

Veritas Storage Foundation™ and High Availability Solutions Virtualization Guide

Solaris

5.1

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: 5.1

Document version: 5.1.0

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas, Veritas Storage Foundation 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 maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

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

Contacting Technical Support

Customers with a current maintenance 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 non-technical support Web page at the following URL:

customercare.symantec.com

Customer service

Customer Care information is available at the following URL:

www.symantec.com/customercare

Customer Service is available to assist with 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 maintenance 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

Documentation feedback

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:

sfha_docs@symantec.com

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance 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	supportolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4	
Chapter 1	Overview of Veritas Storage Foundation™ and High Availability Virtualization Solutions	13
	Overview	13
	Reference documentation	14
	Reference online	15
	About Veritas Storage Foundation and High Availability Virtualization Solutions	15
Chapter 2	Storage Foundation and High Availability Solutions support for Solaris Zones	17
	About Solaris Zones	17
	About VCS support for zones	18
	Overview of how VCS works with zones	18
	About the ContainerInfo service group attribute	19
	About the ContainerOpts resource type attribute	19
	Zone-aware resources	20
	About the Mount agent	21
	About the Zone agent	22
	About configuring failovers among physical and virtual servers	22
	Configuring VCS in zones	22
	Prerequisites for configuring VCS in zones	22
	Deciding on the zone root location	24
	Creating a zone with root on local disk	24
	Creating a zone root on shared storage	25
	About installing applications in a zone	27
	Configuring the service group for the application	27
	Modifying the service group configuration	30
	Verifying the zone configuration	30
	Performing maintenance tasks	31
	Troubleshooting zones	31
	Configuring for physical to virtual and virtual to physical failovers—a typical setup	32

Adding VxFS file systems to a non-global zone	32
Delegating VxFS file systems to a non-global zone	33
Mounting a VxFS file system in a non-global zone	34
Adding a direct mount to a zone's configuration	34
Veritas Storage Foundation Cluster File System mounts	35
Concurrent I/O access in non-global zones	36
Veritas extension for Oracle Disk Manager	36
Exporting VxVM volumes to a non-global zone	38
Volume physical nodes in Solaris 10	39
Removing a VxVM volume from a non-global zone	39
Software limitations of Storage Foundation support of non-global zones	40
Localization support for VxFS commands in non-global zone	40
Solaris patches required for Veritas Volume Manager support in non-global zone	40
Cannot remove a volume configured in a non-global zone	40
Administration commands are not supported in non-global zone	40
A cluster-mounted VxFS file system is not supported as the root of a non-global zone	41
QIO and CQIO are not supported	41
Package and patch installation in non-global zones	41
Package and patch removal with non-global zone configurations	41
Root volume cannot be added to non-global zones	41
Some Veritas Volume Manager operations can cause volume device names to go out of sync	42
 Chapter 3	
Storage Foundation and High Availability Solutions support for Branded Zones	43
About branded zones	43
System requirements	44
Veritas Storage Foundation support for branded zone	44
About migrating VCS clusters on Solaris 9 systems	45
Preparing to migrate a VCS cluster	46
Configuring VCS/SF in a branded zone environment	46
 Chapter 4	
Storage Foundation and High Availability Solutions support for Solaris Logical Domains	51
Terminology for Logical Domains	52
Solaris Logical Domains deployment models	53

Split Storage Foundation stack	54
Guest-based Storage Foundation stack	54
Benefits in a Logical Domain environment	54
Standardization of tools	55
Array migration	55
Moving storage between physical and virtual environments	55
Boot Image Management	55
New features	55
Storage Foundation new features	55
LDom new features	58
Split Storage Foundation stack model	59
How Storage Foundation and High Availability Solutions works in the Solaris Logical Domains	60
Veritas Storage Foundation features restrictions	61
Guest-based Storage Foundation stack model	63
How Storage Foundation and High Availability Solutions works in the guest Logical Domains	63
Supported configurations with SFCFS and multiple I/O Domains	65
Veritas Storage Foundation features restrictions	68
System requirements	69
Solaris operating system requirements	69
Solaris patch requirements	70
Sun hardware requirements	71
Component product release notes	71
Veritas Storage Foundation and High Availability	71
Veritas Storage Foundation Cluster File System	71
Product licensing	72
Installing Storage Foundation in a Logical Domain environment	72
Installing and configuring the Logical Domain software and domains	73
Installing Storage Foundation in the control domain or guest	73
Installing Veritas File System in the guest domain using pkgadd	74
Verifying the configuration	74
Migrating a Veritas Volume Manager disk group from a non-Logical Domain environment	74
Provisioning storage for a guest Logical Domain	76
Provisioning Veritas Volume Manager volumes as data disks for guest Logical Domains	76
Provisioning Veritas Volume Manager volumes as boot disks for guest Logical Domains	78

	Using Veritas Volume Manager snapshots for cloning Logical Domain boot disks	79
	Software limitations	84
	Memory Corruption in the guest Logical Domain during SCSI commands	84
	Exporting the raw volume device node fails	84
	Resizing a Veritas Volume Manager volume (exported as a slice or full disk) does not dynamically reflect the new size of the volume in the guest	85
	Known issues	85
	Guest-based known issues	86
	Split Storage Foundation stack known issues	87
Chapter 5	Veritas Cluster Server support for using multiple nodes in a Logical Domain environment	89
	Clustering using Cluster Volume Manager	89
	Installing Storage Foundation on multiple nodes in a Logical Domain	90
	Reconfiguring the clustering agents for Cluster Volume Manager	90
	Cluster Volume Manager in the control domain for providing high availability	92
Chapter 6	Veritas Cluster Server: Configuring Logical Domains for high availability	95
	About Veritas Cluster Server in a Logical Domain environment	95
	Veritas Cluster Server prerequisites	96
	Veritas Cluster Server requirements	97
	Veritas Cluster Server limitations	97
	Veritas Cluster Server known issues	98
	About Veritas Storage Foundation Cluster File System in a Logical Domain environment	99
	Veritas Storage Foundation Cluster File System limitations	99
	About Veritas Cluster Server configuration models in a Logical Domain environment	99
	Veritas Cluster Server setup to fail over a Logical Domain on a failure	100
	Veritas Cluster Server installed in the control domain to monitor the Logical Domains	100
	Veritas Cluster Server installed in the control domain to monitor the applications in the guest domain	101

Veritas Cluster Server setup to fail over an application on a failure	103
Configuring Veritas Cluster Server to fail over a Logical Domain on a failure	104
Configuration scenarios	105
Creating the Veritas Cluster Server service groups for Logical Domain	108
Configuring Veritas Cluster Server for application monitoring	108
Verifying a Logical Domain service group failover	111
Configuring Veritas Cluster Server to fail over an application on a failure	111
LDom migration in a VCS environment	112

Overview of Veritas Storage Foundation™ and High Availability Virtualization Solutions

This chapter includes the following topics:

- [Overview](#)
- [About Veritas Storage Foundation and High Availability Virtualization Solutions](#)

Overview

This document provides information about Veritas Storage Foundation and High Availability Virtualization Solutions. Review this entire document before you install Veritas Storage Foundation and High Availability products in zones, branded zones, and logical domains.

This book provides many high-level examples and information. As such, you should be a skilled user of Veritas products and knowledgeable concerning Sun's virtualization technologies.

This guide is divided into five main chapters excluding this overview. Each chapter presents information on using a particular Sun virtualization technology with Veritas products. These chapters follow:

- Storage Foundation and High Availability Solutions support for Solaris Zones
- Storage Foundation and High Availability Solutions support for Branded Zones

- Storage Foundation and High Availability Solutions support for Solaris Logical Domains
- Using multiple nodes in a Logical Domain environment
- Configuring Logical Domains for high availability

Reference documentation

The following documentation provides information on installing, configuring, and using Veritas Cluster Server:

- *Veritas Cluster Server Release Notes*
- *Veritas Cluster Server Installation Guide*
- *Veritas Cluster Server Bundled Agents Reference Guide*
- *Veritas Cluster Server Agent for DB2 Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Sybase Installation and Configuration Guide*

The following documentation provides information on installing, configuring, and using Veritas Storage Foundation products:

- *Veritas Storage Foundation Release Notes*
- *Veritas Storage Foundation Installation Guide*
- *Veritas Volume Manager Administrator's Guide*
- *Veritas File System Administrator's Guide*

The following documentation provides information on installing, configuring, and using Veritas Storage Foundation Cluster File System:

- *Veritas Storage Foundation Cluster File System Release Notes*
- *Veritas Storage Foundation Cluster File System Installation Guide*
- *Veritas Storage Foundation Cluster File System Administrator's Guide*

Note: Storage Foundation Cluster File System does not support branded zones.

For Solaris Logical Domains, Branded Zone, and Zone installation and configuration information, refer to the Sun Microsystems site: www.sun.com.

Sun Microsystems provides regular updates and patches for the Solaris Logical Domains, Branded Zones, and Zone features. Contact Sun Microsystems for details.

Reference online

For the latest information about this guide, see:

<http://seer.entsupport.symantec.com/docs/vascont/278.html>

About Veritas Storage Foundation and High Availability Virtualization Solutions

Veritas Storage Foundation and High Availability Virtualization Solutions includes support for non-global and Branded zones, and Solaris Logical Domains.

Solaris Zones, also known as non-global zones is an operating system-level virtualization technology, which provides a means of virtualizing operating system services to create an isolated environment for running applications. Non-global zones function as completely isolated virtual servers with a single operating system instance.

Branded zones are an extension of the Solaris Zone infrastructure. A Branded zone is a non-native zone that allows individual zones to emulate an operating system environment other than the native environment of the global operating system.

Solaris Logical Domains (LDoms) is a virtualization technology on the Solaris SPARC platform that enables the creation of independent virtual machine environments on the same physical system. LDoms are a virtualized computing environment abstracted from all physical devices, which allow you to consolidate and centrally manage your workloads on one system. The logical domains can be specified roles such as a Control domain, Service domain, I/O domain , and Guest domain. Each domain is a full virtual machine where the operating systems can be started, stopped, and rebooted independently.

Storage Foundation and High Availability Solutions support for Solaris Zones

This chapter includes the following topics:

- [About Solaris Zones](#)
- [About VCS support for zones](#)
- [Configuring VCS in zones](#)
- [Adding VxFS file systems to a non-global zone](#)
- [Delegating VxFS file systems to a non-global zone](#)
- [Veritas Storage Foundation Cluster File System mounts](#)
- [Concurrent I/O access in non-global zones](#)
- [Veritas extension for Oracle Disk Manager](#)
- [Exporting VxVM volumes to a non-global zone](#)
- [Software limitations of Storage Foundation support of non-global zones](#)

About Solaris Zones

Solaris Zones is a software partitioning technology, which provides a means of virtualizing operating system services to create an isolated environment for running applications. This isolation prevents processes that are running in one zone from monitoring or affecting processes running in other zones.

See the *System Administration Guide: Solaris Containers--Resource Management and Solaris Zones* Solaris operating environment document.

Sun Microsystems, Inc. provides regular updates and patches for the Solaris Zones feature. Contact Sun Microsystems, Inc. for more information.

About VCS support for zones

VCS provides application management and high availability to applications running in zones.

Overview of how VCS works with zones

You can use VCS to perform the following:

- Start, stop, monitor, and fail over a non-global zone.
- Start, stop, monitor, and fail over an application that runs in a zone.

How VCS models containers

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

You can configure VCS to use Symantec Product Authentication Service to run in a secure environment. Communication from non-global zones to global zones is secure in this environment.

Installing and configuring zones in VCS environments

Install and configure the zone. Create the service group with the standard application resource types (application, storage, networking) and the Zone resource. VCS manages the zone as a resource. You then configure the service group's ContainerInfo attribute.

Configuring the ContainerInfo attribute

The service group attribute ContainerInfo specifies information about the zone. When you have configured and enabled the ContainerInfo attribute, you have enabled the zone-aware resources in that service group to work in the zone environment.

VCS defines the zone information at the level of the service group so that you do not have to define it for each resource. You need to specify a per-system value for the ContainerInfo attribute.

About the ContainerInfo service group attribute

The ContainerInfo attribute has the Name key, Type key, and Enabled key. The Name key defines the name of the container. The Type key lets you select the type of container that you plan to use. The Enabled key enables the Zone-aware resources within the service group. The ContainerInfo attribute specifies if you can use the service group with the container.

Assign the following values to the ContainerInfo attribute:

- Name
The name of the container.
- Type
The type of container. You can set this to Zone.
- Enabled
Specify the value as 0, if you want to disable the container. Specify the value as 1, if you want to enable the container. Specify the value as 2, to enable physical to virtual and virtual to physical failovers. When the value is 2, the Zone resource mimics a non-existent entity.

You can set a per-system value for this attribute.

About the ContainerOpts resource type attribute

The ContainerOpts resource attribute is pre-set for Zone-aware resource types. It determines the following:

- Whether the zone-aware resource can run in the zone.
- Whether the container information that is defined in the service group's ContainerInfo attribute is passed to the resource.

These values are only effective when you configure the ContainerInfo service group attribute.

attribute's keys follow:

The ContainerOpts resource type attribute's definitions for Zone-aware types contain the following values:

- RunInContainer
When the value of the RunInContainer key is 1, the agent function (entry point) for that resource runs inside of the local container.

When the value of the RunInContainer key is 0, the agent function (entry point) for that resource runs outside the local container (in the global environment). A limitation for the RunInContainer value is that only script agent functions (entry points) can run inside a container.

■ **PassCInfo**

When the value of the PassCInfo key is 1, the agent function receives the container information that is defined in the service group’s ContainerInfo attribute. An example use of this value is to pass the name of the container to the agent.

Zone-aware resources

[Table 2-1](#) lists the ContainerOpts attributes default values for resource types. Zone-aware resources have predefined values for the ContainerOpts attribute.

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

See [“About the Mount agent”](#) on page 21.

Table 2-1 ContainerOpts attribute default values for applications and resource types

Resource type	RunInContainer	PassCInfo
Apache	1	0
Application	1	0
ASMinst	1	0
ASMDG	1	0
Db2udb	1	0
IP	0	1
IPMultiNIC	0	1
IPMultiNICB	0	1
Process	1	0
Zone	0	1
Oracle	1	0

Table 2-1 ContainerOpts attribute default values for applications and resource types (*continued*)

Resource type	RunInContainer	PassCInfo
Netlsnr	1	0
Sybase	1	0
SybaseBk	1	0

About the Mount agent

You may need to modify the ContainerOpts values for the Mount resource in certain situations.

