

hscadmin1@hmc2.veritas.com:~> lssyscfg -m PServer1-SN100129A \
-r lpar -F name
Pserver1_VIO1
lpar_test
```

The virtual Ethernet adapter's VLAN ID should match that of VIO server in order for connectivity to outside network, the virtual scsi adapter's remote-lpar-ID/remote-lpar-name/remote-slot-number should match with that of VIO's partition ID, VIO's name and VIO's virtual SCSI Host adapter ID that has the OS disk mapped for this LPAR. Note: The VIO's virtual SCSI Host adapter that is assigned for this LPAR should have any partition and any slot option if this LPAR is capable and might be used for LPM in future (in addition to VCS failover capability).

- 4 Create the same profile on all physical servers where the LPAR can fail over.
- 5 Verify that the LPAR can boot on the physical servers where the profile has been created.

Bundled agents for managing the LPAR

The LPAR agent can be used to manage and provide high availability for LPARs.

The LPAR agent performs the following functions using HMC CLIs:

- **Open:** Blocks migration of the management LPAR. Get the information required by the LPAR agent.
- **Monitor:** Monitors the status of LPAR.
- **Online:** Starts the LPAR.
- **Offline:** Shuts down the LPAR.
- **Clean:** Stops the LPAR forcefully.
- **Shutdown:** Unblock migration of the management LPAR.

Configuring VCS service groups to manage the LPAR

You must configure a VCS service group to manage the LPAR.

To configure LPAR service groups

- 1 Create a failover service group for LPAR.
- 2 Set SysDownPolicy = { "AutoDisableNoOffline" } for this group.
- 3 Configure all the cluster nodes (management LPARs) in the SystemList attribute where the managed LPAR can fail over.
- 4 Configure LPAR resource for the managed LPAR.

Sample configuration for LPAR service group

Sample configuration

```
include "types.cf"
cluster lparcluster1 (
)

system sysa (
)

system sysb (
)

group virtgrp (
  SystemList = { sysa = 0, sysb = 1 }
  SysDownPolicy = { "AutoDisableNoOffline" }
)

LPAR lpar1 (
  LPARName = lpar1
```

```
MCUser = sr90901
MCName = { "thorhmc1.xyz.com" }
OnlineWaitLimit = 2
OfflineWaitLimit = 2
)
```

Failure scenarios handling

VCS handles the LPAR failures in the following cases.

- Physical server is down
- Management LPAR is down but physical server is up

Physical server is down

When the physical server is down, the management LPAR as well as managed LPARs will be down. In this case, the managed LPARs will be failed over by the sysoffline trigger using HMC. Ensure that HMC access is setup on all nodes of the cluster even if the node is not managing any LPAR.

Management LPAR is down but physical server is up

When the management LPAR is down, the physical server may not be down. The managed LPARs might be running. In this case, it is not desirable to automatically failover the managed LPARs. To ensure this the group that has LPAR resource should have SysDownPolicy = { "AutoDisableNoOffline" }. With this the groups will remain autodisabled on system fault. You can online the LPAR on any other system by setting autoenable for the group, after ensuring that the LPAR is down on the faulted system.

Veritas Cluster Server Solutions for IBM LPARs with Virtual Ethernet

This chapter includes the following topics:

- [About IBM Virtual Ethernet](#)
- [VCS configuration in the Virtual Ethernet environment](#)
- [Virtual Ethernet and Cluster Management Software](#)

About IBM Virtual Ethernet

Virtual Ethernet enables communication between inter-partitions on the same server, without requiring each partition to have a physical network adapter. You can define in-memory connections between partitions that are handled at the system level (for example, interaction between POWER Hypervisor and the operating systems). These connections exhibit characteristics similar to physical high-bandwidth Ethernet connections and support the industry standard protocols (such as IPv4, IPv6, ICMP, or ARP). Virtual Ethernet also enables multiple partitions to share physical adapters for access to external networks using Shared Ethernet Adapter (SEA).

Shared Ethernet Adapter (SEA)

A Shared Ethernet Adapter is a layer-2 network bridge to securely transport network traffic between virtual Ethernet networks and physical network adapters. The SEA also enables several client partitions to share one physical adapter. The SEA is hosted in the Virtual I/O Server.

To bridge network traffic between the internal virtual network and external networks, configure the Virtual I/O Server with at least one physical Ethernet adapter. Multiple virtual Ethernet adapters can share one SEA. Each virtual Ethernet adapter can support multiple VLANs.

The SEA has the following characteristics:

- Virtual Ethernet MAC addresses of virtual Ethernet adapters are visible to outside systems (using the `arp -a` command).
- Supports unicast, broadcast, and multicast. Protocols such as Address Resolution Protocol (ARP), Dynamic Host Configuration Protocol (DHCP), Boot Protocol (BOOTP), and Neighbor Discovery Protocol (NDP) can work across an SEA.

VCS configuration in the Virtual Ethernet environment

To use VCS in the Virtual Ethernet environment, configure VCS according to the following sections:

- Configure the LLT private links.
See [“LLT Private links configuration”](#) on page 32.
- Configure the VCS Agents.
See [“VCS Agents”](#) on page 35.

LLT Private links configuration

LLT heartbeats

LLT uses standard Ethernet networks to provide communication for its heartbeats. These networks can be provided through physical ports or virtual Ethernet interfaces. These interfaces do not require IP addresses to be configured since LLT heartbeats are based on layer 2 protocols. The best practice includes two independent paths for heartbeats to eliminate any single point of failure. This scenario includes redundant VIO servers with each providing a virtual Ethernet to each client LPAR participating in the VCS cluster.

LLT Private Links connections

The diagrams illustrate LLT Heartbeat connections in an IBM VIO environment with Virtual Ethernet, Shared Ethernet Adapters, and LPARs. The three node cluster consists of two VIO Client Partitions in System A and one LPAR in System B. POWER6 based systems that are controlled by the same Hardware Management Console (HMC).

Figure 3-1 shows an example of an environment with a single Virtual I/O Server

Figure 3-1 An environment with a single Virtual I/O Server

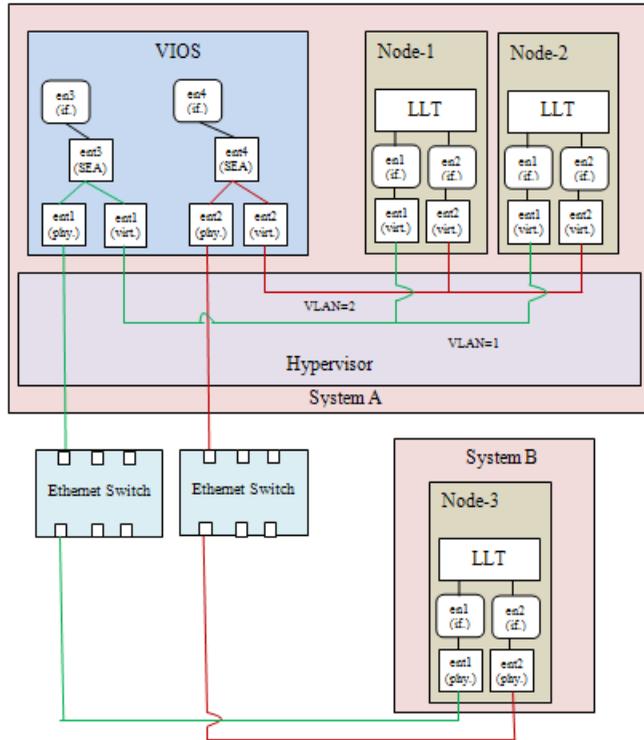


Figure 3-2 shows an example of an environment with two Virtual I/O Servers


```
link en1 /dev/dlpi/en:1 - ether - 1500  
link en2 /dev/dlpi/en:2 - ether - 1500
```

After you configure the LLT, use the below command on all the nodes of your cluster to be sure that the overall MTU size is less than 1500.

```
# lltstat -c | grep mtu  
mtu: 1460
```

VCS Agents

All VCS network-related agents support Virtual Ethernet environment. You can configure any of the bundled networking agents to monitor the network inside an LPAR.

Virtual Ethernet and Cluster Management Software

Virtual Ethernet environment offers various advantages and flexibility, but you should be aware of the challenges. The various independent clusters consisting of VIO client partitions in the same physical computer can be configured with the heartbeat routed through the same physical Ethernet adapters to additional nodes outside the physical computer. Ensure that each cluster has a unique cluster ID. Unique cluster IDs eliminate conflict and allow the Virtual Ethernet environment to greatly reduce the required number of physical Ethernet adapters. According to IBM, there are issues to be aware that are not the fault of the applicable Cluster Management Software or the configuration. Rather, the issues arise as a direct consequence of I/O virtualization.

To reiterate, although some of these may be viewed as configuration restrictions, many are direct consequences of I/O Virtualization.

The issues and recommendation are as follows:

- If two or more Clustered nodes use a VIO server or servers in the same frame, the Cluster Management Software cannot detect and react to single physical interface failures. This behavior does not limit the availability of the entire cluster because VIOS itself routes traffic around the failure. The behavior of the VIOS is analogous to AIX the EtherChannel. Notification of individual Adapter failures must use other methods (not based on the VIO server).
- All Virtual Ethernet interfaces that are defined to the Cluster Management Software should be treated as “single-Adapter networks” according to IBM. To correctly monitor and detect failure of the network interface, you must create a file that includes a list of clients to ping. Due to the nature of Virtual

Ethernet, other mechanisms to detect the failure of network interfaces are not effective.

- If the VIO server has only a single physical interface on a network, then the Cluster Management Software can detect a failure of that interface. However, that failure isolates the node from the network.

Check the IBM documentation for detailed information on the Virtual Ethernet and various configuration scenarios using virtual I/O Server. For information about the above issues, see the following link:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10390>

Storage Foundation and High Availability Virtualization Solutions for IBM LPARs with virtual SCSI Devices

This chapter includes the following topics:

- [About IBM LPARs with virtual SCSI devices](#)
- [Using Storage Foundation in the VIO client with virtual SCSI devices](#)
- [Using Veritas Cluster Server with virtual SCSI devices](#)

About IBM LPARs with virtual SCSI devices

This discussion of vSCSI devices applies only to SAN-based LUNs presented through VIO. Internal devices, volumes, and files presented by VIO as vSCSI devices are not recommended for use with Storage Foundation.

Virtual SCSI uses a client/server model. A Virtual I/O server partition owns the physical I/O devices, and exports the devices as virtual SCSI (vSCSI) resources to the client partitions. The Virtual I/O client is a logical partition that has a virtual client adapter node defined in its device tree. The VIO client uses the vSCSI resources provided by the Virtual I/O Server partition to access the storage devices.

If redundant SAN connections exist to the VIO server, the VIO server provides multi-pathing to the array. Client partitions can also perform multi-pathing between VIO servers in an active/standby configuration. This configuration provides extended protection from VIO configuration and maintenance. Redundant VIO servers are recommended for production workloads.

What is a virtual SCSI (vSCSI) disk?

A virtual SCSI (vSCSI) disk is a resource which can be a SCSI disk, or a volume or file in a VIO Server (VIOS) that is exported to a virtual IO client (VIOC). IBM vSCSI LUNs implement a sub-set of the SCSI protocol. The two main limitations are:

- Persistent reservations (SCSI3 – PGR) are not implemented.
The lack of SCSI reservations means that I/O Fencing is not supported. Storage Foundation Cluster File System (SFCFS) and Storage Foundation for Oracle RAC (SFRAC) do not support vSCSI disks, because SFCFS and SFRAC require I/O fencing.
- Device inquiry limitations.
Veritas Storage Foundation (SF) cannot directly fetch the inquiry data, as is done from a physical SCSI disk. However, if the vSCSI disk in VIOC is backed by a dmpnode in VIOS, then all the inquiry data that can be fetched from a physical disk can be fetched.
Starting with release SF 5.1, cross-platform data sharing (CDS) functionality is supported.

Using Storage Foundation in the VIO client with virtual SCSI devices

Storage Foundation provides support for virtual SCSI (vSCSI) devices on the VIO client. You can create and manage Veritas Volume Manager (VxVM) volumes on vSCSI devices, as for any other devices. Storage Foundation provides Dynamic Multi-Pathing (DMP) for vSCSI devices, by default. Storage Foundation can also co-exist with MPIO for multi-pathing. If you choose to use MPIO to multipath the vSCSI devices, DMP works in pass-through mode.

Use the `vxddladm` utility and the `vxdmppadm` utility to administer DMP for vSCSI devices. The `vxddladm` utility controls enabling and disabling DMP on vSCSI devices, adding and removing supported arrays, and listing supported arrays. The `vxdmppadm` utility controls the I/O policy and the path policy for vSCSI devices.

Using Storage Foundation with virtual SCSI devices

Versions of SF that support vSCSI disks are:

Prior to Storage Foundation 5.1, Portable Data Containers (disk type CDS) were not supported. With extensions included in Storage Foundation 5.1, CDS type devices are now supported.

Storage Foundation can be used in the following ways:

- use DMP in the VIO server to provide multi-pathing to the array. DMP presents a dmpnode as a vSCSI device to the VIO client.
- use Storage Foundation in the VIO client to provide volume management on the vSCSI devices, and multi-pathing through the VIO servers with DMP.
- use SF in the VIO client to provide volume management on the vSCSI devices, and use MPIO to provide multi-pathing.

Setting up DMP for vSCSI devices in the Virtual I/O Client

In this release of Storage Foundation, Veritas Dynamic Multi-Pathing (DMP) is enabled on VIO clients by default. After you install or upgrade Storage Foundation in the Virtual IO client, any vSCSI devices are under DMP control. MPIO is disabled.

If you have already installed or upgraded Storage Foundation in the Virtual I/O client, use the following procedure to enable DMP support for vSCSI devices. This procedure is only required if you have previously disabled DMP support for vSCSI devices.

To enable vSCSI support within DMP and disable MPIO

- 1 Enable vSCSI support.

```
# vxddladm enablevscsi
```

- 2 You are prompted to reboot the system, if required.

DMP takes control of the devices, for any array that has DMP support to use the array for vSCSI devices. You can add or remove DMP support for vSCSI for arrays.

See [“Adding and removing DMP support for vSCSI devices for an array”](#) on page 41.

About disabling DMP multi-pathing for vSCSI devices in the Virtual IO Client

Storage Foundation can co-exist with MPIO multi-pathing in the Virtual I/O client. If you prefer to use MPIO for multi-pathing, you can override the default behavior, which enables Dynamic Multi-Pathing (DMP) in the Virtual I/O client.

There are two ways to do this:

- Before you install or upgrade Storage Foundation in the Virtual I/O client
See [“Preparing to install or upgrade Storage Foundation with DMP disabled for vSCSI devices in the Virtual I/O client”](#) on page 40.
- After Storage Foundation is installed in the Virtual I/O client
See [“Disabling DMP multi-pathing for vSCSI devices in the Virtual IO Client, after installation”](#) on page 40.

Preparing to install or upgrade Storage Foundation with DMP disabled for vSCSI devices in the Virtual I/O client

Before you install or upgrade Storage Foundation, you can set an environment variable to disable DMP use for the vSCSI devices. Storage Foundation is installed with DMP in pass-through mode. MPIO is enabled for multi-pathing.

Note: When you upgrade an existing VxVM installation that has DMP enabled, then DMP remains enabled regardless of whether or not the environment variable `__VXVM_DMP_VSCSI_ENABLE` is set to no.

To disable DMP before installing or upgrading SF in the Virtual I/O Client

- 1 Before you install or upgrade VxVM, set the environment variable `__VXVM_DMP_VSCSI_ENABLE` to no.

```
# export __VXVM_DMP_VSCSI_ENABLE=no
```

Note: The environment variable name `__VXVM_DMP_VSCSI_ENABLE` begins with two underscore (`_`) characters.

- 2 Install Storage Foundation, as described in the *Storage Foundation High Availability Installation Guide*

Disabling DMP multi-pathing for vSCSI devices in the Virtual IO Client, after installation

After VxVM is installed, use the `vxd1adm` command to switch vSCSI devices between MPIO control and DMP control.

To return control to MPIO, disable vSCSI support with DMP. After DMP support has been disabled, MPIO takes control of the devices. MPIO implements

multi-pathing features such as failover and load balancing; DMP acts in pass-through mode.

To disable vSCSI support within DMP and enable MPIO

- 1 Disable vSCSI support.

```
# vxddladm disablevscsi
```

- 2 You are prompted to reboot the system, if required.

Adding and removing DMP support for vSCSI devices for an array

Veritas Dynamic Multi-Pathing (DMP) controls the devices for any array that has DMP support to use the array for vSCSI devices.

To add or remove DMP support for an array for use with vSCSI devices

- 1 To determine if DMP support is enabled for an array, list all of the arrays that DMP supports for use with vSCSI devices:

```
# vxddladm listvscsi
```

- 2 If the support is not enabled, add support for using an array as a vSCSI device within DMP:

```
# vxddladm addvscsi array_vid
```

- 3 If the support is enabled, you can remove the support so that the array is not used for vSCSI devices within DMP:

```
# vxddladm rmvscsi array_vid
```

- 4 You are prompted to reboot the system, if required.

How DMP handles I/O for vSCSI devices

On the VIO client, DMP uses the Active/Standby array mode for the vSCSI devices. Each path to the vSCSI device is through a VIO server. One VIO server is Active and the other VIO servers are Standby. An Active/Standby array permits I/O through a single Active path, and keeps the other paths on standby. During failover, I/O is scheduled on one of the standby paths. After failback, I/Os are scheduled back onto the original Active path. The Active/Standby mode is a variation of an active/active array; only one path is active at a time.

The DMP I/O policy for vSCSI enclosures is always Single-Active. It is not possible to change the DMP I/O policy for the vSCSI enclosure, because only one VIO server can be Active.

The following command shows the vSCSI enclosure:

```
# vxdmpadm listenclosure all
ENCLR_NAME    ENCLR_TYPE  ENCLR_SNO    STATUS      ARRAY_TYPE  LUN_COUNT
=====
ibm_vscsi0    IBM_VSCSI   VSCSI        CONNECTED   VSCSI       9
```

The following command shows the I/O policy for the vSCSI enclosure:

```
# vxdmpadm getattr enclosure ibm_vscsi0 iopolicy
ENCLR_NAME    DEFAULT      CURRENT
=====
ibm_vscsi0    Single-Active Single-Active
```

For vSCSI devices, DMP balances the load between the VIO servers, instead of balancing the I/O on paths. By default, the `iopolicy` attribute of the vSCSI array is set to `lunbalance`. When `lunbalance` is set, the vSCSI LUNs are distributed so that the I/O load is shared across the VIO servers. For example, if you have 10 LUNs and 2 VIO servers, 5 of them are configured so that VIO Server 1 is Active and VIO Server 2 is Standby. The other 5 are configured so that the VIO Server 2 is Active and VIO Server 1 is Standby. To turn off load sharing across VIO servers, set the `iopolicy` attribute to `nolunbalance`.

DMP dynamically balances the I/O load across LUNs. When you add or remove disks or paths in the VIO client, the load is rebalanced. Temporary failures like enabling or disabling paths or controllers do not cause the I/O load across LUNs to be rebalanced.

Setting the vSCSI I/O policy

By default, DMP balances the I/O load across VIO servers. This behavior sets the I/O policy attribute to `lunbalance`.

To display the current I/O policy attribute for the vSCSI array

- ◆ Display the current I/O policy for a vSCSI array:

```
# vxdmpadm getattr vscsi iopolicy
VSCSI        DEFAULT      CURRENT
=====
IOPolicy    lunbalance   lunbalance
```

To turn off the LUN balancing, set the I/O policy attribute for the vSCSI array to `noLunbalance`.

To set the I/O policy attribute for the vSCSI array

- ◆ Set the I/O policy for a vSCSI array:

```
# vxdmpadm setattr vscsi iopolicy={lunbalance|noLunbalance}
```

Note: The DMP I/O policy for each vSCSI device is always Single-Active. You cannot change the DMP I/O policy for the vSCSI enclosure. Only one VIO server can be Active for each vSCSI device.

Using Veritas Cluster Server with virtual SCSI devices

Veritas Cluster Server (VCS) supports disk groups and volume groups created on virtual SCSI devices. The VCS DiskGroup agent supports disk groups. The VCS LVMVG agent supports volume groups.

Due to lack of SCSI3 persistent reservations, I/O Fencing is not supported with virtual SCSI devices.

Veritas Dynamic Multi-Pathing for the Virtual I/O Server

This chapter includes the following topics:

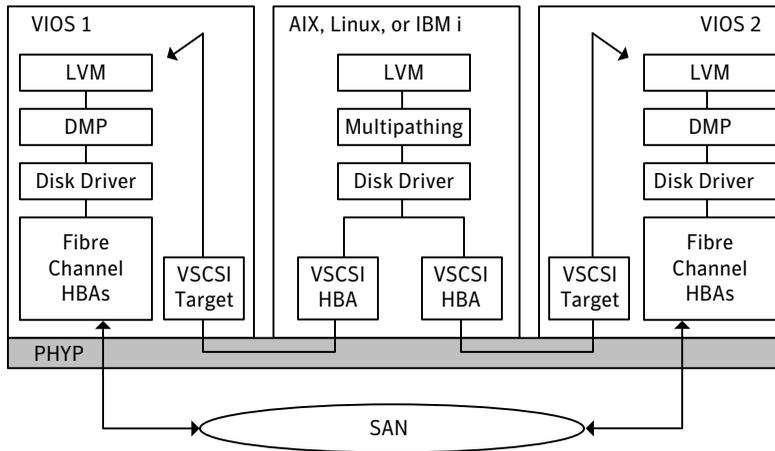
- [Virtual I/O server overview](#)
- [Virtual I/O Server \(VIOS\) requirements](#)
- [DMP administration and management on Virtual I/O Server](#)
- [Veritas Volume Manager \(VxVM\) administration and management](#)
- [Configuring DMP on Virtual I/O Server](#)
- [Configuring DMP pseudo devices as virtual SCSI devices](#)
- [Extended attributes in VIO client for a virtual SCSI disk](#)
- [Virtual IO client adapter settings for Dynamic Multi-Pathing in dual-VIOS configurations](#)

Virtual I/O server overview

Virtual I/O (VIO) server is a virtualization technology by IBM. A Virtual I/O server is a logical partition (LPAR) that runs a trimmed-down version of the AIX operating system. Virtual I/O servers have APV support, which allows sharing of physical I/O resources between virtual I/O clients.

[Figure 5-1](#) illustrates DMP enablement in the Virtual I/O Server.

Figure 5-1 Veritas Dynamic Multi-Pathing in the Virtual I/O Server



For support information concerning running Dynamic Multi-Pathing (DMP) in Virtual I/O Server (VIOS), see the *Veritas Dynamic Multi-Pathing Release Notes*.

See the PowerVM wiki for more in-depth information about VIO server and virtualization:

<http://www.ibm.com/developerworks/wikis/display/virtualization/VIO>

For more information, see the *PowerVM Virtualization on IBM System p redbook*:

<http://www.redbooks.ibm.com/redpieces/abstracts/sg247940.html>

Virtual I/O Server (VIOS) requirements

To run DMP in VIOS, the minimum VIOS level that is required is 2.1.3 10-FP-23 or later.

DMP administration and management on Virtual I/O Server

DMP is fully functional in the Virtual I/O server. DMP administration and management commands (`vxddmpadm`, `vxddladm`, `vxdisk`, etc.) must be invoked from the non-restricted root shell.

```
$ oem_setup_env
```

Some example commands:

```
dmpvios1$ vxddmpadm getsubpaths dmpnodename=ibm_ds8x000_0337
```

NAME	STATE[A]	PATH-TYPE[M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
hdisk21	ENABLED(A)	-	fscsi0	IBM_DS8x00	ibm_ds8x000	-
hdisk61	ENABLED(A)	-	fscsi0	IBM_DS8x00	ibm_ds8x000	-
hdisk80	ENABLED(A)	-	fscsi1	IBM_DS8x00	ibm_ds8x000	-
hdisk99	ENABLED(A)	-	fscsi1	IBM_DS8x00	ibm_ds8x000	-

```
dmpvios1$ vxddmpadm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE	LUN_COUNT
disk	Disk	DISKS	CONNECTED	Disk	1
ibm_ds8x000	IBM_DS8x00	75MA641	CONNECTED	A/A	6

For complete information about managing Dynamic Multi-Pathing, see the *Veritas Dynamic Multi-Pathing Administrator's Guide*.

Veritas Volume Manager (VxVM) administration and management

Veritas Volume Manager (VxVM) functionality is disabled in Virtual I/O Server. VxVM commands that manage volumes or disk groups are disabled in the VIO server.

In the VIOS, VxVM does not detect disk format information, so the disk status for VxVM disks is shown as unknown. For example:

```
dmpvios1$ vxddisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
disk_0	auto	-	-	unknown
ibm_ds8x000_02c1	auto	-	-	unknown
ibm_ds8x000_0288	auto	-	-	unknown
ibm_ds8x000_029a	auto	-	-	unknown
ibm_ds8x000_0292	auto	-	-	unknown
ibm_ds8x000_0293	auto	-	-	unknown
ibm_ds8x000_0337	auto	-	-	unknown

In the VIOS, VxVM displays an error if you run a command that is disabled, as follows:

```
dmpvios1$ vxdisk -f init ibm_ds8x000_0288
VxVM vxdisk ERROR V-5-1-5433 Device ibm_ds8x000_0288: init failed:
Operation not allowed. VxVM is disabled.

dmpvios1$ vxdg import bootdg
VxVM vxdg ERROR V-5-1-10978 Disk group bootdg: import failed:
Operation not allowed. VxVM is disabled.
```

Configuring DMP on Virtual I/O Server

In this release, you can install DMP in the virtual I/O server (VIOS). This enables the VIO server to export dmpnodes to the VIO clients. The VIO clients access the dmpnodes in the same way as any other vSCSI devices. DMP handles the I/O to the disks backed by the dmpnodes.

Installing Veritas Dynamic Multi-Pathing (DMP) on Virtual I/O Server

Veritas Dynamic Multi-Pathing (DMP) can operate in the Virtual I/O server. Install DMP on the Virtual I/O server.

To install DMP on the Virtual I/O Server

- 1 Log into the VIO server partition.
- 2 Use the `oem_setup_env` command to access the non-restricted root shell.
- 3 Install Veritas Dynamic Multi-Pathing on the Virtual I/O Server.
 See the *Veritas Dynamic Multi-Pathing Installation Guide*.
- 4 Installing DMP on the VIO server enables the `dmp_native_support` tunable. Do not set the `dmp_native_support` tunable to off.