Bringing a Mount resource online in the zone

The Mount resource is brought online in the global zone by default (RunInContainer = 0). If you want to bring a mount resource online inside the non-global zone, perform the following:

- Make sure that the resource is in a service group that has the ContainerInfo attribute configured.
- Override the ContainerOpts attribute at the resource level.
- Set the value of the RunInContainer key to 1.

For information on overriding resource type static attributes, refer to the *Veritas Cluster Server Administrator's Guide*.

Selecting the proper attribute values for a Mount resource for the zone's root file system for NFS mounts

For NFS mounts, you must mount in the non-global zone.

For this, set RIC=1. When you set RIC=1, specify the value of the MountPoint attribute relative to the zone root.

For example:

BlockDevice = abc:/fs1

MountPoint = /mnt1

The file system is mounted on /zone root/mnt1

About the Zone agent

The Zone agent monitors zones, brings them online, and takes them offline. For more information about the agent, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

The agent creates a user account with group administrative privileges, if such an account doesn't exist. The user account enables communication between the global zone and the zone. In secure clusters, it also renews the authentication certificate before the certificate expires.

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.

See [“Configuring for physical to virtual and virtual to physical failovers—a typical setup”](#) on page 32.

Configuring VCS in zones

Configuring VCS in zones involves the following tasks:

- | | |
|--------|--|
| First | Review the prerequisites.
See “Prerequisites for configuring VCS in zones” on page 22. |
| Second | Decide on the location of the zone root, which is either on local storage.
See “Deciding on the zone root location” on page 24. |
| Third | Install the application in the zone.
See “About installing applications in a zone” on page 27. |
| Fourth | Create the application service group and configure its resources.
See “Configuring the service group for the application” on page 27. |

Prerequisites for configuring VCS in zones

Review the following prerequisites for configuring VCS in zones:

- VCS supports UFS and VxFS mounts for the zone root. VCS does not support CFS mounts for the zone root.
- Mounts and CFS mounts must meet one of the following two conditions:

- Use a loopback file system. All mounts that the application uses must be part of the zone configuration and must be configured in the service group. For example, you can create a zone, `z-ora`, and define the file system containing the application's data to have the mount point as `/oradata`. When you create the zone, you can define a path in the global zone. An example is `/export/home/oradata`, which the mount directory in the non-global zone maps to. The `MountPoint` attribute of the `Mount` resource for the application is set to `/export/home/oradata`. Confirm that `/export/home/oradata` maps to `/oradata`.
- Use a direct mount file system. All file system mount points that the application uses that run in a zone must be set relative to the zone's root. For example, if the Oracle application uses `/oradata`, and you create the zone with the `zonpath` as `/z_ora`, then the mount must be `/z_ora/root/oradata`. The `MountPoint` attribute of the `Mount` resource must be set to this path. The `Mount` resource depends on the `Zone` resource.

Using custom agents in zones

If you use custom agents, review the following information for their use in zones:

- If you use custom agents to monitor the applications that run in zones, make sure that the agents use script-based entry points. VCS does not support running C++ entry points inside a zone.
- If you want the custom agent to monitor an application in the zone, for the custom agent type, set the following values for the `ContainerOpts` attribute: `RunInContainer = 1` and the `PassCInfo = 0`.
- If you don't want the custom agent to monitor an application in the zone, for the custom agent type, set the following values for the `ContainerOpts` attribute: `RunInContainer = 0` and the `PassCInfo = 0`.
- Two main cases exist where you might want to use a `RunInContainer = 0` and `PassCInfo = 1`, descriptions of these follow.
 - The first is the `Zone` agent's use of these values. The `Zone` agent's entry points cannot run inside of the non-global zone but the agent itself manages the zone. `RunInContainer` requires a value of 0 because the agent must run in the global zone. `PassCInfo` has a value of 1 because the `Zone` agent requires the name of the container from the `ContainerInfo` service group attribute.
 - The second case is how the `IP` agent uses `RunInContainer` and `PassCInfo`. The `IP` agent's entry points must run outside of the non-global zone because the networking stack is not completely in the non-global zone. You cannot perform an `ifconfig` command and then plumb the IP from inside of a

non-global zone. When you run the `ifconfig` command in the global zone with the `zone` option--it plumbs the IP and makes it available to the zone that you specify. The need for the container's name comes from the use of this command, even though it cannot run in the container.

Deciding on the zone root location

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

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

- Zone root on local storage

In this configuration, you must configure and install a zone on each node in the cluster.

- Zone root on shared storage

In this configuration, configure and install a zone in shared storage from one system and duplicate the configuration on each node in the cluster.

Setting the zone root on shared storage means you need to install the non-global zone on shared storage from one system only. The zone root can fail over to the other systems. To do this, the system software, including the patches, must be identical on each system during the existence of the zone. Symantec recommends using Solaris 10 Update 3 or later if you decide to configure the zone root on shared storage.

Creating a zone with root on local disk

Create a zone root on the local disk on each node in the cluster. The file system for application data is on a shared device and is either the loopback type or the direct mount type. For a direct mount file system, run the mount command from the global zone with the mount point specified as the complete path that starts with the zone root. For a loopback file system, add it into the zone's configuration before you boot the zone.

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

- 1 Configure the zone with the `zonecfg` command.

```
zonecfg -z newzone
zonecfg:newzone> create
```

- 2 Set the `zonepath` parameter to specify a location for the zone root.

```
zonecfg:newzone> set zonepath=/export/home/newzone
```


3 If your application data resides on a loopback mount file system, create the loopback file system in the zone.

4 Exit the zonecfg configuration.

```
zonecfg> exit
```

5 Create the zone root directory.

```
mkdir zonepath
```

6 Set permissions for the zone root directory.

```
chmod 700 zonepath
```

7 Install the non-global zone.

```
zoneadm -z zonename install
```

8 Repeat step 1 to step 7 on each system in the service group's SystemList.

9 If the application data is on a loopback file system, mount the file system containing the application's data on shared storage.

10 Boot the zone.

```
zoneadm -z zonename boot
```

11 If the application data is on a direct mount file system, mount the file system from the global zone with the complete path that starts with the zone root.

Creating a zone root on shared storage

Create a zone root which points to the shared disk's location on each node in the cluster. The file system for application data is on a shared device and is either the loopback type or the direct mount type. For a direct mount file system, run the mount command from the global zone with the mount point specified as the complete path that starts with the zone root. For a loopback file system, add it into the zone's configuration before you boot the zone.

To create a zone root on shared disks on each node in the cluster

- 1 Create a file system on shared storage for the zone root. The file system that is to contain the zone root may be in the same disk group as the file system that contains the application data.

- 2 Configure the zone with the `zonecfg` command.

```
zonecfg -z newzone  
zonecfg:newzone> create
```

- 3 Set the `zonpath` parameter to specify a location for the zone root.

```
zonecfg:newzone> set zonpath=/export/home/newzone
```

- 4 If your application data resides on a loopback mount file system, create the loopback file system in the zone.

- 5 Exit the `zonecfg` configuration.

```
zonecfg> exit
```

- 6 Create the zone root directory.

```
mkdir zonpath
```

- 7 Set permissions for the zone root directory.

```
chmod 700 zonpath
```

- 8 Repeat step 2 to step 7 on each system in the service group's SystemList.

- 9 Mount the file system that contains the shared storage on one of the systems that share the storage to the directory specified in `zonpath`.

- 10 Run the following command to install the zone on the system where the zone path is mounted.

```
zoneadm -z zonename install
```

- 11 If the application data is on a loopback file system, mount the file system containing the application's data on shared storage.

12 Boot the zone.

```
zoneadm -z zonename boot
```

- 13 If the application data is on a direct mount file system, mount the file system from the global zone with the complete path that starts with the zone root.

About installing applications in a zone

Perform the following tasks to install the application in a zone:

- If you have created zones locally on each node in the cluster, install the application identically in all zones 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 zone and the currently existing zones. The operating system installs the agents in future zones when they are installed.
- You must define all the mount points that the application uses that are configured in the zone in the service group's configuration.

Configuring the service group for the application

You need to configure the application service group and the required resource dependencies. The following diagrams illustrates different examples of resource dependencies. In one case the zone root is set up on local storage. In the other, zone root is set up on shared storage.

Resource dependency diagrams: zone root on local disks

The following resource dependency diagrams show zone configurations on local disks configured for loopback and direct mounted file systems.

[Figure 2-1](#) depicts the dependency diagram when the zone root is set up on local storage with the loopback file system for the application. You can replace the Mount resource with the CFSSMount resource and the DiskGroup resource with the CVMVolDg resource in the following diagram. In this configuration, decide if you want the service group to be a parallel service group. If so, you may need to localize certain attributes for resources in the service group. For example, you have to change the IP resource's Address attribute for each node.

Figure 2-1 Zone root on local disks with loopback file system





Figure 2-2 depicts the dependency diagram when the zone root is set up on local storage with a direct mount file system for the application. You can replace the Mount resource with the CFSMount resource and the DiskGroup resource with the CVMVolDg resource in the following diagram. In this configuration, decide if you want the service group to be a parallel service group. If so, you may need to localize certain attributes for resources in the service group. For example, you have to change the IP resource's Address attribute for each node.

Figure 2-2 Zone root on local disks with direct mount file system



Resource dependency diagrams: zone root on shared disks

The following resource dependency diagrams show zone configurations on shared disks configured for loopback and direct mounted file systems.

Figure 2-3 depicts the dependency diagram when a zone root is set up on shared storage with the loopback file system. You can replace the Mount resource with the CFSMount resource and the DiskGroup resource with the CVMVolDg resource in the following diagram. In this configuration, decide if you want the service group to be a parallel service group. If so, you may need to localize certain

attributes for resources in the service group. For example, you have to change the IP resource's Address attribute for each node.

Figure 2-3 Zone root on shared storage with loopback file system





Figure 2-4 depicts the dependency diagram when a zone root is set up on shared storage with the direct mount file system for the application. You can replace the Mount resource with the CFSMount resource and the DiskGroup resource with the CVMVolDg resource in the following diagram. In this configuration, decide if you want the service group to be a parallel service group. If so, you may need to localize certain attributes for resources in the service group. For example, you have to change the IP resource's Address attribute for each node.

Figure 2-4 Zone root on shared storage a direct mounted file system



Use the following principles when you create the service group:

- Set the MountPoint attribute of the Mount resource to the mount path.
- If the application requires an IP address, configure the IP resource in the service group.

- If the zone root file system is on shared storage, you can configure separate mounts for the zone and the application (as shown in the illustration), but you can configure the same disk group for both.

Modifying the service group configuration

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

To create the configuration to manage a zone

- 1 Run the `hazonesetup` script to set up the zone configuration.

```
# hazonesetup servicegroup_name zonerres_name zone_name \  
password autostart systems
```

servicegroup_name Name of the application service group.

zonerres_name Name of the resource configured to monitor the zone.

zone_name Name of the zone.

password Password to be assigned to VCS or Security (Symantec Product Authentication Service) user created by the command.

autostart If you adding a value of 1 for autostart, it populates the AutoStartList for the service group. Use a value of 1 or 0 for Autostart.

systems List of systems on which the service group will be configured. Use this option only when creating the service group.

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

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

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

Verifying the zone configuration

Run the `hazoneverify` command to verify the zone configuration.

The command verifies the following requirements:

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

To verify the zone configuration

- 1 If you use custom agents make sure the resource type is added to the APP_TYPES or SYS_TYPES environment variable.
See “[Using custom agents in zones](#)” on page 23.
- 2 Run the `hazoneverify` command to verify the zone configuration.

```
# hazoneverify servicegroup_name
```

Performing maintenance tasks

Perform the following maintenance tasks as necessary:

- Make sure that the zone configuration files are consistent on all the nodes at all times. The file is located at `/etc/zones/zone_name.xml`.
- 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 zones

Use following information to troubleshoot VCS and zones:

- VCS HA commands do not work.

Recommended actions:

- Verify the VCS packages are installed.
- Run the `halogin` command from the zone.
For more information on the `halogin` command, refer to the *Veritas Cluster Server User's Guide*.
- Verify your VCS credentials. Make sure the password is not changed.
- Verify the VxSS certificate is not expired.
- Resource does not come online in the zone.


Recommended actions:

- Verify VCS and the agent packages are installed correctly.
- Verify the application is installed in the zone.
- Verify the configuration definition of the agent.

Configuring for physical to virtual and virtual to physical failovers—a typical setup

In this configuration, you have two physical nodes. One node runs Solaris 9 (sysA) and another node that runs Solaris 10 or Solaris 10 without zones configured (sysB).

Figure 2-5 An application service group that can fail over into a zone and back



In the `main.cf` configuration file, define the container name, type of container, and whether it is enabled or not.

```
ContainerInfo@sysA = {Name = Z1, Type = Zone, Enabled = 2}  
ContainerInfo@sysB = {Name = Z1, Type = Zone, Enabled = 1}
```

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

Adding VxFS file systems to a non-global zone

VxFS file systems that were previously created in the global zone can be made available in the non-global zone using a loopback file system mount. This functionality is especially useful when the sole purpose of making the file system

available in the non-global zone is to share access of this file system with one or more non-global zones. For example, if a configuration file is available in a particular file system and this configuration file is required by the non-global zone, then the file system can be shared with the non-global zone using a loopback file system mount.

The following commands share access of file system `/mnt1` as a loopback file system mount with the non-global zone `myzone`:

```
global# zonecfg -z myzone
zonecfg:myzone> add fs
zonecfg:myzone:fs> set dir=/mnt1
zonecfg:myzone:fs> set special=/oradata
zonecfg:myzone:fs> set type=lofs
zonecfg:myzone:fs> end
zonecfg:myzone> commit
zonecfg:myzone> exit
global# zoneadm -z myzone reboot
```

The value of `dir` is a directory in the non-global zone. The value of `special` is a directory in the global zone to be mounted in the non-global zone.

Caution: Sharing file systems with non-global zones through a loopback file system mount makes the file system available for simultaneous access from all the non-global zones. This method should be used only when you want shared read-only access to the file system.