```
dmpvios1$ vxddmoadm gettune dmp_native_support
Tunable                Current Value  Default Value
-----                -
dmp_native_support     on            off
```

Migrating from other multi-pathing solutions to DMP on Virtual I/O Server

DMP supports migrating from AIX MPIIO and EMC PowerPath multi-pathing solutions to DMP on Virtual I/O Server.

To migrate from other multi-pathing solutions to DMP on Virtual I/O Server

- 1 Before migrating, back up the Virtual I/O Servers to use for reverting the system in case of issues.
- 2 Shut down all VIO client partitions that are serviced by the VIOS.
- 3 Log into the VIO server partition. Use the following command to access the non-restricted root shell. All subsequent commands in this procedure must be invoked from the non-restricted shell.

```
$ oem_setup_env
```

- 4 For each Fibre Channel (FC) adapter on the system, verify that the following attributes have the recommended settings:

```
fc_err_recov          fast_fail
dyntrk                yes
```

If required, use the `chdev` command to change the attributes.

The following example shows how to change the attributes:

```
dmpvios1$ chdev -a fc_err_recov=fast_fail -a dyntrk=yes -l \
  fscsi0 -P
fscsi0 changed
```

The following example shows the new attribute values:

```
dmpvios1$ lsattr -El fscsi0

attach      switch  How this adapter is CONNECTED  False
dyntrk      yes     Dynamic Tracking of FC Devices  True
fc_err_recov fast_fail FC Fabric Event Error RECOVERY
Policy True
scsi_id     0xd0c00 Adapter SCSI ID                 False
sw_fc_class 3       FC Class for Fabric             True
```

- 5 Use commands like `lsdev` and `lsmmap` to view the configuration.
- 6 Unconfigure all VTD devices from all virtual adapters on the system:

```
dmpvios1$ rmdev -p vhost0
```

Repeat this step for all other virtual adapters.

7 Migrate from the third-party device driver to DMP.

Note that you do not need to do turn on the `dmp_native_support` again, because it is turned on for VIOS by default. You can use the `vxddmpadm gettune dmp_native_support` command to verify that the tunable parameter is turned on.

For the migration procedure, see the *Veritas Dynamic Multi-Pathing Administrator's Guide*.

8 Reboot the VIO Server partition.

9 Use the following command to verify that all Virtual SCSI mappings of TPD multi-pathing solution have been correctly migrated to DMP:

```
dmpvios1$ /usr/ios/cli/ioscli lsmmap -all
```

10 Repeat step 1 through step 9 for all of the other VIO server partitions of the managed system.

11 After all of the VIO Server partitions are successfully migrated to DMP, start all of the VIO client partitions.

Example: migration from MPIO to DMP on Virtual I/O Server for a dual-VIOS configuration

This section shows an example of a migration from MPIO to DMP on the Virtual I/O Server, in a configuration with two VIO Servers.

Example configuration:

```
Managed System: dmpviosp6  
VIO server1: dmpvios1  
VIO server2: dmpvios2  
VIO clients: dmpvioc1  
SAN LUNs: IBM DS8K array  
Current multi-pathing solution on VIO server: IBM MPIO  
  
ODM definition fileset required to disable MPIO support  
for IBM DS8K array LUNs:  
devices.fcp.disk.ibm.rte
```

To migrate dmpvios6 from MPIO to DMP

- 1 Before migrating, back up the Virtual I/O Server to use for reverting the system in case of issues.

See the IBM website for information about backing up Virtual I/O Server.

- 2 Shut down all of the VIO clients that are serviced by the VIO Server.

```
dmpvioc1$ halt
```

- 3 Log into the VIO server partition. Use the following command to access the non-restricted root shell. All subsequent commands in this procedure must be invoked from the non-restricted shell.

```
$ oem_setup_env
```

- 4 Verify that the FC adapters have the recommended settings. If not, change the settings as required.

For example, the following output shows the settings:

```
dmpvios1$ lsattr -El fscsi0
attach      switch  How this adapter is CONNECTED  False
dyntrk     yes     Dynamic Tracking of FC Devices  True
fc_err_recov fast_fail FC Fabric Event Error RECOVERY
Policy True
scsi_id     0xd0c00 Adapter SCSI ID                False
sw_fc_class 3       FC Class for Fabric             True
```

5 The following command shows `lsmmap` output before migrating MPIO VTD devices to DMP:

```
dmpvios1$ /usr/ios/cli/iosctl lsmmap -all
SVSA                Physloc                Client Partition ID
-----
vhost0              U9117.MMA.0686502-V2-C11  0x00000004

VTD                 vtscsi0
Status              Available 8100000000000000
Backing device      hdisk21
LUN                 0x
Physloc             U789D.001.DQD04AF-P1-C5-T1-W500507630813861A-L4
0034037000000000

VTD                 vtscsi1
Status              Available
LUN                 0x8200000000000000
Backing device      hdisk20
Physloc             U789D.001.DQD04AF-P1-C5-T1-W500507630813861A-L4
00240C1000000000

VTD                 vtscsi2
Status              Available
LUN                 0x8300000000000000
Backing device      hdisk18
Physloc             U789D.001.DQD04AF-P1-C5-T1-W500507630813861A-L4
002409A000000000
```

The VIO Server has MPIO providing multi-pathing to these hdisks. The following commands show the configuration:

```
dmpvios1$ lsdev -Cc disk | egrep "hdisk21|hdisk20|hdisk18"

hdisk18 Available 02-08-02 MPIO Other FC SCSI Disk Drive
hdisk20 Available 02-08-02 MPIO Other FC SCSI Disk Drive
hdisk21 Available 02-08-02 MPIO Other FC SCSI Disk Drive
```

6 Unconfigure all VTD devices from all virtual adapters on the system:

```
dmpvios1 $ rmdev -p vhost0  
vtscsi0 Defined  
vtscsi1 Defined  
vtscsi2 Defined
```

Repeat this step for all other virtual adapters.

7 Migrate the devices from MPIO to DMP.

Unmount the file system and varyoff volume groups residing on the MPIO devices.

Display the volume groups (vgs) in the configuration:

```
dmpvios1$ lsvg
rootvg
brunovg

dmpvios1 lsvg -p brunovg

brunovg:
PV_NAME PV STATE TOTAL PPs FREE PPs FREE DISTRIBUTION
hdisk19 active 511 501 103..92..102..102..102
hdisk22 active 511 501 103..92..102..102..102
```

Use the varyoffvg command on all affected vgs:

```
dmpvios1$ varyoffvg brunovg
```

Install the IBMDS8K ODM definition fileset to remove IBM MPIO support for IBM DS8K array LUNs.

```
dmpvios1$ installp -aXd . devices.fcp.disk.ibm.rte
```

```
+-----+
                Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...
Installation Summary
-----
Name                               Level  Part  Event  Result
-----
devices.fcp.disk.ibm.rte  1.0.0.2  USR   APPLY  SUCCESS
devices.fcp.disk.ibm.rte  1.0.0.2  ROOT  APPLY  SUCCESS
```

8 Reboot VIO server1

```
dmpvios1$ reboot
```

- 9** After the VIO server1 reboots, verify that all of the existing volume groups on the VIO server1 and MPIO VTDs on the VIO server1 are successfully migrated to DMP.

```
dmpvios1 lsvg -p brunovg
```

```
brunovg:
```

```
PV_NAME          PV STATE TOTAL PPs FREE PPs FREE DISTRIBUTION
ibm_ds8000_0292  active   511      501   103..92..102..102..102
ibm_ds8000_0293  active   511      501   103..92..102..102..102
```

Verify the vSCSI mappings of IBM DS8KLUNs on the migrated volume groups:

```
dmpvios1 lsmmap -all
```

SVSA	Physloc	Client Partition ID
vhost0	U9117.MMA.0686502-V2-C11	0x00000000
VTD	vtscsi0	
Status	Available	
LUN	0x8100000000000000	
Backing device	ibm_ds8000_0337	
Physloc		
VTD	vtscsi1	
Status	Available	
LUN	0x8200000000000000	
Backing device	ibm_ds8000_02c1	
Physloc		
VTD	vtscsi2	
Status	Available	
LUN	0x8300000000000000	
Backing device	ibm_ds8000_029a	
Physloc		

- 10** Repeat step 1 through step 9 for VIO server2.
- 11** Start all of the VIO clients using HMC.

Example: migration from PowerPath to DMP on Virtual I/O Server for a dual-VIOS configuration

This section shows an example of a migration from PowerPath to DMP on the Virtual I/O Server, in a configuration with two VIO Servers.

Example configuration:

```
Managed System: dmpviosp6  
VIO server1: dmpvios1  
VIO server2: dmpvios2  
VIO clients: dmpviocl  
SAN LUNs: EMC Clariion array  
Current multi-pathing solution on VIO server: EMC PowerPath
```

To migrate dmpviosp6 from PowerPath to DMP

- 1 Before migrating, back up the Virtual I/O Server to use for reverting the system in case of issues.

See the IBM website for information about backing up Virtual I/O Server.

- 2 Shut down all of the VIO clients that are serviced by the VIO Server.

```
dmpviocl$ halt
```

- 3 Log into the VIO server partition. Use the following command to access the non-restricted root shell. All subsequent commands in this procedure must be invoked from the non-restricted shell.

```
$ oem_setup_env
```

- 4 Verify that the FC adapters have the recommended settings. If not, change the settings as required.

For example, the following output shows the settings:

```
dmpvios1$ lsattr -El fscsi0  
attach      switch      How this adapter is CONNECTED  False  
dyntrk      yes         Dynamic Tracking of FC Devices  True  
fc_err_recov fast_fail   FC Fabric Event Error RECOVERY Policy  
True  
scsi_id      0xd0c00    Adapter SCSI ID                  False  
sw_fc_class  3          FC Class for Fabric              True
```

- 5 The following command shows `lsmmap` output before migrating PowerPath VTD devices to DMP:

```
dmpvios1$ /usr/ios/cli/iosctl lsmmap -all
```

SVSA	Physloc	Client Partition ID
-----	-----	-----
vhost0	U9117.MMA.0686502-V2-C11	0x00000004
VTD	P0	
Status	Available	
LUN	0x8100000000000000	
Backing device	hdiskpower0	
Physloc	U789D.001.DQD04AF-P1-C5-T1-W500507630813861A-L4	
0034037		
00000000		
VTD	P1	
Status	Available	
LUN	0x8200000000000000	
Backing device	hdiskpower1	
Physloc	U789D.001.DQD04AF-P1-C5-T1-W500507630813861A-L40	
0240C10		
00000000		
VTD	P2	
Status	Available	
LUN	0x8300000000000000	
Backing device	hdiskpower2	
Physloc	U789D.001.DQD04AF-P1-C5-T1-W500507630813861A-L40	
02409A00000000		

- 6 Unconfigure all VTD devices from all virtual adapters on the system:

```
dmpvios1 $ rmdev -p vhost0  
P0 Defined  
P1 Defined  
P2 Defined
```

Repeat this step for all other virtual adapters.

7 Migrate the devices from PowerPath to DMP.

Unmount the file system and varyoff volume groups residing on the PowerPath devices.

Display the volume groups (vgs) in the configuration:

```
dmpvios1$ lsvg  
rootvg  
brunovg
```

```
dmpvios1 lsvg -p brunovg
```

```
brunovg:  
PV_NAME      PV STATE TOTAL PPs FREE PPs FREE DISTRIBUTION  
hdiskpower3 active   511      501   103..92..102..102..102
```

Use the varyoffvg command on all affected vgs:

```
dmpvios1$ varyoffvg brunovg
```

Unmanage the EMC Clariion array from PowerPath control

```
# powermt unmanage class=clariion  
hdiskpower0 deleted  
hdiskpower1 deleted  
hdiskpower2 deleted  
hdiskpower3 deleted
```

8 Reboot VIO server1

```
dmpvios1$ reboot
```

- 9 After the VIO server1 reboots, verify that all of the existing volume groups on the VIO server1 and MPIO VTDs on the VIO server1 are successfully migrated to DMP.

```
dmpviosl lsvg -p brunovg
```

```
brunovg:
```

```
PV_NAME          PV STATE TOTAL PPs FREE PPs FREE DISTRIBUTION
emc_clari0_138 active   511      501    103..92..102..102..102
```

Verify the mappings of the LUNs on the migrated volume groups:

```
dmpviosl lsmap -all
```

SVSA	Physloc	Client Partition ID
vhost0	U9117.MMA.0686502-V2-C11	0x00000000
VTD	P0	
Status	Available	
LUN	0x8100000000000000	
Backing device	emc_clari0_130	
Physloc		
VTD	P1	
Status	Available	
LUN	0x8200000000000000	
Backing device	emc_clari0_136	
Physloc		
VTD	P2	
Status	Available	
LUN	0x8300000000000000	
Backing device	emc_clari0_137	
Physloc		

- 10 Repeat step 1 to step 9 for VIO server2.
- 11 Start all of the VIO clients.

Configuring DMP pseudo devices as virtual SCSI devices

DMP in the VIO server supports the following methods to export a device to the VIO client:

- DMP node method
See “Exporting DMP devices as virtual SCSI disks ” on page 60.
- Logical partition-based method
See “Exporting a Logical Volume as a virtual SCSI disk” on page 63.
- File-based method
See “Exporting a file as a virtual SCSI disk” on page 65.

Exporting DMP devices as virtual SCSI disks

DMP supports disks backed by DMP as virtual SCSI disks. Export the DMP device as a vSCSI disk to the VIO client.

To export a DMP device as a vSCSI disk

- 1 Log into the VIO server partition. Use the following command to access the non-restricted root shell. All subsequent commands in this procedure must be invoked from the non-restricted shell.

```
$ oem_setup_env
```

- 2 Use the `oem_setup_env` command to access the non-restricted root shell.
- 3 The following command displays the DMP devices on the VIO server:

```
dmpvios1$ lsdev -t dmpdisk

ibm_ds8000_0287 Available Veritas DMP Device
ibm_ds8000_0288 Available Veritas DMP Device
ibm_ds8000_0292 Available Veritas DMP Device
ibm_ds8000_0293 Available Veritas DMP Device
ibm_ds8000_029a Available Veritas DMP Device
ibm_ds8000_02c1 Available Veritas DMP Device
ibm_ds8000_0337 Available Veritas DMP Device
```

- 4 Assign the DMP device as a backing device. Exit from the non-restricted shell to run this command from the VIOS default shell.

```
dmpvios1$ exit
```

```
$ mkvdev -vdev ibm_ds8000_0288 -vadapter vhost0
vtscsi3 Available
```

5 Use the following command to display the configuration.

```
$ lsmmap -all
```

SVSA	Physloc	Client Partition ID
vhost0	U9117.MMA.0686502-V2-C11	0x00000000
VTD	vtscsi0	
Status	Available	
LUN	0x8100000000000000	
Backing device	ibm_ds8000_0337	
Physloc		
VTD	vtscsi1	
Status	Available	
LUN	0x8200000000000000	
Backing device	ibm_ds8000_02c1	
Physloc		
VTD	vtscsi2	
Status	Available	
LUN	0x8300000000000000	
Backing device	ibm_ds8000_029a	
Physloc	V	
TD	vtscsi3	
Status	Available	
LUN	0x8400000000000000	
Backing device	ibm_ds8000_0288	
Physloc		

6 For a dual-VIOS configuration, export the DMP device corresponding to the same SAN LUN on the second VIO Server in the configuration. To export the DMP device on the second VIO server, identify the DMP device corresponding to the SAN LUN as on the VIO Server1.

- If the array supports the AVID attribute, the DMP device name is the same as the DMP device name on the VIO Server1.
- Otherwise, use the UDID value of the DMP device on the VIO Server1 to correlate the DMP device name with same UDID on the VIO Server2.
On VIO Server1:

```
$ oem_setup_env
```

```
dmpvios1$ lsattr -El ibm_ds8000_0288
```

```
attribute value          description          user_settable
dmpname   ibm_ds8x000_0288 DMP Device name   True
pvid      none              Physical volume identifier True
unique_id IBM%5F2107%5F75MA641%5F6005076308FFC61A0000000000000288
Unique device identifier  True
```

On VIO Server2:

```
$ oem_setup_env
```

```
dmpvios2$ odmgget -q "attribute = unique_id and
value = 'IBM%5F2107%5F75MA641%5F6005076308FFC61A0000000000000288'" CuAt
```

CuAt:

```
name = "ibm_ds8000_0288"
attribute = "unique_id"
value = "IBM%5F2107%5F75MA641%5F6005076308FFC61A000000000000288"
type = "R"
generic = "DU"
rep = "s"
nls_index = 4
```

- 7 Use the DMP device name identified in step 6 to assign the DMP device as a backing device. Exit from the non-restricted shell to run this command from the VIOS default shell.

```
dmpvios1$ exit

$ mkvdev -vdev ibm_ds8000_0288 -vadapter vhost0
vtscsi3 Available
```

- 8 Use the following command to display the configuration.

```
$ lsmmap -all
```

SVSA	Physloc	Client Partition ID
vhost0	U9117.MMA.0686502-V2-C11	0x00000000
VTD	vtscsi0	
Status	Available	
LUN	0x8100000000000000	
Backing device	ibm_ds8000_0337	
Physloc		
VTD	vtscsi1	
Status	Available	
LUN	0x8200000000000000	
Backing device	ibm_ds8000_02c1	
Physloc		
VTD	vtscsi2	
Status	Available	
LUN	0x8300000000000000	
Backing device	ibm_ds8000_029a	
Physloc	V	
TD	vtscsi3	
Status	Available	
LUN	0x8400000000000000	
Backing device	ibm_ds8000_0288	
Physloc		

Exporting a Logical Volume as a virtual SCSI disk

DMP supports vSCSI disks backed by a Logical Volume. Export the Logical Volume as a vSCSI disk to the VIO client.

To export a Logical Volume as a vSCSI disk

1 Create the volume group.

```
$ mkvg -vg brunovg ibm_ds8000_0292 ibm_ds8000_0293  
brunovg
```

The following command displays the new volume group:

```
$ lsvg -pv brunovg  
brunovg:  
PV_NAME          PV STATE TOTAL PPs FREE PPs FREE DISTRIBUTION  
ibm_ds8000_0292 active    494      494      99..99..98..99..99  
ibm_ds8000_0293 active    494      494      99..99..98..99..99
```

2 Make a logical volume in the volume group.

```
$ mklv -lv brunovg_lv1 brunovg 1G  
brunovg_lv1
```

The following command displays the new logical volume:

```
$ lsvg -lv brunovg  
brunovg:  
LV NAME          TYPE   LPs   PPs   PVs  LV STATE      MOUNT POINT  
brunovg_lv1     jfs    256   256   1    closed/syncd  N/A
```

3 Assign the logical volume as a backing device.

```
$ mkvdev -vdev brunovg_lv1 -vadapter vhost0  
vtscsi4 Available
```

4 Use the following command to display the configuration.

```
$ lsmmap -all
```

SVSA	Physloc	Client Partition ID
-----	-----	-----
vhost0	U9117.MMA.0686502-V2-C11	0x00000000
VTD	vtscsi0	
Status	Available	
LUN	0x8100000000000000	
Backing device	ibm_ds8000_0337	
Physloc		
VTD	vtscsi1	
Status	Available	
LUN	0x8200000000000000	
Backing device	ibm_ds8000_02c1	
Physloc		
VTD	vtscsi2	
Status	Available	
LUN	0x8300000000000000	
Backing device	ibm_ds8000_029a	
Physloc		
VTD	vtscsi3	
Status	Available	
LUN	0x8400000000000000	
Backing device	ibm_ds8000_0288	
Physloc		
VTD	vtscsi4	
Status	Available	
LUN	0x8500000000000000	
Backing device	brunovg_lv1	
Physloc		

Exporting a file as a virtual SCSI disk

DMP supports vSCSI disks backed by a file. Export the file as a vSCSI disk to the VIO client.

To export a file as a vSCSI disk

1 Create the storage pool.

```
$ mksp brunospool ibm_ds8000_0296
brunospool
0516-1254 mkvg: Changing the PVID in the ODM.
```

2 Create a file system on the pool.

```
$ mksp -fb bruno_fb -sp brunospool -size 500M
bruno_fb
File system created successfully.
507684 kilobytes total disk space.
New File System size is 1024000
```

3 Mount the file system.

```
$ mount
```

node	mounted	mounted over	vfs	date	options
/dev/hd4	/	jfs2	Jul 02 14:47	rw,log=/dev/hd8	
/dev/hd2	/usr	jfs2	Jul 02 14:47	rw,log=/dev/hd8	
/dev/hd9var	/var	jfs2	Jul 02 14:47	rw,log=/dev/hd8	
/dev/hd3	/tmp	jfs2	Jul 02 14:47	rw,log=/dev/hd8	
/dev/hd1	/home	jfs2	Jul 02 14:48	rw,log=/dev/hd8	
/dev/hd11admin	/admin	jfs2	Jul 02 14:48	rw,log=/dev/hd8	
/proc	/proc	procfs	Jul 02 14:48	rw	
/dev/hd10opt	/opt	jfs2	Jul 02 14:48	rw,log=/dev/hd8	
/dev/livedump	/var/adm/ras/livedump	jfs2	Jul 02 14:48	rw,log=/dev/hd8	
/dev/bruno_fb	/var/vio/storagepools/bruno_fb	jfs2	Jul 02 15:38	rw,log=INLINE	

4 Create a file in the storage pool.

```
$ mkbdsp -bd bruno_fbdev -sp bruno_fb 200M
Creating file "bruno_fbdev" in storage pool "bruno_fb".
bruno_fbdev
```

5 Assign the file as a backing device.

```
$ mkbdsp -sp bruno_fb -bd bruno_fbdev -vadapter vhost0
Assigning file "bruno_fbdev" as a backing device.
vtscsi5 Available
bruno_fbdev
```

6 Use the following command to display the configuration.

```
$ lsmap -all
```

SVSA	Physloc	Client Partition ID
vhost0	U9117.MMA.0686502-V2-C11	0x00000000
...		
...		
VTD	vtscsi5	
Status	Available	
LUN	0x8600000000000000	
Backing device	/var/vio/storagepools/bruno_fb/bruno_fbdev	
Physloc		

Extended attributes in VIO client for a virtual SCSI disk

Using DMP in the Virtual I/O server enables the DMP in the VIO Client to receive the extended attributes for the LUN. This enables the client LPAR to view back-end LUN attributes such as thin, SSD, and RAID levels associated with the vSCSI devices.

For more information about extended attributes and the prerequisites for supporting them, see the following tech note:

<http://seer.entsupport.symantec.com/docs/337516.htm>

Configuration prerequisites for providing extended attributes on VIO client for virtual SCSI disk

DMP in VIO client will provide extended attributes information of backend SAN LUN. The following conditions are prerequisites for using extended attributes on the VIO client:

- VIO client has vSCSI disks backed by SAN LUNs.
- In the VIO Server partition, DMP is controlling those SAN LUNs.
- On VIO client, DMP is controlling the vSCSI disks.

Displaying extended attributes of virtual SCSI disks

When a VIO client accesses a virtual SCSI disk that is backed by a DMP device on the Virtual I/O Server, the VIO client can access the extended attributes associated with the virtual SCSI disk.

The following commands can access and display extended attributes information associated with the vSCSI disk backed by DMP device on Virtual I/O Server.

- `vxdisk -e list`
- `vxdmppadm list dmpnodename=<daname>`
- `vxdmppadm -v getdmpnode dmpnodename=<daname>`
- `vxdisk -p list <daname>`

For example, use the following command on the VIO client `dmpvioc1`:

```
# vxdisk -e list
```

DEVICE	TYPE	DISK	GROUP	STATUS	OS_NATIVE_NAME	ATTR
ibm_ds8x000_114f	auto:LVM	-	-	LVM	hdisk83	std
3pardata0_3968	auto:aixdisk	-	-	online thin	hdisk84	tp

```
# vxdmppadm list dmpnode dmpnodename=3pardata0_3968
```

```
dmpdev          = 3pardata0_3968
state           = enabled
enclosure       = 3pardata0
cab-sno         = 744
asl             = libvxvscsi.so
vid             = AIX
pid            = VDASD
array-name      = 3PARDATA
array-type      = VSCSI
iopolicy        = Single-Active
avid           = 3968
lun-sno         = 3PARdata%5FVV%5F02E8%5F2AC00F8002E8
udid           = AIX%5FVDASD%5F%5F3PARdata%255FVV%255F02E8%255F2AC00F8002E8
dev-attr        = tp
###path        = name state type transport ctlr hwpath aportID aportWWN attr
path           = hdisk84 enabled(a) - SCSI vscsil vscsil 3 - -
```

Virtual IO client adapter settings for Dynamic Multi-Pathing in dual-VIOS configurations

Symantec recommends the following Virtual I/O client (VIO client) adapter settings when using Dynamic Multi-Pathing (DMP) in dual-VIOS configurations:

- Set the `vscsi_err_recov` attribute to `fast_fail`.
The virtual SCSI (vSCSI) adapter driver uses the `vscsi_err_recov` attribute, which is similar to the attribute `fc_error_recov` for physical fibre channel (FC) adapters. When this parameter is set to `fast_fail`, the VIO client adapter sends a `FAST_FAIL` datagram to the VIO server so that the I/O fails immediately, rather than being delayed.
- Enable the `vscsi_path_to` attribute.
This attribute allows the virtual client adapter driver to determine the health of the VIO Server and improve path failover processing. The value of this attribute defines the number of seconds that the vSCSI client adapter waits for commands sent to the vSCSI server adapter to be serviced. If that time is exceeded, the vSCSI client adapter fails the outstanding requests. If DMP is present, another path to the disk will be tried to service the requests. A value of 0 (default) disables this functionality.

To set the VIO client adapter settings

- 1 Set the `vscsi_err_recov` attribute to `fast_fail`, and the `vscsi_path_to` attribute to a non-zero number. For example:

```
# chdev -a vscsi_err_recov=fast_fail -a vscsi_path_to=30 -l vscsi0
```

- 2 Verify the settings.