The loopback file system mount mode of sharing file systems in the non-global zones is supported in Veritas File System 4.1 and later.

Delegating VxFS file systems to a non-global zone

Exclusive access of a VxFS file system can be delegated to a non-global zone by direct mounting the file system in the non-global zone. Using direct mounts limits the visibility of and access to the file system to only the non-global zone that has direct mounted this file system.

Delegating file systems to non-global zone using direct mounts is supported in Veritas File System 4.1 Maintenance Pack 1 and later.

See [“Mounting a VxFS file system in a non-global zone”](#) on page 34.

See [“Adding a direct mount to a zone's configuration”](#) on page 34.

Mounting a VxFS file system in a non-global zone

To direct mount a VxFS file system in a non-global zone, the directory to mount must be in the non-global zone and the mount must take place from the global zone. The following procedure mounts the directory `dirmnt` in the non-global zone `myzone` with a mount path of `/zonedir/myzone/root/dirmnt`.

Note: VxFS entries in the global zone `/etc/vfstab` file for non-global zone direct mounts are not supported, as the non-global zone may not yet be booted at the time of `/etc/vfstab` execution.

Once a file system has been delegated to a non-global zone through a direct mount, the mount point will be visible in the global zone through the `mount` command, but not through the `df` command.

To direct mount a VxFS file system in a non-global zone

- 1 Log in to the zone and make the mount point:

```
global# zlogin myzone
myzone# mkdir dirmnt
myzone# exit
```

- 2 Mount the VxFS file system:

- Non-cluster file system:

```
global# mount -F vxfs /dev/vx/dsk/dg/voll /zonedir/zone1/root/dirmnt
```

- Cluster file system:

```
global# mount -F vxfs -o cluster /dev/vx/dsk/dg/voll /zonedir/zone1/root
```

- 3 Log in to the non-global zone and ensure that the file system is mounted:

```
global# zlogin myzone
myzone# df | grep dirmnt
/dirmnt (/dirmnt):142911566 blocks 17863944 files
```

Adding a direct mount to a zone's configuration

A non-global zone can also be configured to have a VxFS file system direct mount automatically when the zone boots using `zonecfg`. The `fsck` command is run,

before the file system is mounted. If the `fsck` command fails, the zone fails to boot.

To add a direct mount to a zone's configuration

1 Check the status and halt the zone:

```
global# zoneadm list -cv
  ID NAME          STATUS    PATH                      BRAND  IP
   0 global          running   /                        native shared
   2 myzone          running   /zones/myzone           native shared
global# zoneadm -z myzone halt
```

2 Add a new FS entry to the zone's configuration and set type to `vxfs`. For a non-cluster file system, for example:

```
global# zonecfg -z myzone
zonecfg:myzone> add fs
zonecfg:myzone:fs> set dir=/dirmnt
zonecfg:myzone:fs> set type=vxfs
zonecfg:myzone:fs> set special=/dev/vx/dsk/dg/voll
zonecfg:myzone:fs> set raw=/dev/vx/rdisk/dg/voll
zonecfg:myzone:fs> end
zonecfg:myzone> commit
```

3 Boot the zone:

```
global# zoneadm -z myzone boot
```

4 Log in to the non-global zone and ensure that the file system is mounted:

```
global# zlogin myzone
myzone# df | grep dirmnt
/dirmnt (/dirmnt):142911566 blocks 17863944 files
```

Veritas Storage Foundation Cluster File System mounts

Veritas Storage Foundation Cluster File System (SFCFS) provides support for the same file system to be made available from multiple nodes that have been grouped together as a cluster. VxFS supports the sharing or delegation of cluster-mounted file systems in the non-global zone.

See [“Adding VxFS file systems to a non-global zone”](#) on page 32.

See [“Delegating VxFS file systems to a non-global zone”](#) on page 33.

The requirements to support SFCFS in non-global zones parallels that of SFCFS support in global zones. Some key points are as follows:

- Both lofs and direct mount are supported; Symantec recommends direct mount
- The device must be visible and shared on all nodes
- The mount point path must be the same on all nodes
- The zone name and configuration should be the same on all nodes

Support for SFCFS in a non-global zone is available in Veritas File System 5.0 Maintenance Pack 1 and later.

Concurrent I/O access in non-global zones

Concurrent I/O allows multiple processes to read from or write to the same file without blocking other `read(2)` or `write(2)` calls. POSIX semantics requires `read` and `write` calls to be serialized on a file with other `read` and `write` calls.

Concurrent I/O is generally used by applications that require high performance for accessing data and do not perform overlapping writes to the same file.

Veritas Storage Foundation supports concurrent I/O for applications running in the non-global zones as well. This implies that a process in a non-global zone can access the file concurrently with other processes in the global or non-global zone. The application or running threads are responsible for coordinating the write activities to the same file when using Concurrent I/O.

An application must perform the following activities to enable the concurrent I/O advisory on a file:

```
fd=open(filename, oflag)
ioctl(fd, VX_SETCACHE, VX_CONCURRENT)
write(fd, buff, numofbytes)
```

Veritas extension for Oracle Disk Manager

The Veritas extension for Oracle Disk Manager (ODM) is specifically designed for Oracle9i or later to enhance file management and disk I/O throughput. The features of Oracle Disk Manager are best suited for databases that reside in a file system contained in Veritas File System. Oracle Disk Manager allows Oracle9i or later users to improve database throughput for I/O intensive workloads with special I/O optimization.

The Veritas extension for Oracle Disk Manager is supported in non-global zones. To run Oracle 10g on a non-global zone and use Oracle Disk Manager, the Oracle software version must be 10.1.0.3 or higher.

Care must be taken when installing and removing packages and patches when working with the VRTSodm package, for more information refer to the following:

- See “[Package and patch installation in non-global zones](#)” on page 41.
- See “[Package and patch removal with non-global zone configurations](#)” on page 41.

The following procedure enables Oracle Disk Manager file access from non-global zones with Veritas File System.

To enable Oracle Disk Manager file access from non-global zones with Veritas File System

- 1 Make global zone licenses visible to the non-global zone by exporting the `/etc/vx/licenses/lic` directory to the non-global zone as a lofs:

```
global# zonecfg -z myzone
zonecfg:myzone> add fs
zonecfg:myzone:fs> set dir=/etc/vx/licenses/lic
zonecfg:myzone:fs> set special=/etc/vx/licenses/lic
zonecfg:myzone:fs> set type=lofs
zonecfg:myzone:fs> end
zonecfg:myzone> commit
```

- 2 Create the `/dev/odm` directory in the non-global zonepath from the global zone:

```
global# mkdir -p myzonepath/dev/odm
```

- 3 Log in to the non-global zone and mount `/dev/odm` either manually or use the startup script. Use one of the following:

- `global# zlogin myzone`
`myzone# mount -F odm /dev/odm /dev/odm`

Or:

- `global# zlogin myzone`
`myzone# /lib/svc/method/odm start`

Exporting VxVM volumes to a non-global zone

A volume device node can be exported for use in non-global zone using the `zonecfg` command. The following procedure makes a volume `vol1` available in the non-global zone `myzone`.

Caution: Exporting raw volumes to non-global zones has implicit security risks. It is possible for the zone administrator to create malformed file systems that could later panic the system when a mount is attempted. Directly writing to raw volumes, exported to non-global zones, and using utilities such as `dd` can lead to data corruption in certain scenarios.

To export VxVM volumes to a non-global zone

- 1 Create a volume `vol1` in the global zone:

```
global# ls -l /dev/vx/rdisk/rootdg/vol1
crw-----  1 root    root      301, 102000 Jun  3
12:54 /dev/vx/rdisk/rootdg/vol1crw-----  1 root    sys      301, 10200
0 Jun  3 12:54 /devices/pseudo/vxio@0:rootdg,vol1,102000,raw
```

- 2 Add the volume device `vol1` to the non-global zone `myzone`:

```
global# zonecfg -z myzone
zonecfg:myzone> add device
zonecfg:myzone:device> set match=/dev/vx/rdisk/rootdg/vol1
zonecfg:myzone:device> end
zonecfg:myzone> commit
```

- 3 Ensure that the devices will be seen in the non-global zone:

```
global# zoneadm -z myzone halt
global# zoneadm -z myzone boot
```

- 4 Verify that `/myzone/dev/vx` contains the raw volume node and that the non-global zone can perform I/O to the raw volume node.

The exported device can now be used for performing I/O or for creating file systems.

Note: A VxFS file system can only be constructed and mounted from the global zone.

Volume physical nodes in Solaris 10

On the Solaris 10 operating environment, there are two physical nodes corresponding to each volume node entry, `/devices` and `/dev`, respectively, with the same major and minor number. The physical nodes appear as follows:

```
/devices raw volume node : /devices/pseudo/vxio@0:
    dname,volname,minor_number,raw
/devices block volume node : /devices/pseudo/vxio@0:
    dname,volname,minor_number,blk
/dev raw volume node : /dev/vx/rdisk/dgname/volumename
/dev block volume node : /dev/vx/dsk/dgname/volumename
```

The following example provides sample values in `/devices`:

```
vm240v1:/-> ls -l /devices/pseudo/vxio*vol1*
brw----- 1 root    sys      302, 66000 Mar 25
17:21 /devices/pseudo/vxio@0:mydg,vol1,66000,blk
crw----- 1 root    sys      302, 66000 Mar 25
17:21 /devices/pseudo/vxio@0:mydg,vol1,66000,raw
```

The following example provides sample values in `/dev`:

```
vm240v1:/-> ls -l /dev/vx/*dsk/mydg/vol1
brw----- 1 root    root     302, 66000 Mar 25 17:21 /dev/vx/dsk/mydg/vol1
crw----- 1 root    root     302, 66000 Mar 25 17:21 /dev/vx/rdisk/mydg/vol1
```

Removing a VxVM volume from a non-global zone

The following procedure removes a VxVM volume from a non-global zone.

To remove a VxVM volume from a non-global zone

- ◆ Remove the volume device `vol13` from the non-global zone `myzone`:

```
global# zonecfg -z myzone
zonecfg:myzone> remove device match=/dev/vx/rdisk/pddg/vol13
zonecfg:myzone> end
zonecfg:myzone> commit
```

Software limitations of Storage Foundation support of non-global zones

This section describes the software limitations of Storage Foundation support of non-global zones in this release.

Localization support for VxFS commands in non-global zone

To ensure localization support for VxFS commands in non-global zone, you need to inherit `/etc/fs/vxfs` directory when setting up non-global zone configuration. For example:

```
zonecfg:myzone1>add inherit-pkg-dir
zonecfg:myzone1:inherit-pkg-dir>set dir=/etc/fs/vxfs
zonecfg:myzone1:inherit-pkg-dir>end
```

Solaris patches required for Veritas Volume Manager support in non-global zone

The *Veritas Storage Foundation Release Notes* contains a list of patches required for the proper operation of VxVM in non-global zones.

See the *Veritas Storage Foundation Release Notes*.

Cannot remove a volume configured in a non-global zone

After a volume node is configured in a Solaris non-global zone, the node cannot be removed in certain Solaris releases. Solaris patch 122660-10 resolves this issue.

Administration commands are not supported in non-global zone

All administrative tasks, such as resizing a volume, adding a volume to a volume set, and file system reorganization, are supported only in the global zone. Consequently, administrative commands, such as `fsadm`, `fsvoladm`, and `vxassist`, and administrative `ioctl`s are not supported in the non-global zone by both VxFS and VxVM.

A cluster-mounted VxFS file system is not supported as the root of a non-global zone

The root of a non-global zone cannot currently be on a cluster-mounted VxFS file system.

QIO and CQIO are not supported

Quick I/O and Cached Quick I/O are not supported by VxFS in non-global zones.

Package and patch installation in non-global zones

To ensure that the package and patch updates applied to Veritas products in the global zone are also propagated to the non-global zone, ensure that the non-global zones are in a bootable state. For example, to update the software installed in a system, the file system housing the root directory of the non-global zone needs to be mounted and the disk group housing the non-global zone needs to be online at the time that you run `patchadd`, `pkgadd`, or the CPI installation scripts.

If `VRTSodm` is part of the package installation, all non-global zones must be booted and in a running state at the time of package installation. If the non-global zones are not booted, you may need to reinstall the `VRTSodm` package manually after booting the non-global zones.

Package and patch removal with non-global zone configurations

If non-global zones are part of the system configuration and the `VRTSodm` package is installed, ensure that `/dev/odm` is unmounted in each non-global zone prior to `VRTSodm` package removal or product uninstallation. This ensures there are no non-global zone `odm` module references that might prevent the global zone from unloading the `odm` module.

You can unmount the non-global zone with the following commands:

```
global# zlogin myzone  
myzone# umount /dev/odm
```

Root volume cannot be added to non-global zones

The root volume cannot be added to non-global zones.

Some Veritas Volume Manager operations can cause volume device names to go out of sync

If a volume is exported to a non-global zone, some Veritas Volume Manager operations can cause the global and non-global volume device names to go out of sync, which can create data corruption. This is because the Solaris operating environment zone support is not aware of the `devfsadm(1M)` command, and thus the zone configuration is not updated with any changes to the `/dev` or `/devices` namespaces.

The following operations can cause device names to go out of sync:

- Removing a volume
- Importing a disk group
- Deporting a disk group
- Renaming a disk group or volume
- Reminoring a disk group
- Restarting `vxconfigd` or resetting the kernel

To prevent device names from to going out of sync, if a volume is exported to a non-global zone and an operation that can cause device names to go out of sync occurs on that volume, remove the volume from the zone configuration using the `zonecfg` command and reboot the zone using the `zoneadm` command.

See the `zonecfg(1M)` and `zoneadm(1M)` manual pages.

Note: This issue applies to any Solaris device for which the `/dev` or `/devices` device node is changed and has been configured in the non-global zone before the change, including Solaris Volume Manager volumes.

Storage Foundation and High Availability Solutions support for Branded Zones

This chapter includes the following topics:

- [About branded zones](#)
- [System requirements](#)
- [Veritas Storage Foundation support for branded zone](#)
- [About migrating VCS clusters on Solaris 9 systems](#)
- [Preparing to migrate a VCS cluster](#)
- [Configuring VCS/SF in a branded zone environment](#)

About branded zones

Branded zones (BrandZ) is a framework that extends the Solaris Zones infrastructure to create branded zones. Branded zone is a non-native zone that allows you to emulate an operating system environment other than the native operating system. Each operating system is plugged into the BrandZ framework with a brand associated to the operating system.