```
# lsattr -El vscsi0
vscsi_err_recov      fast_fail
vscsi_path_to        30
```

- 3 Repeat step 1 and step 2 for each vSCSI client adapter.

Storage Foundation and High Availability Virtualization Solutions for IBM LPARs with N_Port ID Virtualization

This chapter includes the following topics:

- [About IBM LPARs with N_Port ID Virtualization \(NPIV\)](#)
- [Support for Storage Foundation in NPIV environment](#)
- [Installation, patching, and configuration requirements](#)

About IBM LPARs with N_Port ID Virtualization (NPIV)

N_Port ID Virtualization or NPIV is a Fibre Channel industry standard technology that allows multiple N_Port IDs to share a single physical N_Port. NPIV provides the capability to take a single physical Fibre Channel HBA port and divide it such that it appears, to both the host and to the SAN, as though there are multiple World Wide Port Names (WWPNs).

NPIV provides direct access to the Fibre Channel adapters from multiple virtual machine (client partitions), simplifying zoning and storage allocation. Resources can be zoned directly to the virtual client, which has its own World Wide Port Name (WWPN).

The use of NPIV with IBM VIO provides the capability to use a single Fibre Channel port and overlay multiple WWPNs so that it appears to the SAN as both the VIO server and client partitions. NPIV enables the AIX VIO server to provision entire dedicated logical ports to client LPARs rather than individual LUNs. Client partitions with this type of logical port operates as though the partition has its own dedicated FC protocol adapter. To utilize the NPIV functionality, a new type of virtual Fibre Channel (VFC) adapter is defined on both the VIO and Client. A server VFC adapter can only be created on a VIO server partition; a client VFC adapter can only be created on client partitions. WWPNs are allocated to client VFC adapters when they are defined in the profile, based upon an assignment pool generated from the backing physical adapter.

There is always corresponding one-to-one mapping relationship between VFC adapters on client logical partitions and VFC on the VIOS. That is, each VFC that is assigned to a client logical partition must connect to only one VFC adapter on VIOS, and each VFC on VIOS must connect to only one VFC on the client logical partition.

NPIV support is included with PowerVM Express, Standard, and Enterprise Edition and supports AIX 6.1 and AIX 7.1.

More details on NPIV and how to configure IBM VIO environment is available in IBM documentation.

Characteristics of a LUN through NPIV

To the operating system, multi-pathing drivers and system tools, a LUN presented through NPIV has all the characteristics of a LUN presented through a dedicated HBA. Device inquiry and probing works as with physical HBAs. When a VFC interface is created, two World Wide Port Names (WWPNs) are assigned. This information is available in the HMC as part of the virtual HBA properties.

All SCSI device inquiry operations work, allowing for array identification functions, visibility of LUN Device Identifiers, and discovery of such attributes as thin and thin re-claim capability. SCSI-3 persistent reservation functionality is also supported, enabling the use of SCSI-3 I/O Fencing if the underlying storage supports.

When Zoning/LUN mapping operations occur, care should be made to ensure that storage is assigned to both WWPNs. During normal operation, only one of the WWPN identifiers is in use, but during a Live Partition migration event, the WWPN identifier not previously used will be configured on the appropriate backing HBA on the target system, log into the SAN, and then become the active WWPN. The previously used WWPN will become inactive until the next Live Partition Mobility operation.

VIO requirements

NPIV requires Power6 systems, VIOS 2.1, and 8GB HBA adapters (model number xxxxx). NPIV also requires NPIV aware switches. The end storage devices need not be NPIV aware.

Hardware requirements

NPIV requires extended functionality on the HBA. Currently IBM sells this as an 8GB HBA, part number XXXXX. The SAN Switch ports must also support NPIV as well, Brocade and Cisco make products that provide this functionality.

Support for Storage Foundation in NPIV environment

Starting with Storage Foundation 5.0MP3 RP1, Storage Foundation supports NPIV in IBM Virtual I/O Server (VIOS) environments. The VIOS is configured with NPIV capable FC adapters that are connected to a SAN switch that is NPIV capable. The LUNs mapped to the VIO client behave like an LPAR having a dedicated FC adapter. The devices in the VIO client appear as regular SCSI disks. Storage Foundation can access these LUNS, and treat these devices as if they came from a regular SAN storage array LUN. Unlike in the classic VIO environment without NPIV, SF treats these devices as if they came from a regular SAN storage array LUN. With NPIV, the VIO client environment is transparent to SF. All of the SF commands would have the same output as in a regular physical AIX server. SF identifies the vSCSI LUNs through the array properties of the LUNs. Otherwise, the devices in the VIO client appear as regular SCSI disks. You can import the disk group, which provides access to volumes and file systems.

Symantec has qualified NPIV support with Storage Foundation, starting with 5.0MP3 RP1.

Symantec has also qualified migration of storage used by Storage Foundation from the AIX physical server environment to the IBM VIO environment.

See [“Migrating from Physical to VIO environment”](#) on page 104.

Storage Foundation

Storage Foundation 6.0 supports all functionality available with dedicated HBAs when using LUNs presented through NPIV. All IBM supported NPIV enabled HBAs are supported by Storage Foundation.

Storage Foundation functionality is fully supported with NPIV.

Cluster File System

Cluster File System is supported with NPIV.

Veritas Cluster Server

VCS supports NPIV. With NPIV, the VIOS client environment is transparent to VCS and the LUNs are treated as regular SAN storage array LUNs. Since SCSI3 persistent reserve is available, I/O fencing is also supported.

Installation, patching, and configuration requirements

Symantec strongly recommends that you use Storage Foundation 6.0 with the latest patches. No other configuration is required. Refer to the following website for the latest patches for Storage Foundation 6.0 on AIX:

<https://sort.symantec.com/checklist/install/>

Storage Foundation and High Availability support for Live Partition Mobility

This chapter includes the following topics:

- [About Live Partition Mobility \(LPM\)](#)
- [SFHA supported configuration](#)
- [SFHA requirements for the Live Partition Mobility](#)
- [About VCS support for Live Partition Mobility](#)
- [Overview of partition migration process](#)
- [Performance considerations](#)

About Live Partition Mobility (LPM)

The Live Partition Mobility enables you to migrate an entire logical partition from one physical system to another. Live Partition Mobility transfers the configuration from source to destination without disrupting the hosted applications or the setup of the operating system and applications. Live Partition Mobility gives you a greater control over the usage of resources in the data center.

Live Partition Mobility allows a level of reconfiguration that in the past was not possible due to complexity or because of service level agreements that do not allow an application to be stopped for an architectural change. The migration process can be performed in the following ways:

- Inactive migration

The logical partition is powered off and moved to the destination system.

- **Active migration**

The migration of the partition is performed while service is provided, without disrupting user activities. During an active migration, the applications continue to handle their normal workload. Disk data transactions, running network connections, user contexts, and the complete environment are migrated without any loss and migration can be activated any time on any production partition.

SFHA supported configuration

All SFHA stacks support LPM including Fencing configured with NPIV disks.

Note: Please check the IBM documentation for the detailed information on the LPM requirements and LPM process.

SFHA requirements for the Live Partition Mobility

The main requirements for the migration of a logical partition are:

System requirements

- Two POWER6 or POWER7 systems controlled by the same Hardware Management Console (HMC).
- The destination system must have enough CPU and memory resources to host the mobile partition.

Network requirements

- The migrating partition must use the virtual LAN (VLAN) for all LLT links and public network access. The VLAN must be bridged to a physical network using a shared Ethernet adapter in the Virtual I/O Server partition. If there is more than one VLAN, each VLAN must be bridged.
- The Virtual I/O Servers on both systems must have a shared Ethernet adapter configured to bridge to the same Ethernet network used by the mobile partition.
- Your LAN must be configured such that migrating partitions can continue to communicate with the other nodes after a migration is completed.

Storage requirements

- The operating system, applications, and data of the mobile partition must reside on virtual storage on an external storage subsystem since the mobile partition's disk data must be available after the migration to the destination system is completed. An external, shared access storage subsystem is required.
- The mobile partition's virtual disks must be mapped to LUNs; they cannot be part of a storage pool or logical volume on the Virtual I/O Server. The LUNs must be zoned and masked to the Virtual I/O Servers on both systems.

About VCS support for Live Partition Mobility

You can use Live Partition Mobility to perform a stateful migration of an LPAR 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, 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.

For more information, refer to the *Veritas Cluster Server Administrator's Guide*.

Overview of partition migration process

The partition migration, either inactive or active, is divided into the following stages:

- Preparing the infrastructure to support Live Partition Mobility.
- Checking the configuration and readiness of the source and destination systems.
- Transferring the partition state from the source to destination. The same command is used to launch inactive and active migrations. The HMC determines the appropriate type of migration to use based on the state of the mobile partition.
- Completing the migration by freeing unused resources on the source system and the HMC.

Performance considerations

Active partition migration involves moving the state of a partition from one system to another while the partition is still running. The mover service partitions working with the hypervisor use partition virtual memory functions to track changes to partition memory state on the source system while it is transferring memory state to the destination system.

During the migration phase, there is an initial transfer of the mobile partition's physical memory from the source to the destination. Since the mobile partition is still active, a portion of the partition's resident memory will almost certainly have changed during this pass. The hypervisor keeps track of these changed pages for retransmission to the destination system in a dirty page list. It makes additional passes through the changed pages until the mover service partition detects that a sufficient number of pages are clean or the timeout is reached. The speed and load of the network that is used to transfer state between the source and destination systems influence the time that is required for both the transfer of the partition state and the performance of any remote paging operations. The amount of changed resident memory after the first pass is controlled more by write activity of the hosted applications than by the total partition memory size. Nevertheless, it is reasonable to assume that partitions with a large memory requirement will have higher numbers of changed resident pages than smaller ones.

To ensure that active partition migrations are truly non-disruptive, even for large partitions, the POWER Hypervisor resumes the partition on the destination system before all the dirty pages have been migrated over to the destination. If the mobile partition tries to access a dirty page that has not yet been migrated from the source system, the hypervisor on the destination sends a demand paging request to the hypervisor on the source to fetch the required page.

Providing a high-performance network between the source and destination mover partitions and reducing the partition's memory update activity before migration will improve the latency of the state transfer phase of migration. We suggest using a dedicated network for state transfer, with a bandwidth of at least 1 Gbps.

Storage Foundation and High Availability support for IBM Workload Partitions

This chapter includes the following topics:

- [About IBM Workload Partitions](#)
- [When to use WPARs](#)
- [Storage Foundation support for WPARs](#)
- [WPAR mobility](#)
- [About VCS support for WPARs](#)
- [About configuring VCS in WPARs](#)
- [Configuring AIX WPARs for disaster recovery using VCS](#)

About IBM Workload Partitions

IBM Workload Partitions (WPARs) are implemented starting with AIX 6.1. Workload Partitions allow administrators to virtualize the AIX operating system, by partitioning an AIX operating system instance into multiple environments. Each environment within the AIX operating system instance is called a workload partition (WPAR). One WPAR can host applications and isolate the applications from applications executing in other WPARs. WPAR is a pure software solution and has no dependencies on hardware features.

The WPAR solution allows for fewer operating system images on your IBM System p partitioned server. Prior to WPARs, you had to create a new Logical Partition

(LPAR) for each new "isolated" environment. Starting with AIX 6.1, you can instead use multiple WPARs within one LPAR, in many circumstances.

In an LPAR environment, each LPAR requires its own operating system image and a certain number of physical resources. While you can virtualize many of these resources, some physical resources must be allocated to the system for each LPAR. Furthermore, you need to install patches and technology upgrades to each LPAR. Each LPAR requires its own archiving strategy and DR strategy. It also takes some time to create an LPAR; you also need to do this outside of AIX, through a Hardware Management Console (HMC) or the Integrated Virtualization Manager (IVM).

In contrast, WPARs are much simpler to manage and can be created from the AIX command line or through SMIT. WPARs allow you to avoid the biggest disadvantage of LPARs: maintaining multiple images, and therefore possibly over-committing expensive hardware resources, such as CPU and RAM. While logical partitioning helps you consolidate and virtualize hardware within a single box, operating system virtualization through WPAR technology goes one step further and allows for an even more granular approach of resource management.

The WPAR solution shares operating system images and is clearly the most efficient use of CPU, RAM, and I/O resources. Rather than a replacement for LPARs, WPARs are a complement to them and allow one to further virtualize application workloads through operating system virtualization. WPARs allow for new applications to be deployed much more quickly.

WPARs have no real dependency on hardware and can even be used on POWER4 systems that do not support IBM's PowerVM (formerly known as APV). For AIX administrators, the huge advantage of WPARs is the flexibility of creating new environments without having to create and manage new AIX partitions.

On the other hand, it's important to understand the limitations of WPARs. For example, each LPAR is a single point of failure for all WPARs that are created within the LPAR. In the event of an LPAR problem (or a scheduled system outage), all underlying WPARs are also affected.

The following sections describe the types of WPARs:

- System workload partition
- Application workload partition

System workload partition

The system WPAR is much closer to a complete version of AIX. The system WPAR has its own dedicated, completely writable file-systems along with its own inetd and cron. You can define remote access to the System workload partition.

Application workload partition

Application WPARs are lightweight versions of virtualized OS environments. They are extremely limited and can only run application processes, not system daemons such as `inetd` or `cron`. You cannot even define remote access to this environment. These are only temporarily objects; they actually disintegrate when the final process of the application partition ends, and as such, are more geared to execute processes than entire applications.

When to use WPARs

You can use WPARs when you need an isolated environment, especially if you do not want to create new LPARs because of the limitation of the available resources. Here are a few recommended scenarios:

- Application/workload isolation
- Quickly testing an application

WPARs share the global resources with other WPARs in the same LPAR, which limits the usefulness of WPARs in some situations.

We recommend not using WPARs in the following situations:

- Security: WPAR processes can be seen by the global environment from the central LPAR. If you are running a highly secure type of system, this may be a problem for you from a security standpoint. Further, the root administrator of your LPAR will now have access to your workload partition, possibly compromising the security that the application may require.
- Performance: Each WPAR within the LPAR uses the same system resources of the LPAR. You need to be more careful when architecting your system and also when stress testing the system.
- Availability: If you are in an environment where it is very difficult to bring a system down, it's important to note that when performing maintenance on an LPAR that every WPAR defined will be affected. At the same time, if there is a system panic and AIX crashes, every WPAR has now been brought down.
- Production: Avoid using WPARs in the production environment. LPARs provide more granularity and complete OS isolation.
- Physical devices: Physical devices are not supported within a WPAR. More details on WPAR administration can be found in the IBM red book on WPARs at

<http://www.redbooks.ibm.com/abstracts/sg247431.html>

Storage Foundation support for WPARs

This section describes Veritas File System (VxFS) support for workload partitions (WPARs). Currently, there are no VxVM operations available within a system WPAR, so any VxFS file system that is needed for data use must be created in the global environment, then set up so the WPAR can access it. The VxFS (local mount only) is supported inside the workload partition (WPAR) environment. Cluster mount is not yet supported inside a WPAR. WPAR can have both root and non-root partitions as VxFS file system.

In Storage Foundation, there is limited support for WPARs, as follows:

- All the Storage Foundation packages must be installed and configured in the global partition of AIX.
- Storage Foundation can only be administered from the global partition.

There are two ways to use a local mount VxFS file system inside WPAR environment.

- Using a VxFS file system within a single system WPAR
- Using VxFS as a shared file system

Using a VxFS file system within a single system WPAR

The following procedure describes how to set up a WPAR with VxFS for non-root partition.

To set up WPAR with VxFS for non-root partition

- 1 Create a vxfs filesystem in the global environment:

```
# /opt/VRTS/bin/mkfs -V vxfs /dev/vx/rdisk/testvg/vol1
```

- 2 Create a WPAR. For example, use the following command.

```
# mkwpar -n devpayrollWPAR01
```

For other options while creating WPARs, refer to the IBM Redbook for WPAR.

- 3 List the WPAR.

```
# lswpar
Name                State Type Hostname                Directory
-----
devpayrollWPAR01 D      S      devpayrollWPAR01 /wpars/devpayrollWPAR01
```

- 4 The above output shows that WPAR does not have the devices. To get the vxfs file system in WPAR, create the file system in the global environment. Then mount it to the WPAR directories which are located at `/wpar/wparname/`

```
# mkdir /wpars/devpayrollWPAR01/vxfs_dir
# mount -V vxfs /dev/vx/dsk/testdg/vol1 /wpars/devpayrollWPAR01/vxfs_dir
```

- 5 Start the WPAR:

```
# startwpar -Dv devpayrollWPAR01 2>/startwpar_t12
```

- 6 Log in to the WPAR.

```
# clogin hostname
```

For example, to log in to the WPAR devpayrollWPAR01:

```
# clogin devpayrollWPAR01
```

- 7 The following output shows the VxFS mount point in the WPAR.

```
# mount
```

node	mounted	mounted over	vfs	date	options
Global	/		jfs2	Jun 23 03:15	rw,log=INLINE
Global	/home		jfs2	Jun 23 03:15	rw,log=INLINE
Global	/opt		namefs	Jun 23 03:15	ro
Global	/proc		namefs	Jun 23 03:15	rw
Global	/tmp		jfs2	Jun 23 03:15	rw,log=INLINE
Global	/usr		namefs	Jun 23 03:15	ro
Global	/var		jfs2	Jun 23 03:15	rw,log=INLINE
Global	/vxfs_dir		vxfs	Jun 23 03:14	rw,delaylog, suid,ioerror=mwdisable,qio,largefiles

- 8 To stop the WPAR, use the following command:

```
# stopwpar -Dv devpayrollWPAR01 2>/devpayrollWPAR01_t12
```

WPAR with root (/) partition as vxfs

Starting with the Storage Foundation 5.1SP1 release, the / (root) partition of any WPAR can be created as vxfs. Previous to the 5.1SP1 release, it was mandatory

to have the / partition as JFS2. Other mount points appear as before but / can be vxfs.

- 1 Create the / (root) partition of the WPAR as vxfs.

```
# mkwpar -n fsqawpar -M directory=/ dev=/dev/vx/rdisk/testdg/vol2 vfs=vxfs
```

- 2 Start the WPAR.

```
# startwpar -v fsqawpar 2>/fsqawpar_t12
```

- 3 Login to the WPAR.

```
# clogin fsqawpar
```

- 4 Other mount points appear as before but / can be vxfs.

```
# mount
node   mounted  mounted over  vfs    date           options
-----
Global /          vxfs    Jun 23 03:30 rw, delaylog,
suid,ioerror=mwdisable,qio,largefiles
Global /home     jfs2    Jun 23 03:30 rw,log=INLINE
Global /opt      namefs  Jun 23 03:30 ro
Global /proc    namefs  Jun 23 03:30 rw
Global /tmp     jfs2    Jun 23 03:30 rw,log=INLINE
Global /usr     namefs  Jun 23 03:30 ro
Global /var     jfs2    Jun 23 03:30 rw,log=INLINE
```

Using VxFS as a shared file system

VxFS is also supported as “namefs” in the WPAR, so a VxFS file system can also be shared between the global environment and WPARs.

To use VxFS as a shared file system

- 1 Mount vxfs on some directory in the global environment.

```
# mount -V vxfs /dev/vx/dsk/testdg/vol1 /mnt
```

- 2 Mount that directory in /wpar/ devpayrollWPAR01/vxfs_dir.

```
# mount /mnt /wpars/devpayrollWPAR01/vxfs_dir/
```

3 Start the WPAR.

```
# startwpar -Dv devpayrollWPAR01 2>/devpayrollWPAR01_t12
```

4 Login to the WPAR.

```
# clogin devpayrollWPAR01
```

5 After login to devpayrollWPAR01, /vxfs_dir will appear as namefs.

```
# mount
node      mounted mounted over vfs      date              options
-----
Global /          jfs2   Jun 23 03:30 rw,log=INLINE
Global /home      jfs2   Jun 23 03:30 rw,log=INLINE
Global /opt       namefs Jun 23 03:30 ro
Global /proc      namefs Jun 23 03:30 rw
Global /tmp       jfs2   Jun 23 03:30 rw,log=INLINE
Global /usr       namefs Jun 23 03:30 ro
Global /var       jfs2   Jun 23 03:30 rw,log=INLINE
Global /vxfs_dir  namefs Jun 23 03:29 rw
```

WPAR mobility

Live application mobility allows for planned migrations of workload from one system to another without interrupting the application. This technology can be used to perform a planned firmware installation on the server. Most workloads do not need to be aware of the WPAR relocation. WPAR mobility, also referred to as relocation, applies to both types of WPARs: application and system. The relocation of a WPAR consists of moving its executable code from one LPAR to another one while keeping the application data on the same storage devices. It is therefore mandatory that these storage devices are accessible from both the source and target LPARs hosting the WPAR. The hosting global environment hides the physical and logical device implementations from the hosted WPARs. The WPAR only works with data storage at the file system level. All files that need to be written by the application must be hosted on an NFS file system. All other files, including the AIX operating system files, can be stored in file systems local to the hosting global environment. The NFS server must provide access to both the global environment and the WPAR in order for the WPAR to work at all. In a mobility scenario, access must be provided to the WPAR and all global environments to which the WPAR might be moved.

About VCS support for WPARs

VCS provides application management and high availability to applications that run in WPARs. VCS supports only system WPARs, application WPARs are not supported.

Overview of how VCS works with WPARs

You can use VCS to perform the following:

- Start, stop, monitor, and failover a local WPAR.
- Start, stop, monitor, and failover an application that runs in a WPAR.

Installing and configuring WPARs in VCS environments

Install and configure the WPAR. Create the service group with the standard application resource types (application, storage, networking) that need to be run inside the WPAR, and the WPAR resource. VCS represents the WPAR and its state using the WPAR resource. You then configure the service group's ContainerInfo attribute.

Configuring the ContainerInfo attribute

The service group attribute ContainerInfo specifies information about the WPAR. When you have configured and enabled the ContainerInfo attribute, you have enabled the WPAR-aware resources in the service group to work in the WPAR environment. VCS defines the WPAR information at the level of the service group so that you do not have to define it for each resource. You can specify a per-system value for the ContainerInfo attribute.

Running VCS, its resources, and your applications

VCS and the necessary agents run in the global environment. For applications that run in a WPAR, the agents can run some of their functions (entry points) inside the WPAR. If any resource faults, VCS fails over the service group with the WPAR to another node.

The ContainerInfo attribute

The ContainerInfo attribute has the Name key, Type key, and Enabled key. The Name key defines the name of the WPAR. The Type key lets you select the type of container that you plan to use (WPAR or Zone). The Enabled key enables the WPAR-aware resources within the service group. To configure the ContainerInfo attribute, use the `hawparsetup.pl` command.

You can specify a per-system value for the ContainerInfo attribute. For more information, refer to the *Veritas Cluster Server Administrator's Guide*.

The ContainerOpts attribute

The ContainerOpts attribute has the RunInContainer key and PassCInfo key. If the resource type has the RunInContainer and PassCInfo keys defined in ContainerOpts, the resource type is WPAR-aware. WPAR-aware indicates that VCS can monitor and control a resource of that type inside a WPAR.

The ContainerOpts attribute determines the following:

- The RunInContainer key determines whether the entry points of a WPAR-aware resource type can run in the WPAR.
- The PassCInfo key determines whether the container information is passed to the entry points of the resource type. The container information is defined in the service group's ContainerInfo attribute. An example use of the PassCInfo key is to pass the agent the name of the WPAR.

For more information, refer to the *Veritas Cluster Server Administrator's Guide*.

Note: Symantec recommends that you do not modify the value of the ContainerOpts attribute, with the exception of the Mount agent.

WPAR-aware resource types

The following are the ContainerOpts attribute default values for resource types. WPAR-aware resources have predefined default values for the ContainerOpts attribute.

Table 8-1 ContainerOpts attribute default values for resource types

Resource Type	RunInContainer	PassCInfo
Application	1	0
DB2	1	0
IP	0	1
IPMultiNICB	0	1
Netlsnr	1	0
Mount	0	0
Oracle	1	0

Table 8-1 ContainerOpts attribute default values for resource types (*continued*)

Resource Type	RunInContainer	PassCInfo
Process	1	0
WPAR	0	1

About the Mount agent

You may need to modify the ContainerOpts values for the Mount resource in certain situations. Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.

About the WPAR agent

The WPAR agent monitors WPARs, brings WPARs online, and takes them offline.

For more information about the agent, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

The agent requires a user account with group administrative privileges to enable communication between the global environment and the WPAR. To create a user account, use the `hawparsetup.pl` command to configure the service group.

See [“Configuring the service group for the application”](#) on page 93.

In secure clusters, the agent renews the authentication certificate before the certificate expires.

About configuring VCS in WPARs

Configuring VCS in WPARs involves the following tasks:

- Review the prerequisites.
See [“Prerequisites for configuring VCS in WPARs”](#) on page 89.
- Decide on the location of the WPAR root, which is either on local storage or NFS. The WPAR root is the topmost directory in a section of the file system hierarchy in which the WPAR is configured.
See [“Deciding on the WPAR root location”](#) on page 90.
- Install the application in the WPAR.

See “Installing the application” on page 93.

- Create the application service group and configure its resources.
See “Configuring the service group for the application” on page 93.

Prerequisites for configuring VCS in WPARs

- In a WPAR configuration, all nodes that host applications must run the same version of the operating system.
- The WPAR root must be installed on JFS, JFS2, NFS, or VxFS.
- Mounts must meet one of the following two conditions:
 - Use a namefs file system. All mounts that the application uses must be part of the WPAR configuration and must be configured in the service group. For example, you can create a WPAR, `w_ora`, and define the file system containing the application’s data to have the mount point as `/oradata`. When you create the WPAR, you can define a path in the global environment, for example `/export/home/oradata`, which maps to the mount directory in the WPAR. The `MountPoint` attribute of the `Mount` resource for the application is set to `/export/home/oradata`.
 - Use a direct mount file system. All file system mount points that the application uses that run in a WPAR must be set relative to the WPAR’s root. For example, if the Oracle application uses `/oradata`, and you create the WPAR with the WPAR path as `/w_ora`, then the mount must be `/w_ora/oradata`. The `MountPoint` attribute of the `Mount` resource must be set to this path.

For more information about how to configure `Mount` resource inside WPAR, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

About using custom agents in WPARs

- If you use custom agents to monitor applications running in WPARs, make sure the agents use script-based entry points. VCS does not support running C++ entry points inside a WPAR.
- If the custom agent monitors an application that runs in a WPAR, add the resource type to the `APP_TYPES` environment variable. If the custom agent monitors an application running in the global environment, add the resource type to the `SYS_TYPES` environment variable.
- If you want the custom agent to monitor an application in the WPAR, for the custom agent type, set the following values for the `ContainerOpts` attribute: `RunInContainer = 1` and the `PassCInfo = 0`.

- If you do not want the custom agent to monitor an application in the WPAR, for the custom agent type, set the following values for the ContainerOpts attribute: RunInContainer = 0 and the PassCInfo= 0.

Deciding on the WPAR root location

Each WPAR has its own section of the file system hierarchy in the WPAR root directory. Processes that run in the WPAR can access files only within the WPAR root.

You can set the WPAR root in the following two ways:

- WPAR root on local storage.
In this configuration, you must create a WPAR on each node in the cluster.
- WPAR root on NFS.
In this configuration, create a WPAR on the NFS storage. You need to duplicate the WPAR configuration across all the nodes in the cluster.
When you set the WPAR root on NFS, install the WPAR from one node only. The WPAR root can fail over to the other nodes in the cluster. The system software, including the patches, must be identical on each node during the existence of the WPAR.

Creating a WPAR root on local disk

Use the following procedure to create a WPAR root on the local disk on each node in the cluster.

To create a WPAR root on local disks on each node in the cluster

- 1 Create the actual WPAR root directory.
- 2 Use the `mkwpar` command to create the WPAR.