See the Sun documentation for more information on branded zones.

System requirements

Veritas Cluster Server (VCS) and Veritas Storage Foundation (SF) requirements in a branded zone environment are as follows:

Solaris requirements

For Solaris 10 10/08 or later

- For Solaris 9 branded zones:
Install Solaris 9 Containers 1.0.1

For Solaris 10 05/08 or prior

- For Solaris 9 branded zones:
Install Solaris 9 Containers 1.0

You can obtain the Containers software bundles from Sun Download Center at: <http://www.sun.com/software/solaris/containers>. For detailed information about the above requirements, read Sun Microsystems' README files from the software bundle.

Symantec recommends that you apply the latest Solaris operating system patches available from Sun. See the following site:

<http://sunsolve.sun.com>

Note: Solaris 8 is not supported in this release.

SF requirements ■ SF 5.1

VCS requirements ■ VCS 5.1

Database support The following Oracle versions in branded zone are supported:

- 9iR2
- 10gR2
- 11gR1

Veritas Storage Foundation support for branded zone

SF provides support for the following in a branded zone environment:

- VxVM volume devices
- VxFS file systems using loopback file system mount or direct mount

You can export a VxVM volume or a VxFS file system to a branded zone and access the volume or the file system in the branded zone. The procedure to export the volumes or the file systems are the same as that for the non-global zones.

Refer to the *Veritas Volume Manager Administrator's Guide* or the *Veritas File System Administrator's Guide* for more details.

About migrating VCS clusters on Solaris 9 systems

You can migrate VCS clusters that run on Solaris 9 systems to solaris9 branded zones on Solaris 10 systems. For example, with BrandZ you can emulate a Solaris 9 operating system as Solaris 9 container in the Solaris 10 branded zone. This Solaris 9 non-global zone acts as a complete runtime environment for Solaris 9 applications on Solaris 10 SPARC systems. You can directly migrate an existing Solaris 9 system into a Solaris 9 container.

[Figure 3-1](#) illustrates the workflow to migrate a VCS cluster on Solaris 9 systems to branded zones on Solaris 10 systems.

Figure 3-1 Workflow to migrate VCS cluster to branded zones

On Solaris 9 systems:

Uninstall VCS/SF

Create a flash archive of the Solaris system image

On Solaris 10 systems:

Install VCS/SF in the global zone

Install the VCS Zone agent patch in the global zone

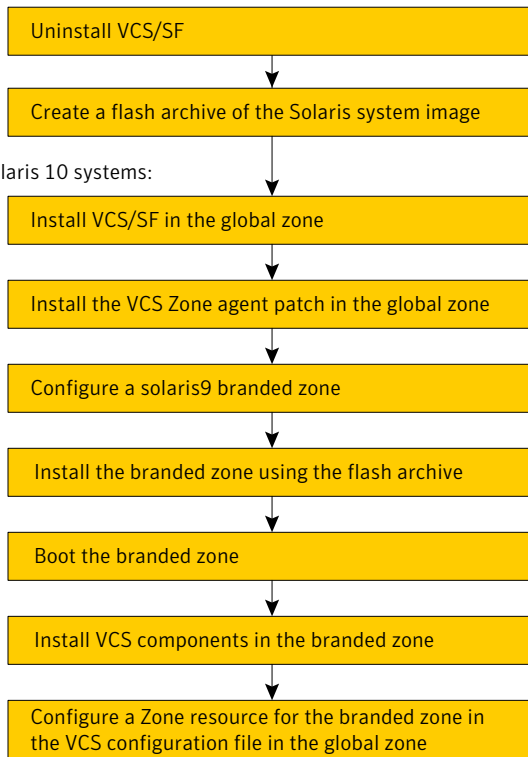
Configure a solaris9 branded zone

Install the branded zone using the flash archive

Boot the branded zone

Install VCS components in the branded zone

Configure a Zone resource for the branded zone in the VCS configuration file in the global zone



Preparing to migrate a VCS cluster

You must perform the following steps on the Solaris 9 systems from where you want to migrate VCS.

To prepare to migrate a VCS cluster

- 1 Uninstall VCS on the systems.

See the *Veritas Cluster Server Installation Guide*.

If SF is installed, uninstall SF on the systems.

See the *Veritas Storage Foundation Installation Guide*.

- 2 Create a flash archive. For example:

```
# flarcreate -S -n sol9image /tmp/sol9image.flar
```

Configuring VCS/SF in a branded zone environment

You must perform the following steps on the Solaris 10 systems.

To configure VCS/SF in a branded zone environment

- 1 Install VCS, SF, or SFHA as required in the global zone.

See the *Veritas Cluster Server Installation Guide*.

See the *Veritas Storage Foundation Installation Guide*.

- 2 Install the VCS Zone agent patch (139146-01) in the global zone if you installed VCS or SFHA.

- 3 Install VCS agent packages for Oracle in the global zone if necessary.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for installation instructions.

- 4 For ODM support, install the VRTSodm (121710-05) 5.1 package in the global zone.

- 5 Configure a solaris9 branded zone. For example, this step configures a solaris9 zone.

- Run the following command in the global zone as the global administrator:

```
# zonecfg -z sol9-zone
sol9-zone: No such zone configured
Use 'create' to begin configuring a new zone.
```

- Create the solaris9 branded zone using the SUNWsolaris9 template.

```
zonecfg:sol9-zone> create -t SUNWsolaris9
```

- Set the zone path. For example:

```
zonecfg:sol9-zone> set zonepath= /zones/sol9-zone
```

Note that zone root for the branded zone can either be on the local storage or the shared storage (raw volumes with VxFS).

- Add a virtual network interface.

```
zonecfg:sol9-zone> add net
zonecfg:sol9-zone:net> set physical=hme0
zonecfg:sol9-zone:net> set address= 192.168.1.20
zonecfg:sol9-zone:net> end
```

- Verify the zone configuration for the zone and exit the zonecfg command prompt.

```
zonecfg:sol9-zone> verify
zonecfg:sol9-zone> exit
```

The zone configuration is committed.

- 6 Verify the zone information for the solaris9 zone you configured.

```
# zonecfg -z sol9-zone info
```

Review the output to make sure the configuration is correct.

- 7 Install the solaris9 zone that you created using the flash archive you created in step 1.

```
# zoneadm -z sol9-zone install -p -a /tmp/sol9image.flar
```

After the zone installation is complete, run the following command to list the installed zones and to verify the status of the zones.

```
# zoneadm list -iv
```

- 8 Boot the solaris9 branded zone.

```
# /usr/lib/brand/solaris9/s9_p2v sol9-zone
```

```
# zoneadm -z sol9-zone boot
```

After the zone booting is complete, run the following command to verify the status of the zones.

```
# zoneadm list -v
```

9 Install VCS in the branded zone:

- Install only the following VCS 5.1 packages:

- VRTSperl
- VRTSat
- VRTSvc
- VRTSvcsg

- For Solaris 10, edit the VCS startup script (/etc/init.d/vcs) to add the following in the beginning of the file:

```
exit 0
```

10 If you configured Oracle to run in the branded zone, then install the VCS agent for Oracle packages (VRTSvcsea) and the patch in the branded zone.

See the *Veritas Cluster Server Agent for Oracle Installation and Configuration Guide* for installation instructions.

11 For ODM support, install the following additional packages and patches in the branded zone:

- Install the following 5.1 packages:

- VRTSvlic
- VRTSodm

12 If using ODM support, relink Oracle ODM library in Solaris 9 branded zones:

- Log into Oracle instance.
- Relink Oracle ODM library.

If you are running Oracle 9iR2:

```
$ rm $ORACLE_HOME/lib/libodm9.so
$ ln -s /opt/VRTSodm/lib/sparcv9/libodm.so \
$ORACLE_HOME/lib/libodm9.so
```

If you are running Oracle 10gR2:


```
$ rm $ORACLE_HOME/lib/libodm10.so
$ ln -s /opt/VRTSodm/lib/sparcv9/libodm.so \
$ORACLE_HOME/lib/libodm10.so
```

If you are running Oracle 11gR1:

```
$ rm $ORACLE_HOME/lib/libodm11.so
$ ln -s /opt/VRTSodm/lib/sparcv9/libodm.so \
$ORACLE_HOME/lib/libodm11.so
```

- Ensure that you have the correct license to run ODM. If you are migrating from a host which has a licence to run ODM, ensure you have the correct license in `/etc/vx/license/vx` directory.

Otherwise, make global zone licenses visible to the non-global zone by exporting the `/etc/vx/licenses/lic` directory to the non-global zone as a lofs:

```
global# zonecfg -z myzone
zonecfg:myzone> add fs
zonecfg:myzone:fs> set dir=/etc/vx/licenses/lic
zonecfg:myzone:fs> set special=/etc/vx/licenses/lic
zonecfg:myzone:fs> set type=lofs
zonecfg:myzone:fs> end
zonecfg:myzone:fs> commit
```

Check if the `/dev/odm` directory exists, if not create the `/dev/odm` directory in the non-global zone from the global zone:

```
global# mkdir -p /myzone/dev/odm
```

- If you are using ODM, first ensure that Storage Foundation works properly. If ODM is not started in the branded zone, start ODM. Log in to the branded zone and mount `/dev/odm` or run `/lib/svc/method/odm start`:

```
global# zlogin myzone
myzone# mount -F odm /dev/odm /dev/odm
```

Or:

```
myzone# /lib/svc/method/odm start
```

`/dev/odm` is not automatically mounted after a zone is booted. Use the `/lib/svc/method/odm start` command to mount `/dev/odm` after a zone is booted.

13 Configure the resources in the VCS configuration file in the global zone. For example:

```
group g1 (
    SystemList = { vcs_sol1 = 0, vcs_sol2 = 1 }
    ContainterInfo@vcs_sol1 {Name = zone1, Type = zone, Enabled = 1 }
    ContainterInfo@vcs_sol2 {Name = zone1, Type = zone, Enabled = 1 }
    AutoStartList = { vcs_sol1 }
    Administrators = { "z_z1@vcs_lzs@vcs_sol2.symantecexample.com" }
)

Process p1 (
    PathName = "/bin/ksh"
    Arguments = "/var/tmp/cont_yoyo"
)

Zone z1 (
)

p1 requires z1
```

See the *Veritas Cluster Server Bundled Agents Reference Guide* for VCS Zone agent details.

Storage Foundation and High Availability Solutions support for Solaris Logical Domains

This chapter includes the following topics:

- [Terminology for Logical Domains](#)
- [Solaris Logical Domains deployment models](#)
- [Benefits in a Logical Domain environment](#)
- [New features](#)
- [Split Storage Foundation stack model](#)
- [Guest-based Storage Foundation stack model](#)
- [System requirements](#)
- [Component product release notes](#)
- [Product licensing](#)
- [Installing Storage Foundation in a Logical Domain environment](#)
- [Migrating a Veritas Volume Manager disk group from a non-Logical Domain environment](#)
- [Provisioning storage for a guest Logical Domain](#)

- [Using Veritas Volume Manager snapshots for cloning Logical Domain boot disks](#)
- [Software limitations](#)
- [Known issues](#)

Terminology for Logical Domains

The following terminology is helpful in configuring the Veritas software in Logical Domains.

Table 4-1 lists the terminology for Logical Domains

Term	Definition
LDom	Logical Domain or Virtual Machine with its own operating system, resources, and identity within the same physical host.
Hypervisor	A firmware layer that provides a set of hardware-specific support functions to the operating systems running inside LDoms through a stable interface, known as the sun4v architecture. The hypervisor is interposed between the operating system and the hardware layer.
Logical Domains Manager	Software that communicates with the Hypervisor and logical domains to sequence changes, such as the removal of resources or creation of a logical domain. The Logical Domains Manager provides an administrative interface and keeps track of the mapping between the physical and virtual devices in a system.
Control domain	The primary domain which provides a configuration platform to the system for the setup and teardown of logical domains. Executes Logical Domains Manager software to govern logical domain creation and assignment of physical resources.

Table 4-1 lists the terminology for Logical Domains (*continued*)

Term	Definition
I/O domain	Controls direct, physical access to input/output devices, such as PCI Express cards, storage units, and network devices. The default I/O domain is the control domain.
Guest domain	Utilizes virtual devices offered by control and I/O domains and operates under the management of the control domain.
Virtual devices	Physical system hardware, including CPU, memory, and I/O devices that are abstracted by the Hypervisor and presented to logical domains within the platform.
Logical Domains Channel (LDC)	A logical domain channel is a point-to-point, full-duplex link created by the Hypervisor. LDCs provide a data path between virtual devices and guest domains and establish virtual networks between logical domains.
Virtual Disk Client	A Solaris kernel module in the guest logical domain which controls the virtual disks visible to that guest, providing standard device interfaces to applications.
Virtual Disk Server	A Solaris kernel module in the control domain which is responsible for exporting various backend devices as virtual disks to guest logical domains.

Solaris Logical Domains deployment models

Logical Domains (LDoms) is a virtualization technology on the Solaris SPARC platform that enables the creation of independent virtual machine environments on the same physical system. This allows you to consolidate and centrally manage your workloads on one system.

Veritas Storage Foundation supports Solaris Logical Domains in the following two deployments models:

- Split Storage Foundation stack
- Guest-based Storage Foundation stack

Veritas Storage Foundation Cluster File System (SFCFS) is supported only in the guest-based Storage Foundation stack.

See “[About Veritas Cluster Server configuration models in a Logical Domain environment](#)” on page 99.

Split Storage Foundation stack

The support for this model was introduced in 5.0 MP1 release and this model continues to be supported in this release.

See “[Split Storage Foundation stack model](#)” on page 59.

See “[Veritas Cluster Server setup to fail over a Logical Domain on a failure](#)” on page 100.

Guest-based Storage Foundation stack

The support for this model is being introduced in 5.0 MP3 release. This support includes Veritas Storage Foundation Cluster File System.

See “[New features](#)” on page 55.

To work inside the guest LDomS you must have 5.0 MP3 RP1 installed.