```
mkwpar -n wpar -h host -N ip_info -d wroot -o /tmp/wpar.log
```

Use the following information to replace the appropriate variables:

<code>wpar</code>	The name of the WPAR.
<code>host</code>	The hostname for the WPAR being created.
<code>ip_info</code>	<p>The information to set the virtual IP address of the system to be the IP address of the WPAR. This value also defines the device name for the NIC associated with the IP address.</p> <p>If you do not specify the value of the interface or netmask, the global partition's values are used.</p> <p>Use the following format to replace <code>ip_info</code>:</p> <pre>interface=<i>interface</i> netmask=<i>netmask</i> address=<i>IPaddress</i></pre> <p>Example: <code>interface='en0' address='172.16.0.0'</code> <code>netmask='255.255.255.0'</code></p>
<code>wroot</code>	The location of the WPAR root directory. For example: <code>/wpar1</code> .

- 3 Repeat the command in step 2 to create the WPAR on each system in the service group's SystemList.
- 4 On one of the systems in the SystemList, mount the shared file system containing the application data.
- 5 Start the WPAR.

Creating WPAR root on shared storage using NFS

Use the following procedure to create a WPAR root on shared storage using NFS.

To create WPAR root on shared storage using NFS

- 1 Create a file system on NFS storage for the WPAR root. The file system that is to contain the WPAR root may be in the same file system as the file system containing the shared data.
- 2 Type the following `mkwpar` command to create the WPAR:

```
mkwpar -n wpar -h host -N ip_info -r -M r_fs -M v_fs -M h_fs -M  
t_fs -d wroot
```

Use the following information to replace the appropriate variables:

Attribute Description

`wpar` The name of the WPAR.

`host` The hostname of the WPAR being created.

`ip_info` The information to set the virtual IP address of the system to be the IP address of the WPAR. This value also defines the device name for the NIC associated with the IP address. Use the following format to replace `ip_info`:

```
interface=interface netmask=netmask address=IPaddress
```

For example: `interface='en0' address='172.16.0.0'
netmask='255.255.255.0'`

If you do not specify the value of the interface or netmask, the global partition's values are used.

`r_fs` The information to specify the NFS volume to use for the root private file system for the WPAR. For example:

```
directory=/ vfs=nfs host=host123 dev=/root01
```

`v_fs` The information to specify the NFS volume to use for the `/var` private file system for the WPAR. For example:

```
directory=/ var vfs=nfs host=host123 dev=/var01
```

`h_fs` The information to specify the NFS volume to use for the `/home` private file system for the WPAR. For example:

```
directory=/home vfs=nfs host=host123 dev=/home01
```

`t_fs` The information to specify the NFS volume to use for the `/tmp` private file system for the WPAR. For example:

```
directory=/tmp vfs=nfs host=host123 dev=/tmp01
```

`wroot` The location of the WPAR root directory, for example, `/wpar1`.

- 3 Use the `lswpar` command to display information about the WPAR's properties and their values.
- 4 On the system where you created the WPAR, run the command:

```
mkwpar -w -o config_file_name -e wparname_just_created
```
- 5 On all the other systems copy the configuration file, run the command:

```
mkwpar -p -f config_file_name -n wparname_just_created
```
- 6 List the WPAR.
- 7 Start the WPAR.
- 8 On one system, mount the shared file system containing the application data.
- 9 Make sure the WPAR created from the first system is in the D state on all other systems in the service group's System List.

Installing the application

Install the application in the WPAR. Perform the following:

- If you have created WPARs on each node in the cluster, install the application identically on all nodes. If you are installing an application that supports a Veritas High Availability agent, see the installation and configuration guide for the agent.
- Install the agent. Agent packages are installed in the global environment and the currently existing WPARs. The operating system installs the agents in future WPARs when they are created.
- In the WPAR, configure all mount points used by the application.
 - If you use `namefs` mounts, verify the global directories are properly mounted inside the WPAR.
 - If you use a direct mount, verify the mount points used by the application have been mounted relative to the WPAR's root. For example, if a WPAR `w_ora` needs to use `/oracle`, mount the drive at `/wpars/w_ora/oracle`.

Configuring the service group for the application

The following diagrams illustrates different examples of resource dependencies. In one case the WPAR root is set up on local storage. In the other, WPAR root is set up on shared storage.

Figure 8-1 WPAR root on local disks (with direct mount file system)

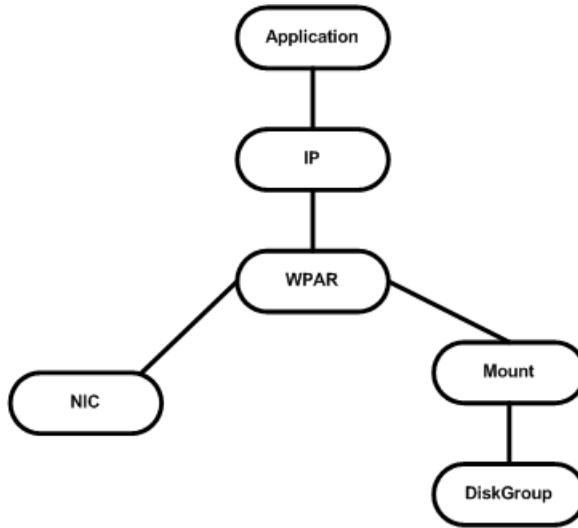


Figure 8-2 WPAR root on local disks (file system mounted from inside WPAR)

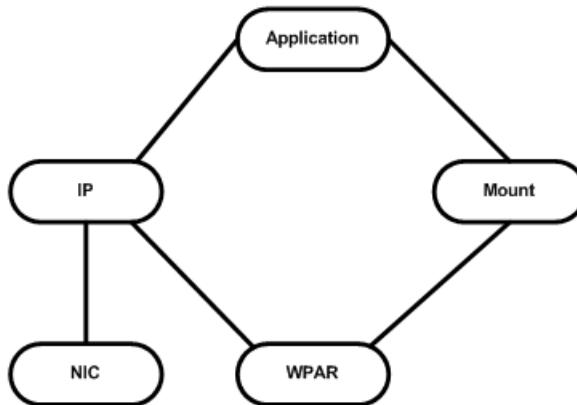
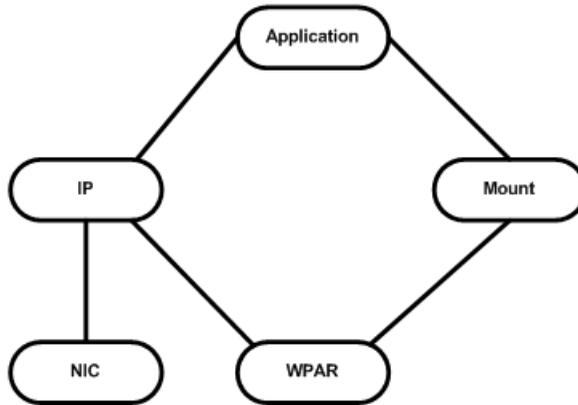


Figure 8-3 WPAR root on shared storage (with namefs file system)



Modifying the service group configuration

Perform the following procedure to modify a service group's configuration.

To modify the service group configuration

- 1 Run the `hawparsetup.pl` command to set up the WPAR configuration.

```
# /opt/VRTSvcs/bin/hawparsetup.pl servicegroup_name WPARres_name WPAR_name
```

<i>servicegroup_name</i>	Name of the application service group.
<i>WPARres_name</i>	Name of the resource configured to monitor the WPAR.
<i>WPAR_name</i>	Name of the WPAR.
<i>password</i>	Password to be assigned to VCS or Security (Symantec Product Authentication Service) user created by the command.
<i>systems</i>	List of systems on which the service group will be configured. Use this option only when creating the service group.

The command adds a resource of type WPAR to the application service group. It also creates a user account with group administrative privileges to enable inter-WPAR communication.

If the application service group does not exist, the command creates a service group with a resource of type WPAR.

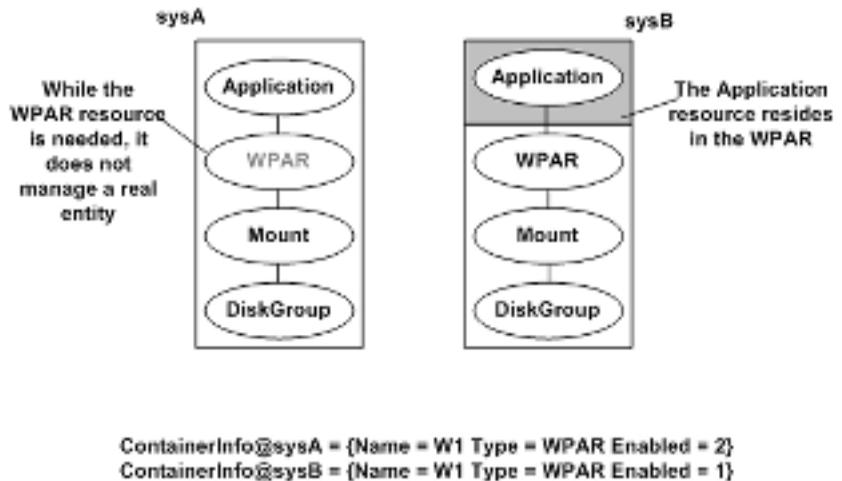
- 2 Modify the resource dependencies to reflect your WPAR configuration. See the resource dependency diagrams for more information.
- 3 Save the service group configuration and bring the service group online.

About configuring failovers among physical and virtual servers

You can configure VCS to fail over from a physical system to a virtual system and vice versa. A physical to virtual failover gives an N + N architecture in an N + 1 environment. For example, several physical servers with applications can fail over to containers on another physical server.

Configuring for failovers—a typical setup

In this configuration, you have two physical nodes. One node runs AIX 7.1 (sysA) and another node that runs AIX 6.1 (sysB). The node that runs AIX 6.1 has WPARs configured.

Figure 8-4 An application service group that can fail over onto a WPAR

In the `main.cf` configuration file, define the container name, type of container, and whether it is enabled or not. The following is an example of the `ContainerInfo` lines in the `main.cf` file:

```
ContainerInfo@sysA = {Name = W1, Type = WPAR, Enabled = 2}
ContainerInfo@sysB = {Name = W1, Type = WPAR, Enabled = 1}
```

On `sysA`, you set the value of `Enabled` to 2 to ignore WPARs so that the application runs on the physical system. When `sysA` fails over to `sysB`, the application runs inside the WPAR after the failover because `Enabled` is set to 1 on `sysB`. The application can likewise fail over to `sysA` from `sysB`.

IMF must be disabled on the node where `Enabled` is set to 2 (`sysA` in this example). To disable IMF, set the mode to 0.

Verifying the WPAR configuration

Run the `hawparverify.pl` command to verify the WPAR configuration. The command verifies the following requirements:

- The systems hosting the service group have the required operating system to run WPARs.
- The service group does not have more than one resource of type WPAR.
- The dependencies of the WPAR resource are correct.

To verify the WPAR configuration

- 1 If you use custom agents make sure the resource type is added to the APP_TYPES or SYS_TYPES environment variable.
See “[About using custom agents in WPARs](#)” on page 89.
- 2 Run the `hawparverify.pl` command to verify the WPAR configuration.