See “[Known issues](#)” on page 85.

For the latest information on updates, patches, documentation, and known issues regarding this 5.0 MP3 RP1 release, see the following TechNote on the Symantec Technical Support website:

For Solaris SPARC:

<http://entsupport.symantec.com/docs/281987>

For Solaris x64:

<http://entsupport.symantec.com/docs/286955>

Note: The SFCFS stack can be installed across multiple I/O domains within or across physical servers.

See “[Veritas Cluster Server limitations](#)” on page 97.

Benefits in a Logical Domain environment

There are several benefits to a Logical Domain environment.

Standardization of tools

Independent of how an operating system is hosted, consistent storage management tools save an administrator time and reduce the complexity of the environment.

Storage Foundation in the control domain provides the same command set, storage namespace, and environment as in a non-virtual environment.

Array migration

Data migration for Storage Foundation can be executed in a central location, migrating all storage from an array utilized by Storage Foundation managed hosts.

This powerful, centralized data migration functionality is available with Storage Foundation Manager 1.1 and later.

Moving storage between physical and virtual environments

Storage Foundation can make painful migrations of data from physical to virtual environments easier and safer to execute.

With Storage Foundation, there is no need to copy any data from source to destination, but rather the administrator reassigns the same storage or a copy of the storage for a test migration, to the virtual environment.

Boot Image Management

Using Storage Foundation in this environment the user can utilize features such as instant snapshots to contain boot images and manage them from a central location in the control domain.

New features

This section describes the new features in Solaris Logical Domains (LDoms) using the products in the Veritas Storage Foundation and High Availability Solutions.

Storage Foundation new features

The following new features apply for Storage Foundation.

The vxloadm utility enables access to a file system contained in a VxVM volume from the Control Domain

The `vxloadm` utility lets you access a file system contained inside a VxVM volume from outside the guest domain, that is from the Control Domain. This is done by

mapping all the partitions contained within that volume using the `vxlo` driver. The partitions can then be mounted if they contain valid file systems.

To use this `vxloadm` utility

- 1 Load the `vxlo` driver in memory:

```
# cd /kernel/drv/sparcv9
# add_drv -m '* 0640 root sys' vxlo
# modload vxlo
```

- 2 Check if the driver is loaded in memory:

```
# modinfo | grep vxlo
226 7b3ec000 3870 306 1 vxlo (Veritas Loopback Driver 0.1)
```

- 3 Run the `vxloadm` utility:

```
# /etc/vx/bin/vxloadm
```

- 4 You can now use the utility.

See [“Examples of using the `vxloadm` utility”](#) on page 56.

- 5 Symantec recommends once you are done using the `vxloadm` utility to unload the `vxlo` driver:

```
# rem_drv vxlo
# modinfo | grep vxlo
226 7b3ec000 3870 306 1 vxlo (Veritas Loopback Driver 0.1)
# modunload -i 226
```

where `226` is the module ID from the `modinfo | grep vxlo` command.

Examples of using the `vxloadm` utility

Use the `vxloadm addall` command to create device(s) mapping the various partition(s) contained in a VxVM volume. For example:

```
# /etc/vx/bin/vxloadm addall vol1 /dev/vx/dsk/testdg/vol1
```

This creates a device node entry for every slice or partition contained within the volume in the `/dev/vxlo/dsk/` and `/dev/vxlo/rdisk/` directories.

```
# ls -l /dev/vxlo/dsk/
lrwxrwxrwx 1 root root 46 Sep 25 14:04 vol1s0
-> ../../../../devices/pseudo/vxlo@0:vol1s0,1,blk
```



```
lrwxrwxrwx 1 root root 46 Sep 25 14:04 volls3
-> ../../../../devices/pseudo/vxlo@0:volls3,2,blk

# ls -l /dev/vxlo/rdsk/
lrwxrwxrwx 1 root root 46 Sep 25 14:04 volls0
-> ../../../../devices/pseudo/vxlo@0:volls0,1,raw
lrwxrwxrwx 1 root root 46 Sep 25 14:04 volls3
-> ../../../../devices/pseudo/vxlo@0:volls3,2,raw
```

Use the `vxloadm get` command to display the list of all currently mapped partition(s) created using the `vxloadm` utility. For example:

```
# /etc/vx/bin/vxloadm get
VxVM  INFO V-5-1-0      NAME      FILENAME
MOUNT  OFFSET  C/H/S
VxVM  INFO V-5-1-15260  volls0    /dev/vx/dsk/testdg/voll
6180      6787/1/618
VxVM  INFO V-5-1-15260  volls3    /dev/vx/dsk/testdg/voll
4326000  50902/1/618
```

Use the appropriate file system commands to access the file system(s) For example:

```
# fstyp /dev/vxlo/rdsk/volls0
ufs
# mount -F ufs /dev/vxlo/dsk/volls0 /mnt
```

Use the `vxloadm delete` to remove the partition mappings of a volume. For example:

```
# /etc/vx/bin/vxloadm delete volls0
# /etc/vx/bin/vxloadm delete volls3
```

Note: This `vxloadm` utility should only be used on volumes that are currently not in use or held open by a guest domain.

The `vxformat` utility automatically relabels the virtual disk backed by a VxVM volume in the guest domain

The `vxformat` utility provides the user the ability to automatically relabel a virtual disk backed by a VxVM volume. This utility is meant to be executed from inside the guest domain only.

The `vxformat` utility is particularly useful when a VxVM volume with existing partitions is grown in size and you need to access the new size from the guest domain.

Requirements for relabeling to succeed

- The relabel succeeds only- if it can find a new cylinder size that is aligned with the start and size of each of the existing partitions.

In case the `vxformat` command cannot find such cylinder size, it displays the following descriptive message and then exits:

```
Cannot re-label device /dev/rdisk/c0t1d2s2 since we failed to  
find new cylinder size that's aligned with all the existing partitions
```

- The relabel succeeds only - if the available blocks is greater than the last sector of each and every non-s2 partition.

Otherwise, the `vxformat` command displays the following message and then exits:

```
Cannot re-label device /dev/rdisk/c0d2s2 since the last sector of a  
non-s2 partition is greater than the available blocks
```

Example of using the `vxformat` utility

Use the `vxformat` command to relabel the virtual disk. For example:

```
# /etc/vx/bin/vxformat c0d1s2  
rawpath: /dev/rdisk/c0d1s2  
Old disk capacity: 2097000 blocks  
New disk capacity: 4194000 blocks  
Device /dev/rdisk/c0d1s2 has been successfully re-labeled.  
Please use prtvtoc(1) to obtain the latest partition table information
```

If the underlying device size has not changed, the `vxformat` command displays the following message without changing the label. For example:

```
# /etc/vx/bin/vxformat c0d1s2  
Old disk capacity: 2343678 blocks  
New disk capacity: 2343678 blocks  
size of device /dev/rdisk/c0d2s2 is unchanged
```

LDom new features

The following new features apply for Logical Domains.

Guest domain migration

LDom 1.1 introduces the guest domain migration feature. The guest domain migration feature is supported for both cold and warm migrations by Storage Foundation with both deployment models:

- Split Storage Foundation stack
- Guest-based Storage Foundation stack

Virtual I/O dynamic reconfiguration

LDom 1.1 introduces the virtual I/O dynamic reconfiguration feature. The virtual I/O dynamic reconfiguration feature is supported with both deployment models:

- Split Storage Foundation stack
- Guest-based Storage Foundation stack

Note: For resizing a volume exported as a single slice: The new size should be visible dynamically in the guest immediately.

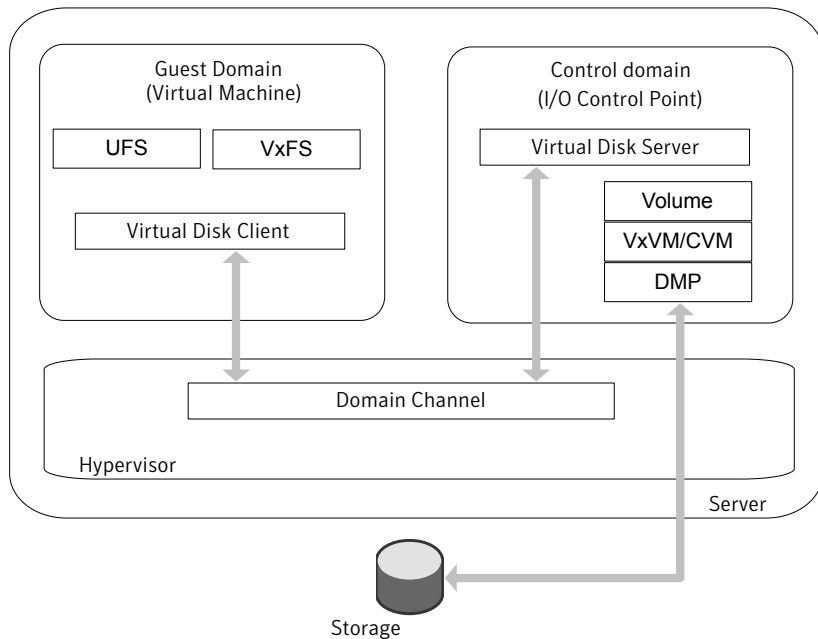
For resizing a volume exported as a full disk: Even though the new size is visible dynamically in the guest, the new space allocated in the volume cannot be utilized unless the label in the vdisk has been adjusted to reflect the new sectors. This adjustment of the label needs to be done carefully.

Symantec has developed a utility to automatically do the relabeling of the vdisk in the guest that will be available in an upcoming release.

Split Storage Foundation stack model

[Figure 4-1](#) illustrates the split Storage Foundation stack model with Solaris Logical Domains.

Figure 4-1 Split Storage Foundation stack model with Solaris Logical Domains



How Storage Foundation and High Availability Solutions works in the Solaris Logical Domains

Veritas Storage Foundation and High Availability Solutions supports Solaris Logical Domains in both single-node, multiple-node, and multiple-node high availability configurations.

[Figure 4-1](#) illustrates the recommended placement of Storage Foundation stack component products in this model.

Following indicates the recommended placement of Storage Foundation stack component products:

- For a single node configuration, Veritas Volume Manager (VxVM) including DMP is placed in the control domain, and Veritas File System (VxFS) is placed in the guest domain.
- For clustered nodes, Cluster Volume Manager (CVM) is placed in the control domain, and VxFS is placed in the guest domain.
See [“Clustering using Cluster Volume Manager”](#) on page 89.
See [“Installing Storage Foundation on multiple nodes in a Logical Domain”](#) on page 90.

See “[Cluster Volume Manager in the control domain for providing high availability](#)” on page 92.

- For clustered nodes in a highly available environment, install Veritas Cluster Server (VCS) in the control domain.
See “[About Veritas Cluster Server in a Logical Domain environment](#)” on page 95.
See “[About Veritas Cluster Server configuration models in a Logical Domain environment](#)” on page 99.
See “[Configuring Veritas Cluster Server to fail over a Logical Domain on a failure](#)” on page 104.
See “[Configuring Veritas Cluster Server to fail over an application on a failure](#)” on page 111.
- VxFS drivers in the guest domain cannot currently interact with the VxVM drivers in the control domain. This renders some features, which require direct VxVM-VxFS coordination, unusable in such a configuration.
See “[Veritas Storage Foundation features restrictions](#)” on page 61.

Note: VxFS can also be placed in the control domain, but there will be no coordination between the two VxFS instances in the guest and the control domain.

Veritas Storage Foundation features restrictions

The following Veritas Storage Foundation software features are restricted in the split Storage Foundation stack model:

- VxVM volume snapshots – Due to the inability of VxFS in the guest domain to coordinate with VxVM in the control domain, taking a data consistent snapshot of a VxVM volume containing a VxFS file system requires shutting down the application and unmounting the file system before taking the snapshot.
- Resizing VxVM volumes and any type of file system on top of the volume with `vxresize` – Resizing any type of file system on the guest whose underlying device is backed by a VxVM volume in the control domain, requires resizing the VxVM volume and the file system in the guest individually.
If you are growing a VxFS file system in the guest whose underlying device is backed by a VxVM volume requires you to first grow the volume in the control domain using the `vxassist` command, and then the file system in the guest LDom using the `fsadm` command.

Shrinking a VxFS file system, on the other hand, requires you to first shrink the file system in the guest LDom using the `fsadm` command, and then the volume in the control domain using the `vxassist` command. Using the `vxassist` command requires you to use the `-f` option of the command, as in the following example.

```
# vxassist -g [diskgroup] -f shrinkto volume length
```

Caution: Do not shrink the underlying volume beyond the size of the VxFS file system in the guest as this can lead to data loss.

- Exporting a volume set to a guest LDom is not supported.
 - Veritas Volume Replicator is not supported in the Split Storage Foundation stack model.
 - Multi-volume filesets/DST
 - File-level Smartsync
 - The following VxFS tunables are not set to their default values based on the underlying volume layout, due to VxFS being in the guest LDom and VxVM being installed in the control domain:
 - `read_pref_io`
 - `write_pref_io`
 - `read_nstream`
 - `write_nstream`
- If desired, you can set the values of these tunables based on the underlying volume layout in the `/etc/vx/tunefstab` file.
- See the *Veritas File System Administrator's Guide* for more information about tuning I/O.
- Storage Foundation Cluster File System is not recommended in this deployment model.

Guest-based Storage Foundation stack model

Figure 4-2 Guest-based Storage Foundation stack model

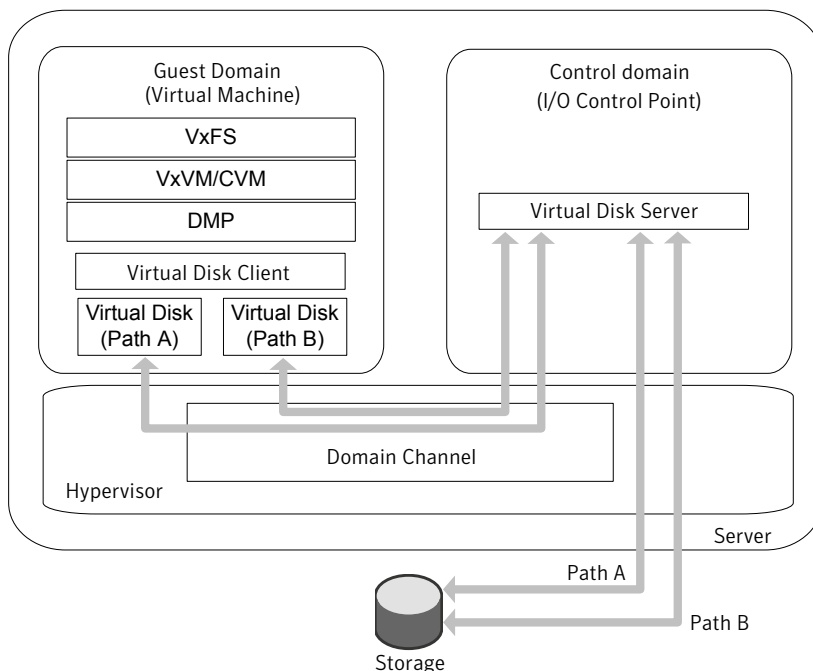


Figure 4-2 illustrates the guest-based Storage Foundation stack model with Solaris Logical Domains.

How Storage Foundation and High Availability Solutions works in the guest Logical Domains

The entire Storage Foundation stack is co-located within the guest in this deployment model.

Symantec recommends that you export all paths to a disk which is being exported to a guest and let Veritas DMP do the multi-pathing of the disk in the guest domain.

Note: Only full SCSI disks can be used under Veritas Volume Manager (VxVM) and DMP in this model. Non-SCSI devices (volume, file, slice, etc) are not supported.

Veritas Storage Foundation and High Availability Solutions and Veritas Storage Foundation Cluster File System supports running in the guest Logical Domains

in both single-node, multiple-node, and multiple-node high availability configurations.

- For a single node configuration, VxVM (including DMP) and VxFS are co-located in the guest LDom.
- For clustered nodes, CVM can also be used inside the guest LDom. As with regular physical nodes, forming a CVM cluster of LDom guests requires shared storage visibility across the guests that are part of the cluster. See the *Veritas Volume Manager Administrator's Guide* for CVM information. See the *Veritas Storage Foundation Installation Guide* for installation and configuration information.
- For clustered nodes in a highly available environment, install Veritas Cluster Server (VCS) in the guest LDom.
See the *Veritas Cluster Server* documentation for more information.

Veritas Volume Replicator (VVR) is supported in the guest-based Storage Foundation stack model in the following configurations:

- A guest domain on one host acting as the VVR primary, and another guest on another host acting as the VVR secondary.
- Two guest LDom on the same physical host, but you must export separate LUNs or disks to the data volumes and Storage Replicator Logs of the two guest domains.

Supported configurations with SFCFS and multiple I/O Domains

Figure 4-3 SFCFS cluster across two guest domains

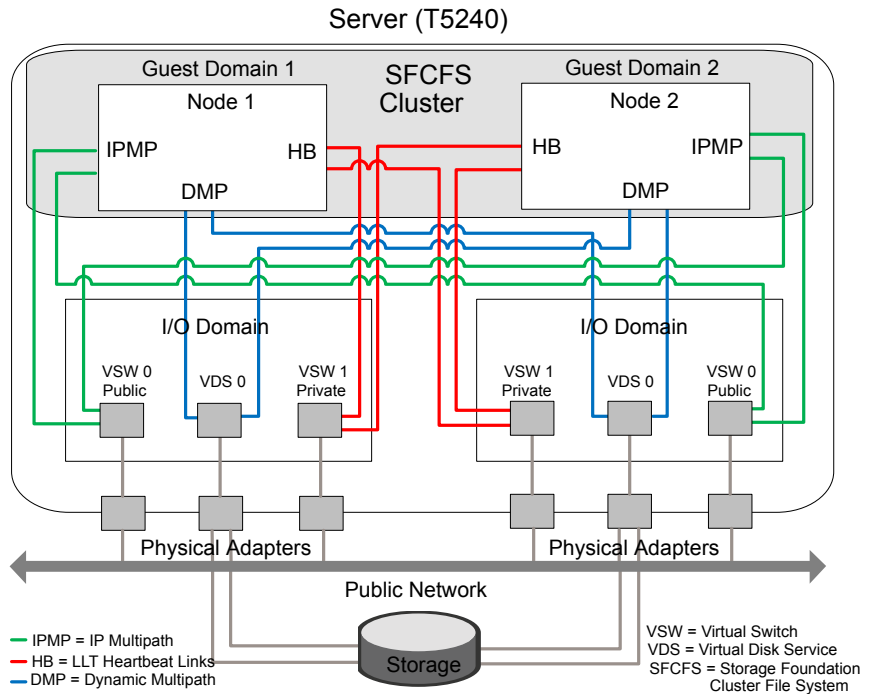


Figure 4-3 illustrates that each guest domain gets network and disk storage redundancy from the two I/O domains.

Figure 4-4 SFCFS cluster across two guest domains

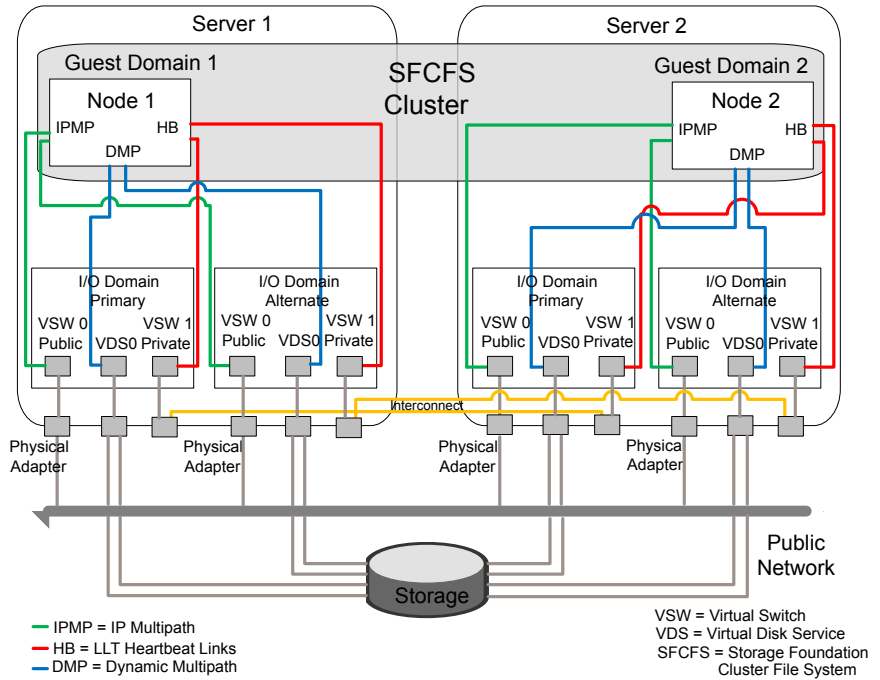


Figure 4-4 illustrates that each guest domain gets network and disk storage redundancy from the two I/O domains on that physical server. The guest cluster spans across two physical servers.

Figure 4-5 SFCFS cluster across four guest domains

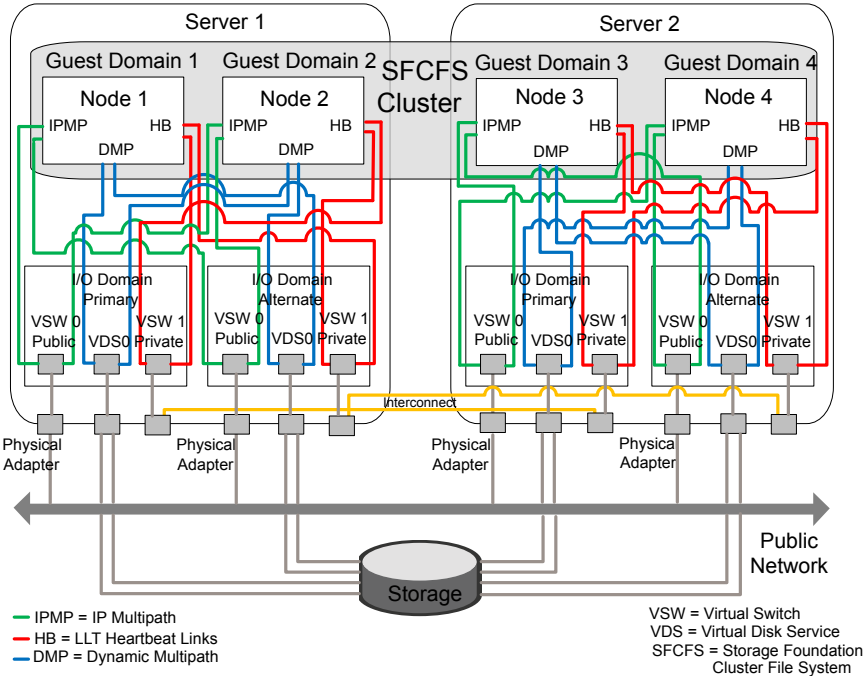


Figure 4-5 illustrates that each guest domain gets network and disk storage redundancy from two I/O domains on that physical server. The guest cluster spans across two physical servers.

Figure 4-6 SFCFS cluster across four guest domains

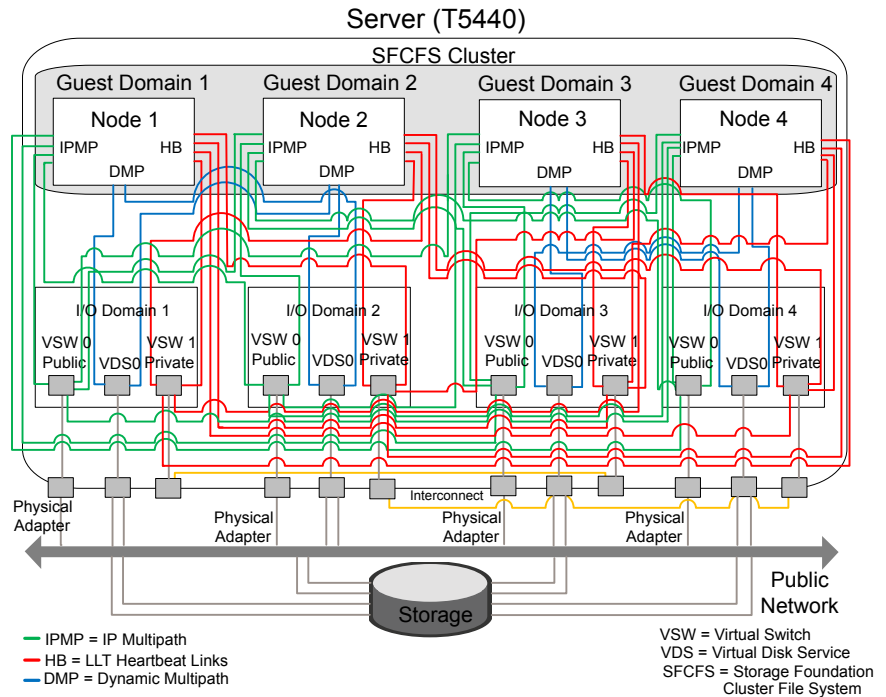


Figure 4-6 illustrates each guest gets its disk storage redundancy from two out of the four I/O domains. Each guest gets its network redundancy from all the four I/O domains.

Veritas Storage Foundation features restrictions

The following Veritas Storage Foundation software features are restricted in the Solaris LDom guest environment.

Veritas Volume Replicator bunker replication

Veritas Volume Replicator (VVR) currently does not support configuring a guest domain as a bunker node in the bunker replication mode. This restriction may be removed in a later release.

Mirroring across controllers using vxassist the mirror=ctrl option

Currently, all virtual disks in the guest are under the same virtual controller c0. When `vxassist` tries to look for a second controller to place the mirror on, it fails and therefore the command fails.

All disks fall under the c0 controller, even if they are exported using different physical paths in the backend and coming from different HBAs.

DMP SCSI bypass

The virtual disk client (VDC) driver controls all virtual disks in the Guest LDom and not the SCSI driver.

Therefore, it is not possible to construct a SCSI packet and send the packet down through DMP. Setting the tunable `dmp_fast_recovery` to on has no effect.

Event Source Daemon (vxesd) fabric monitoring capabilities

One of the features of the `vxesd` daemon is to register with the HBA API to listen to fabric events. Even though the HBA API is loaded in the guest, since there is currently no concept of a directly connected HBA in the guest, this API is not of any use. Therefore, this capability of `vxesd` is not available.

Physical WWN for a path

It is not possible to create a Sub Path Failover Group (SFG) without Physical WWN. Physical World Wide IDs cannot be obtained from inside the guest because they are fetched by DMP using the HBA API which is currently not functional inside the guest.

System requirements

This section describes the system requirements for this release.

Solaris operating system requirements

Veritas Storage Foundation 5.1 with support for Logical Domains is supported on the following Solaris operating systems:

- Solaris 10 (SPARC Platform) Update 7 or later

Note: Symantec recommends installing the latest Solaris 10 Update.

- LDom software from Sun, version 1.2 or later Visit the Sun Microsystems, Inc. Web site for information about Solaris operating environments at the following URL:
<http://www.sun.com/servers/coolthreads/ldoms/get.jsp>

Solaris patch requirements

Some required system patches may already be present in your operating system. Check if your system already contains the required Solaris patches needed before installing Solaris Logical Domains. Use the `showrev -p` command to display the patches included on your system.

For more information, see the `showrev(1M)` manual page.

If the patches shown in the required list are not already installed, go to the SunSolve website to download them. You must install the appropriate patches and then reboot your system.

Warning: Patch version and information are determined at the time of product release. Contact your vendor for the most current patch version and information.