```
# /opt/VRTSvcs/bin/hawparverify servicegroup_name
```

Maintenance tasks

Perform the following maintenance tasks when you use WPARs:

- Whenever you make a change that affects the WPAR configuration, you must run the `hawparsetup` command to reconfigure the WPARs in VCS.
See “[Configuring the service group for the application](#)” on page 93.
- Make sure that the WPAR configuration files are consistent on all the nodes at all times.
- When you add a patch or upgrade the operating system on one node, make sure to upgrade the software on all nodes.
- Make sure that the application configuration is identical on all nodes. If you update the application configuration on one node, apply the same updates to all nodes.

Troubleshooting information

Symptom	Recommended Action
VCS HA commands do not work.	Verify that the VCS filesets are installed. Run the <code>hawparsetup</code> command to set up the WPAR configuration. Run the <code>hawparverify</code> command to verify the configuration. Run the <code>halogin</code> command from the WPAR. For more information, refer to the <i>Veritas Cluster Server Administrator's Guide</i> . Verify your VCS credentials. Make sure the password is not changed. Verify the VxSS certificate is not expired.

Symptom	Recommended Action
Resource does not come online in the WPAR.	Verify VCS and the agent filesets are installed correctly. Verify the application is installed in the WPAR. Verify the configuration definition of the agent. Make sure to define the Name and Type keys in the ContainerInfo attribute.

Configuring AIX WPARs for disaster recovery using VCS

AIX WPARs can be configured for disaster recovery by replicating the base directory using replication methods like Hitachi TrueCopy, EMC SRDF, Veritas Volume Replicator, and so on. The network configuration for the WPAR in the primary site may not be effective in the secondary site if the two sites are in different IP subnets. Hence, you need to make these additional configuration changes to the WPAR resource.

To configure the WPAR for Disaster Recovery, you need to configure VCS on both the sites in the Control Domains with GCO option.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information about global clusters.

To set up the WPAR for disaster recovery

- 1 On the primary site, create the WPAR and configure its network parameters.
- 2 On the primary site, start the WPAR and configure the DNS settings.
- 3 On the primary site, shut down the WPAR.
- 4 Use replication-specific commands to fail over the replication to the secondary site from the primary site.
- 5 Repeat step 1 on the secondary site.
- 6 Perform step 7, step 8, step 9, and step 10 on both the primary cluster and secondary clusters.
- 7 Create a VCS service group with a VCS WPAR resource for the WPAR.

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the WPAR resource.

Configure the DROpts association attribute on the WPAR resource with the following keys and site-specific values for each: DNSServers, DNSSearchPath, and DNSDomain.

- 8** Add the appropriate Mount resources and DiskGroup resources for the file system and disk group on which the WPAR's base directory resides.

Add a resource dependency from the WPAR resource to the Mount resource and another dependency from the Mount resource to the Diskgroup resource.

- 9** Add the appropriate VCS replication resource in the service group.

Examples of hardware replication agents are SRDF for EMC SRDF, HTC for Hitachi TrueCopy, MirrorView for EMC MirrorView, etc.

Refer the appropriate VCS replication agent guide for configuring the replication resource.

For VVR-based replication, add the appropriate RVGPrimary resource to the service group.

Refer to the following manuals for more information:

- For information about configuring VVR-related resources, see the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.
- For information about the VVR-related agents, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

- 10** Add a dependency from the DiskGroup resource to the replication resource.

Figure 8-5 Sample resource dependency diagram for hardware replication based WPARs

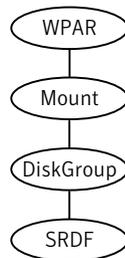
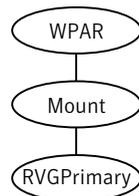


Figure 8-6 Sample resource dependency diagram for VVR replication-based WPARs



When the replication resource is online in a site, the replication agent makes sure of the following:

- The underlying replicated devices are in primary mode and hence the underlying storage and eventually the WPAR's base directory is always in read-write mode.
- The remote devices are in secondary mode.

When the WPAR resource goes online the resource modifies the appropriate files inside the WPAR to apply the disaster recovery-related parameters to the WPAR.

Data migration from Physical to Virtual Clients with NPIV

This chapter includes the following topics:

- [About migration from Physical to VIO environment](#)
- [Migrating from Physical to VIO environment](#)

About migration from Physical to VIO environment

Symantec has qualified migration of storage that is used by Storage Foundation from the physical environment to IBM VIO environment.

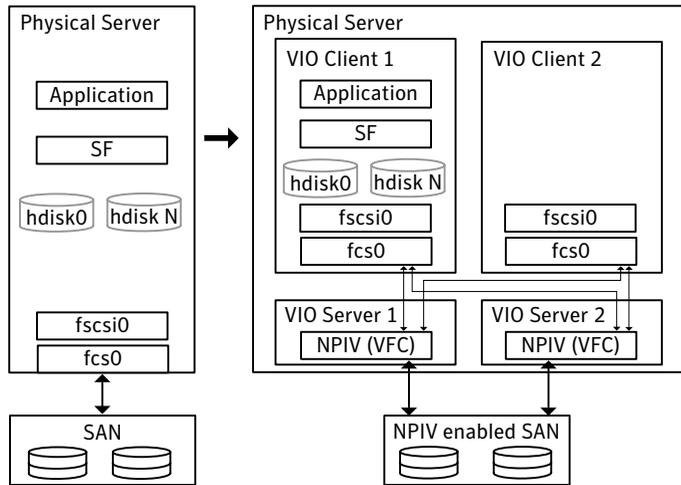
Storage Foundation provides the PDC (Portable Data Container) feature, which enables migrating storage from other platforms (Solaris, AIX, HP or Linux) to AIX VIO environment. You can also use PDC feature to migrate the storage consumed by a AIX physical server to a AIX VIO environment. NPIV helps you migrate the applications along with storage from a AIX physical environment to AIX VIO environment and vice-versa.

When storage is consumed by SF, Veritas Volume Manager (VxVM) initializes the storage LUNs as CDS (Cross-platform Data Sharing) type disks by default. A CDS disk group can be imported in a VIO client which has access to LUN's that are mapped through VFC Adapter on the client.

As part of the migration qualification, an application's storage is migrated from physical server to VIO environment (VIO client 1) which has NPIV capable FC adapter connected to it. This allows the application to access the storage in VIO

client 1. With NPIV capable FC adapter at VIOS, the devices presented to the VIO client would appear as regular AIX hdisk devices. [Figure 9-1](#) shows this migration.

Figure 9-1 SF migration from a physical environment to AIX VIO environment



Migration is an offline task.

Migrating from Physical to VIO environment

Migration of storage from a physical environment to the VIO environment is an offline task. The migration procedure involves stopping the application, unmounting the file systems and deporting the disk group on the physical server. Prior to being deported, you can take a space optimized snapshot, to facilitate fail-back.

Verify that the devices are visible on VIO client and the VFC adapter mapping between VIOS and VIO client is set up correctly. Refer to the IBM documentation for details. After all the required devices are accessible in VIO client 1, import the disk group in the client, mount the file system, and then start the application on the VIO client 1.

Refer to *IBM* documentation on how to configure the VFC adapter mappings between the VIO partition and the Client Partition.

Storage Foundation requirements for migration

Both the source and the target must have the same version of Storage Foundation. The version must be at least 5.0 MP3 RP1.

Boot device management

This chapter includes the following topics:

- [Using DMP to provide multi-pathing for the root volume group \(rootvg\)](#)
- [Boot device on NPIV presented devices](#)

Using DMP to provide multi-pathing for the root volume group (rootvg)

In many cases, the use of MPIO for the rootvg creates a situation with dual multi-pathing tools. To simplify system administration and system reliability, use DMP to provide multi-pathing for the rootvg.

DMP is supported for the rootvg on vSCSI, NPIV, and physical HBAs. DMP is also supported for alternate root disks and root disks with multiple volumes.

To use DMP on the rootvg, DMP requires a vendor-specific ODM predefined fileset. Symantec includes the predefined filesets for vSCSI devices in the Veritas product distribution. For other devices, obtain and install the ODM predefined fileset from the storage vendor. For example, for the IBM DS array, install the `devices.fcp.disk.ibm.rte` fileset.

http://www-1.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D400&q1=host+script&uid=ssg1S4000199&loc=en_US&cs=utf-8&lang=en

Rootability is achieved by using the `vxdmpadm` command, which uses the OS Native stack support feature internally.

Using DMP to provide multi-pathing for the root volume group (rootvg)**To get help about rootability**

- ◆ Run the following command:

```
# vxddmpadm help native

Manage DMP support for AIX boot volume group(rootvg)
Usage:
  vxddmpadm native { enable | disable } vname=rootvg
  vxddmpadm native list [ vname=rootvg ]
  vxddmpadm native { release | acquire } [ vname=rootvg ]
where,
  enable   Enable DMP support for AIX boot volume group(rootvg)
  disable  Disable DMP support for AIX boot volume group(rootvg)
  list     List boot paths on which DMP support is enabled
  release  Giveback pvid to OS device paths corresponding to rootvg
  acquire  Takeover pvid from OS device paths corresponding to rootvg
```

To enable rootability

- 1 Run the following command:

```
# vxddmpadm native enable vname=rootvg
```

- 2 Reboot the system to enable DMP support for LVM bootability.

To disable rootability

- 1 Run the following command:

```
# vxddmpadm native disable vname=rootvg
```

- 2 Reboot the system to disable DMP support for LVM bootability.

To monitor rootability

- ◆ Run the following command:

```
# vxddmpadm native list

PATH                DMPNODENAME
=====
hdisk64             ams_wms0_302
hdisk63             ams_wms0_302
```

For more information about using DMP with rootvg, see the *Veritas Dynamic Multi-Pathing Administrator's Guide*.

Boot device on NPIV presented devices

Dynamic Multi-Pathing (DMP) supports the NPIV presented devices for the rootvg, within the requirements outlined in the vendor support matrix.

Hardware and software requirements

- Any Power 6 or Power 7 based computer
- SAN Switch & FC Adapters should be NPIV capable.
- At least one 8 GB PCI Express Dual Port FC Adapter in VIOS.
- VIO Client Minimum OS-level:
 - AIX 6.1 TL5 or later
 - AIX 7.1 or later
- VIO Server Version 2.1 with Fix Pack 20.1 or later
- HMC 7.3.4

Boot Device Management

All the LUNs presented through NPIV for a client LPAR have the characteristics of a dedicated HBA. Therefore the procedure for using DMP on rootvg devices from NPIV presented devices is similar to using DMP on rootvg devices from physical HBA. Use of DMP on rootvg is supported through `vxdmproot native` command.

NPIV for Data volumes

The behavior of Data volumes presented through NPIV is similar to that of physical HBA. No special handling is required for these volumes. All SCSI device inquiry operations work and SCSI-3 persistent reservation functionality is also supported, enabling the use of SCSI-3 I/O Fencing if the underlying storage supports.

Glossary

Active Memory™ Sharing - Statement of Direction	Provides the ability to pool memory across micro-partitions which can be dynamically allocated based on partition's workload demands to improve memory utilization.
Dynamic Logical Partition (DLPAR)	A virtual server with the ability to add or remove full processors, network, or storage adapters while the server remains online.
Hardware Management Console (HMC)	Dedicated hardware/software to configure and administer a partition capable POWER server.
Integrated Virtualization Manager	Management console which runs in the VIO for partition management of entry level systems.
Live Partition Mobility	Provides the ability to migrate running AIX and Linux partitions across physical servers.
Lx86	Supports x86 Linux applications running on POWER.
Logical Partition (LPAR)	A virtual server running its own operating system instance with dedicated processors and I/O adapters.
Micro-partition	A virtual server with shared processor pools with support for up to 10 micro-partitions per processor core. Depending upon the Power server, you can run up to 254 independent micro-partitions within a single physical Power server. Processor resources can be assigned at a granularity of 1/100th of a core. Also known as shared processor partition.
Multiple Shared Processor Pools	Shared and capped processor resources for a group of micro-partitions.
N_Port ID Virtualization (NPIV)	Virtual HBAs which enable multiple LPARs/micro-partitions to access SAN devices thru shared HBAs providing direct Fibre Channel connections from client partitions to storage. Fibre Channel Host Bus Adapters (HBAs) are owned by VIO Server partition.
POWER Hypervisor	The POWER Hypervisor is responsible for dispatching the logical partition workload across the shared physical processors. The POWER Hypervisor also enforces partition security, and provides inter-partition communication that enables the Virtual I/O Server's virtual SCSI and virtual Ethernet function.

Shared Ethernet Adapter	The Shared Ethernet Adapter (SEA) enables network traffic outside the physical server by routing it through a software-based layer 2 switch running in the VIO Server.
Virtual I/O Server (VIO)	A dedicated LPAR which supports the I/O needs of client partitions (AIX and Linux) without the need to dedicate separate I/O slots for network connections and storage devices for each client partition.
Virtual Ethernet	In-memory network connections between partitions by POWER Hypervisor that reduce or eliminate the need for separate physical Ethernet Adapters in each LPAR.
Virtual SCSI	Virtual disks (vDisks) provided by the VIO server to reduce the need for dedicated physical disk resources for client partitions. HBAs are contained in the VIO server. vDisks can be full LUNs or logical volumes. Dynamic LPARs or micro-partitions can also use dedicated HBAs.
Application WPAR	An application Workload Partition (WPAR) is a light weight partition in which individual applications run. An application WPAR can only run application processes, not system daemons such as <code>inetd</code> or <code>cron</code> . An application WPAR is a temporary object which is removed when app is completed.
System WPAR	A system Workload Partition (WPAR) has a private copy of many of the AIX OS parameters. If desired, it can have its own dedicated, completely writable file systems. Most OS daemons can run, and each system WPAR has its own user privilege space. By default, a system WPAR has no access to physical devices.
WPAR Manager	The WPAR Manager allows an administrator to create, clone, and remove WPAR definitions, or start and stop WPARs. It enables Live Application Mobility which allows relocation of WPARs from one server to another without restarting the application. The WPAR Manager includes a policy engine to automate relocation of WPARs between systems based on system load and other metrics.