Table 4-2 Sun patches

Sun patch number	Notes
141777-01	<p>This patch fixes the following SUN bugs that affects SF functionality:</p> <ul style="list-style-type: none"> ■ SUN bug id: 6795836 – <code>vd_setup_vd()</code> should handle errors from <code>vd_identify_dev()</code> better
139562-02 (obsoleted by 138888-07)	<p>This patch fixes the following SUN bugs that affects SF functionality:</p> <ul style="list-style-type: none"> ■ SUN bug id: 6640564 – <code>vds</code> should not serialize disk image IO ■ SUN bug id: 6699271 – Dynamic virtual disk size management ■ SUN bug id: 6705190 – <code>uscsicmd</code> on <code>vdisk</code> can overflow sense buffer ■ SUN bug id: 6716365 – disk images on volumes should be exported using the <code>ldi</code> interface

Table 4-2 Sun patches (*continued*)

Sun patch number	Notes
138042-05	<p>This patch fixes the following SUN bugs that affects SF functionality:</p> <ul style="list-style-type: none">■ SUN bug id: 6637560 – disks are not correctly exported with vxddmp■ SUN bug id: 6558966 – virtual disks created from files do not support EFI labels <p>Sun Microsystems recommends installing the latest patches in all domains to ensure that you have the most up-to-date drivers.</p>

Sun hardware requirements

Visit the Sun Microsystems, Inc. Web site for information about Sun hardware LDom support at the following URL:

<http://www.sun.com/servers/coolthreads/ldoms/index.xml>

Component product release notes

Read the relevant component product release notes before installing any version of Veritas Storage Foundation and High Availability or Veritas Storage Foundation Cluster File System.

Veritas Storage Foundation and High Availability

Release notes for component products in all versions of Veritas Storage Foundation 5.1 are located under the `storage_foundation/docs` directory and the Veritas Cluster Server 5.1 are located under the `cluster_server/docs` directory of the Veritas Storage Foundation disc.

- *Veritas Storage Foundation Release Notes* (`sf_notes.pdf`)
- *Veritas Cluster Server Release Notes* (`vcs_notes.pdf`)

Veritas Storage Foundation Cluster File System

Release notes for component products in all versions of Veritas Storage Foundation Cluster File System 5.1 are located under the `storage_foundation/docs` directory of the Veritas Storage Foundation disc.

- *Veritas Storage Foundation Release Notes* ([sf_notes.pdf](#)). This document includes information for Veritas Storage Foundation Cluster File System.

Product licensing

Customers running Veritas Storage Foundation or Veritas Storage Foundation Cluster File System in a Solaris LDom environment are entitled to use an unlimited number of logical domains on each licensed server or CPU.

Installing Storage Foundation in a Logical Domain environment

This section describes how to install Storage Foundation in several LDom environments.

To install the Split Storage Foundation stack model environment, you must complete the following sections in order:

- See [“Installing and configuring the Logical Domain software and domains”](#) on page 73.
- See [“Installing Storage Foundation in the control domain or guest”](#) on page 73.
- See [“Installing Veritas File System in the guest domain using pkgadd”](#) on page 74.
- See [“Verifying the configuration”](#) on page 74.

To install the Guest based Storage Foundation stack model environment, which includes Veritas Storage Foundation Cluster File System, you must complete the following sections in order:

- See [“Installing and configuring the Logical Domain software and domains”](#) on page 73.
- See [“Installing Storage Foundation in the control domain or guest”](#) on page 73.
- See [“Verifying the configuration”](#) on page 74.

To install and configure Veritas Cluster Server in a Solaris Logical Domains environment, see the following sections:

- See [“Configuring Veritas Cluster Server to fail over a Logical Domain on a failure”](#) on page 104.
- See [“Configuring Veritas Cluster Server to fail over an application on a failure”](#) on page 111.

Installing and configuring the Logical Domain software and domains

Refer to the Sun Microsystems documentation for instructions about installing and configuring the Logical Domain software and configuring the control and guest domains.

See the *Logical Domains 1.1 Administration Guide*.

Installing Storage Foundation in the control domain or guest

This section describes how to install Storage Foundation in the control domain or guest domain.

Installing the split Storage Foundation stack model

If you are installing the split Storage Foundation stack model, the entire stack must be in the control domain and VxFS must be in the guest domain.

Use the procedures in the Veritas installation documentation and Release Notes to install Storage Foundation in the control domain.

To install the split Storage Foundation stack model

- ◆ Install the product.

See the *Veritas Storage Foundation Installation Guide for Solaris*.

See the *Veritas Storage Foundation Release Notes for Solaris*.

Installing the guest-based Storage Foundation stack model

If you are installing the guest-based Storage Foundation stack model environment, the entire stack must be in the guest domain.

Use the procedures in the SF or SFCFS Installation Guide and Release Notes, as appropriate, to install SF or SFCFS in a guest Logical domain.

To install the guest-based Storage Foundation stack model

- ◆ Install the product.

See the *Veritas Storage Foundation Installation Guide for Solaris* for SF.

See the *Veritas Storage Foundation Release Notes for Solaris* for SF.

See the *Veritas Storage Foundation Cluster File System Installation Guide for Solaris* for SFCFS.

See the *Veritas Storage Foundation Cluster File System Release Notes for Solaris* for SFCFS.

Installing Veritas File System in the guest domain using pkgadd

This section describes how to install VxFS 5.1 in the guest domain using the `pkgadd` command.

To install Veritas File System 5.1 in the guest domain using pkgadd

1 Copy the VxFS packages from the `/packages` directory on the disc to a location in the guest domain where you have write permissions.

2 Change to the `pkgs` directory:

```
# cd pkgs
```

3 Install the packages:

```
# pkgadd -d VRTSvlic.pkg
```

```
# pkgadd -d VRTSvxfs.pkg
```

```
# pkgadd -d VRTSfssdk.pkg
```

4 Reboot the system.

Verifying the configuration

Verify the configuration of Logical Domains in the control domain and the guest domain. Refer to the Sun Microsystems documentation for details.

See *Logical Domains Administration Guide*.

Verify the Storage Foundation installation in both the control domain and the guest domain.

See the *Veritas Storage Foundation Installation Guide for Solaris*.

See the *Veritas Storage Foundation Release Notes for Solaris*.

See the *Veritas Storage Foundation Cluster File System Installation Guide for Solaris*.

See the *Veritas Storage Foundation Cluster File System Release Notes for Solaris*.

Migrating a Veritas Volume Manager disk group from a non-Logical Domain environment

Use the following procedure to migrate a VxVM disk group from a non-LDom environment to an LDom environment.

Note: This section applies to only the Split Storage Foundation model.

In the following example control domain is named “primary” and the guest domain is named “ldom1.” The prompts in each step show in which domain to run the command.

To create virtual disks on top of the Veritas Volume Manager data volumes using the `ldm` command

- 1 The VxVM diskgroup on the target LDom host is imported in the control domain, after which volumes are visible from inside the control domain.

See the *Veritas Volume Manager Administrator’s Guide* to move disk groups between systems.

- 2 In the control domain (primary), configure a service exporting the VxVM volume containing a VxFS or UFS filesystem as a slice using the `options=slice` option:

```
primary# ldm add-vdiskserverdevice options=slice \  
/dev/vx/dsk/dg-name/volume_name  
volume_namevolume_name@primary-vds0
```

Caution: With Solaris 10, Update 5 and LDom 1.1, a volume by default shows up as a full disk in the guest. The Virtual Disk Client driver writes a VTOC on block 0 of the virtual disk, which will end up as a WRITE on block 0 of the VxVM volume. This can potentially cause data corruption, because block 0 of the VxVM volume contains user data. Using `options=slice` exports a volume as a slice to the guest and does not cause any writes to block 0, therefore preserving user data.

- 3 Add the exported disk to a guest LDom:

```
primary# ldm add-vdisk vdisk1 volume_name  
volume_name@primary-vds0 ldom1
```

- 4 Start the guest domain, and ensure that the new virtual disk is visible.

```
primary# ldm bind ldom1  
  
primary# ldm start ldom1
```

- 5 If the new virtual disk device node entries do not show up in the `/dev/[r]dsk` directories, then run the `devfsadm` command in the guest domain:

```
ldom1# devfsadm -C
```

In this example, the new disk appears as `/dev/[r]dsk/c0d1s0`.

```
ldom1# ls -l /dev/dsk/c0d1s0
```

```
lrwxrwxrwx 1 root root 62 Sep 11 13:30 /dev/dsk/c0d1s0 ->
../../../../devices/virtual-devices@100/channel-devices@200/disk@1:a
```

- 6 Mount the file system on the disk to access the application data:

```
ldom1# mount -F vxfs /dev/dsk/c0d1s0 /mnt
```

```
ldom1# mount -F ufs /dev/dsk/c0d1s0 /mnt
```

Provisioning storage for a guest Logical Domain

Use the following procedure to provision storage for a guest LDom. You can provision both boot disks and data disks.

Note: This section applies to the Split Storage Foundation stack model only.

For the guest-based Storage Foundation model:

See [“How Storage Foundation and High Availability Solutions works in the guest Logical Domains”](#) on page 63.

Provisioning Veritas Volume Manager volumes as data disks for guest Logical Domains

The following procedure uses VxVM volumes as data disks (virtual disks) for guest LDom.

VxFS can be used as the file system on top of the data disks.

The example control domain is named “primary” and the guest domain is named “ldom1.” The prompts in each step show in which domain to run the command.

To provision Veritas Volume Manager volumes as data disks

- 1 Create a VxVM disk group (`mydatadg` in this example) with some disks allocated to it:

```
primary# vxvg init mydatadg TagmaStore-USP0_29 TagmaStore-USP0_30
```

- 2 Create a VxVM volume of the desired layout (in this example, creating a simple volume):

```
primary# vxassist -g mydatadg make datavol1 500m
```

- 3 Configure a service exporting the volume `datavol1` as a virtual disk:

```
primary# ldm add-vdiskserverdevice /dev/vx/dsk/mydatadg/datavol1 \
datadisk1@primary-vds0
```

- 4 Add the exported disk to a guest domain.

```
primary# ldm add-vdisk vdisk1 datadisk1@primary-vds0 ldom1
```

- 5 Start the guest domain, and ensure that the new virtual disk is visible:

```
primary# ldm bind ldom1
```

```
primary# ldm start ldom1
```

- 6 If the new virtual disk device node entries do not show up in the `/dev/[r]dsk` directories, then run the `devfsadm` command in the guest domain:

```
ldom1# devfsadm -C
```

- 7 Label the disk using the `format` command to create a valid label before trying to access it.

See the `format(1M)` manual page.

- 8 Create the file system where `c0d1s2` is the disk.

```
ldom1# mkfs -F vxfs /dev/rdsk/c0d1s2
```

9 Mount the file system.

```
ldom1# mount -F vxfs /dev/dsk/c0d1s2 /mnt
```

10 Verify that the file system has been created:

```
ldom1# df -hl -F vxfs
```

```
Filesystem size used avail capacity Mounted on  
/dev/dsk/c0d1s2 500M 2.2M 467M 1% /mnt
```

Provisioning Veritas Volume Manager volumes as boot disks for guest Logical Domains

The following procedure provisions boot disks for a guest domain.

With Solaris 10 Update 5 and LDoms 1.1, a VxVM volume appears as a full disk by default and can be used as a boot disk for a guest LDom.

The following process gives the outline of how a VxVM volume can be used as a boot disk.

The example control domain is named “primary” the guest domain is named “ldom1.” The prompts in each step show in which domain to run the command.

To provision Veritas Volume Manager volumes as boot disks for guest Logical Domains

- 1** On the control domain, create a VxVM volume of a size that is recommended for Solaris 10 installation. In this example, a 7GB volume is created:

```
primary# vxassist -g boot_dg make bootdisk-vol 7g
```

- 2** Configure a service by exporting the /dev/vx/dsk/boot_dg/bootdisk1-vol volume as a virtual disk:

```
primary# ldm add-vdiskserverdevice \  
/dev/vx/dsk/boot_dg/bootdisk1-vol bootdisk1-vol@primary-vds0
```

- 3** Add the exported disk to ldom1:

```
primary# ldm add-vdisk vdisk1 bootdisk1-vol@primary-vds0 ldom1
```

- 4** Follow Sun Microsystems’ recommended steps to install and boot a guest domain, and use the virtual disk vdisk1 as the boot disk during the net install.

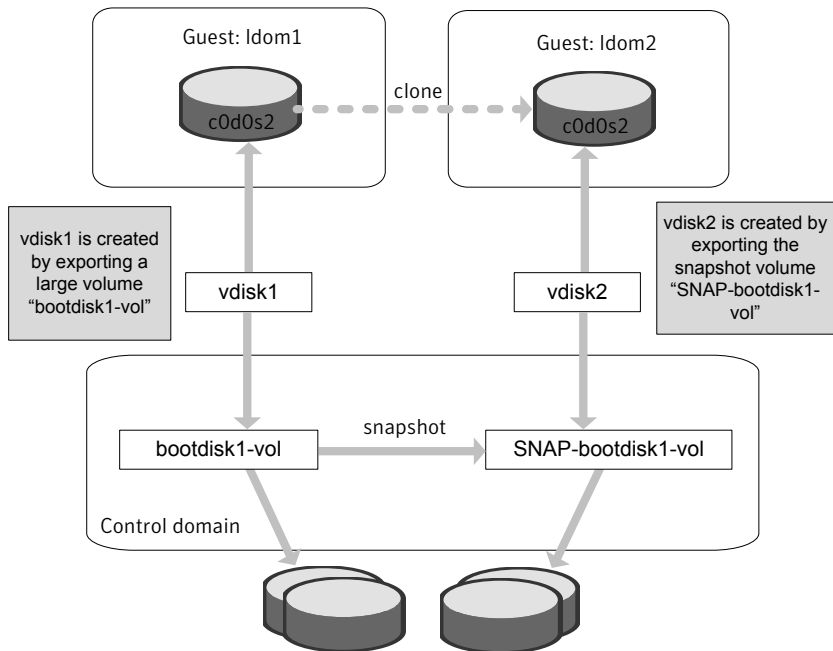
Using Veritas Volume Manager snapshots for cloning Logical Domain boot disks

The following highlights the steps to clone the boot disk from an existing LDom using VxVM snapshots, and makes use of the third-mirror breakoff snapshots.

See “[Provisioning Veritas Volume Manager volumes as boot disks for guest Logical Domains](#)” on page 78.

Figure 4-7 illustrates an example of using Veritas Volume Manager snapshots for cloning Logical Domain boot disks.

Figure 4-7 Example of using Veritas Volume Manager snapshots for cloning Logical Domain boot disks



Before this procedure, `ldom1` has its boot disk contained in a large volume, `/dev/vx/dsk/boot_dg/bootdisk1-vol`.

This procedure involves the following steps:

- Cloning the LDom configuration to form a new LDom configuration. This step is a Solaris LDom procedure, and can be achieved using the following commands:

```
# ldm list-constraints -x
```

```
# ldm add-domain -i
```

Refer to the Sun Microsystems documentation for more information about cloning the LDom configuration to form a new LDom configuration.

See the *Logical Domains 1.1 Administration Guide*.

- After cloning the configuration, clone the boot disk and provision it to the new LDom.

To create a new LDom with a different configuration than that of `ldom1`, skip this step of cloning the configuration and create the desired LDom configuration separately.

To clone the boot disk using Veritas Volume Manager snapshots

- 1 Create a third-mirror breakoff snapshot of the source volume bootdisk1-vol. To create the snapshot, you can either take some of the existing ACTIVE plexes in the volume, or you can use the following command to add new snapshot mirrors to the volume:

```
primary# vxsnap [-b] [-g diskgroup] addmir volume [nmirror=N]  
\ [alloc=storage_attributes]
```

By default, the `vxsnap addmir` command adds one snapshot mirror to a volume unless you use the `nmirror` attribute to specify a different number of mirrors. The mirrors remain in the SNAPATT state until they are fully synchronized. The `-b` option can be used to perform the synchronization in the background. Once synchronized, the mirrors are placed in the SNAPDONE state.

For example, the following command adds two mirrors to the volume, bootdisk1-vol, on disks mydg10 and mydg11:

```
primary# vxsnap -g boot_dg addmir bootdisk1-vol nmirror=2  
\ alloc=mydg10,mydg11
```

If you specify the `-b` option to the `vxsnap addmir` command, you can use the `vxsnap snapwait` command to wait for synchronization of the snapshot plexes to complete, as shown in the following example:

```
primary# vxsnap -g boot_dg snapwait bootdisk1-vol nmirror=2
```

- 2 To create a third-mirror break-off snapshot, use the following form of the `vxsnap make` command.

Caution: Shut down the guest domain before executing the `vxsnap` command to take the snapshot.

```
primary# vxsnap [-g diskgroup] make \ source=volume[/newvol=snapvol]
\ {/plex=plex1[,plex2,...]|/nmirror=number}
```

Either of the following attributes may be specified to create the new snapshot volume, `snapvol`, by breaking off one or more existing plexes in the original volume:

plex Specifies the plexes in the existing volume that are to be broken off. This attribute can only be used with plexes that are in the ACTIVE state.

nmirror Specifies how many plexes are to be broken off. This attribute can only be used with plexes that are in the SNAPDONE state. Such plexes could have been added to the volume by using the `vxsnap addmir` command.

Snapshots that are created from one or more ACTIVE or SNAPDONE plexes in the volume are already synchronized by definition.

For backup purposes, a snapshot volume with one plex should be sufficient.

For example,

```
primary# vxsnap -g boot_dg make
\ source=bootdisk1-vol/newvol=SNAP-bootdisk1-vol/nmirror=1
```

Here `bootdisk1-vol` makes source; `SNAP-bootdisk1-vol` is the new volume and `1` is the `nmirror` value.

The block device for the snapshot volume will be

```
/dev/vx/dsk/boot_dg/SNAP-bootdisk1-vol.
```

- 3 Configure a service by exporting the `/dev/vx/dsk/boot_dg/SNAP-bootdisk1-vol` file as a virtual disk.

```
primary# ldm add-vdiskserverdevice \
/dev/vx/dsk/boot_dg/SNAP-bootdisk1-vol vdisk2@primary-vds0
```

- 4 Add the exported disk to `ldom1` first.

```
primary# ldm bind ldom1
```

```
primary# ldm start ldom1
```

- 5 Start `ldom1` and boot `ldom1` from its primary boot disk `vdisk1`.

```
primary# ldm bind ldom1
```

```
primary# ldm start ldom1
```

- 6 If the new virtual disk device node entries do not show up in the `/dev/[r]dsk` directories, then run the `devfsadm` command in the guest domain:

```
ldom1# devfsadm -C
```

where `vdisk2` is the `c0d2s#` device.

```
ldom1# ls /dev/dsk/c0d2s*
```

```
/dev/dsk/c0d2s0 /dev/dsk/c0d2s2 /dev/dsk/c0d2s4 /dev/dsk/c0d2s6  
/dev/dsk/c0d2s1 /dev/dsk/c0d2s3 /dev/dsk/c0d2s5 /dev/dsk/c0d2s7
```

- 7 Mount the root file system of `c0d2s0` and modify the `/etc/vfstab` entries such that all `c#d#s#` entries are changed to `c0d0s#`. You must do this because `ldom2` is a new LDom and the first disk in the OS device tree is always named as `c0d0s#`.

- 8 After you change the `vfstab` file, unmount the file system and unbind `vdisk2` from `ldom1`:

```
primary# ldm remove-vdisk vdisk2 ldom1
```

- 9 Bind `vdisk2` to `ldom2` and then start and boot `ldom2`.

```
primary# ldm add-vdisk vdisk2 vdisk2@primary-vds0 ldom2
```

```
primary# ldm bind ldom2
```

```
primary# ldm start ldom2
```

After booting `ldom2`, appears as `ldom1` on the console because the other host-specific parameters like `hostname` and IP address are still that of `ldom1`.

```
ldom1 console login:
```

- 10 To change the parameters bring `ldom2` to single-user mode and run the `sys-unconfig` command.

11 Reboot ldom2.

During the reboot, the operating system prompts you to configure the host-specific parameters such as hostname and IP address, which you must enter corresponding to ldom2.

12 After you have specified all these parameters, ldom2 boots successfully.

Software limitations

The following section describes some of the limitations of the Solaris Logical Domains software and how those software limitations affect the functionality of the Veritas Storage Foundation products.

Memory Corruption in the guest Logical Domain during SCSI commands

When running VxVM and DMP the system panics in the guest LDom. The following error messages display on the console:

```
vxldmp: NOTICE: VxVM vxldmp V-5-3-289 DMP: Memory Overrun!  
Caller dmpscsi_sys.c(170) Ptr 0x3000528e580 Size 0x100
```

```
vxldmp: NOTICE: VxVM vxldmp V-5-3-289 DMP: Memory Overrun!  
Caller dmppgrio.c(824) Ptr 0x3002d7f1a80 Size 0xe8
```

These error messages are due to a kernel memory corruption occurring in the Solaris kernel driver stacks (virtual disk drivers). This issue occurs when issuing USCSICMD with the sense request enable (USCSI_RQENABLE) on a virtual disk from the guest.

Symantec has an open escalation with Sun Microsystems and an associated SUN bug id for this issue:

SUN Escalation number: 1-23915211

SUN bug id: 6705190 (ABS: uscsicmd on vdisk can overflow the sense buffer)

<http://sunsolve.sun.com/search/document.do?assetkey=1-1-6705190-1>

This SUN bug has been fixed in Sun patch 139562-02.

See “[Solaris patch requirements](#)” on page 70.

Exporting the raw volume device node fails

Following messages can be observed in the `/var/adm/messages` file:

```
vds: [ID 998409 kern.info] vd_setup_vd():  
/dev/vx/rdisk/testdg/testvol identification failed  
vds: [ID 790452 kern.info] vd_setup_vd():  
/dev/vx/rdisk/testdg/testvol can not be exported as a virtual disk  
(error 5)  
vds: [ID 556514 kern.info] vds_add_vd(): Failed to add vdisk ID 21  
vds: [ID 970787 kern.info] vds_add_vd(): No vDisk entry found for  
vdisk ID 21
```

Workaround: Export VxVM volumes using their block device nodes instead. This issue is under investigation by Sun Microsystems.

SUN bug id: 6716365 (disk images on volumes should be exported using the ldi interface)

This SUN bug is fixed in Sun patch 139562-02.

See [“Solaris patch requirements”](#) on page 70.

Resizing a Veritas Volume Manager volume (exported as a slice or full disk) does not dynamically reflect the new size of the volume in the guest

On resizing a VxVM volume exported to a guest, the virtual disk still shows the old size of the volume. The virtual disk drivers do not update the size of the backend volume after the volume is resized.

Sun Microsystems has an RFE for this issue (CR 6699271 Dynamic virtual disk size management).

Workaround: The guest must be stopped and rebound for the new size to be reflected.

This SUN bug is fixed in Sun patch 139562-02.

See [“Solaris patch requirements”](#) on page 70.

Known issues

The following section describes some of the known issues of the Solaris Logical Domains software and how those known issues affect the functionality of the Veritas Storage Foundation products.

Guest-based known issues

The following are new known issues in this release of Veritas Storage Foundation and High Availability Solutions Support for Solaris Logical Domains.

Encapsulating a non-scsi disk may fail

Trying to encapsulate a non-scsi disk which is a slice of a disk or a disk exported as a slice may fail with the following error:

```
VxVM vxslicer ERROR V-5-1-599 Disk layout does not support swap shrinking  
VxVM vxslicer ERROR V-5-1-5964 Unsupported disk layout.  
Encapsulation requires atleast 0 sectors of unused space either at the  
beginning or end of the disk drive.
```

This is because while installing the OS on such a disk, it is required to specify the entire size of the backend device as the size of slice "s0", leaving no free space on the disk.

Boot disk encapsulation requires free space at the end or the beginning of the disk for it to proceed ahead.

Guest LDom node shows only 1 PGR key instead of 2 after rejecting the other node in the cluster

For configuration information concerning the guest LDom node shows only 1 PGR key instead of 2 after rejecting the other node in the cluster:

See [Figure 4-4](#) on page 66.

This was observed while performing a series of reboots of the primary and alternate I/O domains on both the physical hosts housing the two guests. At some point one key is reported missing on the coordinator disk.

This issue is under investigation. The `vxfen` driver can still function as long as there is 1 PGR key. This is a low severity issue as it will not cause any immediate interruption. Symantec will update this issue when the root cause is found for the missing key.

Disk paths intermittently go offline while performing I/O on a mirrored volume

This was observed while testing the SFCFS stack inside a 4-node guest cluster where each node gets its network and virtual disk resources from multiple I/O domains within the same host.

See [Figure 4-6](#) for the exact configuration.

While performing I/O on a mirrored volume inside a guest, it was observed that a vdisk would go offline intermittently even when at least one I/O domain which provided a path to that disk was still up and running.

This issue is still under investigation. Symantec recommends that you install Solaris 10 Update 7 that contains the fix for Sun bug id 6742587 (vds can ACK a request twice). This fix possibly resolves this issue.

Deadlock between DMP kernel and VDC driver during volume creation

This was observed by Sun Microsystems during their interoperability testing of 5.0 MP3. The deadlock happens when creating a mirrored VxVM volume, while there is ongoing I/O to a UFS file system on another disk which is not under VxVM. This issue has been observed typically in a large configuration such as 14 virtual CPUs and around 10Gig of memory configured for the guest logical domain running VxVM.

Relevant SUN bug id: 6744348

This known issue has been fixed and verified in the 5.0 MP3 RP1 release.

For the latest information on updates, patches, documentation, and known issues regarding this 5.0 MP3 RP1 release, see the following TechNote on the Symantec Technical Support website:

For Solaris SPARC:

<http://entsupport.symantec.com/docs/281987>

For Solaris x64:

<http://entsupport.symantec.com/docs/286955>

Split Storage Foundation stack known issues

The following are new known issues in this release of Veritas Storage Foundation and High Availability Solutions Support for Solaris Logical Domains.

Caching of data writes on the backend volume in the service domain

This was observed by a customer during their evaluation of LDomS with Storage Foundation. This issue occurs because data written to the virtual disk might get cached into the service domain before it is effectively written to the virtual disk backend. This can cause potential data loss if the service domain crashes while the data is still cached.

SUN bug id is: 6684721 (file backed virtual I/O should be synchronous)

This SUN bug is fixed in Sun patch 139562-02 that has been obsoleted by 138888-07.

See [“Solaris patch requirements”](#) on page 70.

A volume can become inaccessible from the guest in the event of control domain reboot

All access to such a volume hangs if the primary domain is rebooted. This is due to the vdisk corresponding to the volume does not come back online after the control domain reboots.

This issue has been identified and fixed under Sun bug id: 6795836 (vd_setup_vd() should handle errors from vd_identify_dev() better)

This SUN bug is fixed in Sun patch 141777-01.

Veritas Cluster Server support for using multiple nodes in a Logical Domain environment

This chapter includes the following topics:

- [Clustering using Cluster Volume Manager](#)
- [Installing Storage Foundation on multiple nodes in a Logical Domain](#)
- [Cluster Volume Manager in the control domain for providing high availability](#)

Clustering using Cluster Volume Manager

The Veritas Volume Manager cluster functionality (CVM) makes logical volumes and raw device applications accessible throughout a cluster.

In the split Storage Foundation model, CVM is placed in the control domain and VxFS is placed in the guest domain. In this model, CVM provides high availability and shared storage visibility at the control domain level across multiple physical nodes in the cluster.

See [“Cluster Volume Manager in the control domain for providing high availability”](#) on page 92.

In the guest-based Storage Foundation stack model, CVM is placed in the guest domain, providing high availability and shared storage visibility at the guest domain level across multiple guest domains that act as the nodes in the cluster.

Installing Storage Foundation on multiple nodes in a Logical Domain

To install Storage Foundation on multiple nodes in a Solaris Logical Domains environment, you must complete the following operations, which are the same as on a single node:

- See [“Installing and configuring the Logical Domain software and domains”](#) on page 73.
- See [“Installing Storage Foundation in the control domain or guest”](#) on page 73.
- See [“Installing Veritas File System in the guest domain using pkgadd”](#) on page 74.
- See [“Verifying the configuration”](#) on page 74.

Reconfiguring the clustering agents for Cluster Volume Manager

This section applies to only the Split Storage Foundation model.

For a Storage Foundation CVM, the following additional configuration steps are necessary:

- See [“Removing the vxfsckd resource”](#) on page 90.
- See [“Creating CVMVolDg in a group”](#) on page 91.

Removing the vxfsckd resource

After configuring Storage Foundation and CVM, remove the vxfsckd resource.

To remove the vxfsckd resource

- 1 Make the configuration writeable:

```
# haconf -makerw
```

- 2 Delete the resource:

```
# hares -delete vxfsckd
```

- 3 Make the configuration read-only:

```
# haconf -dump -makero
```

4 Stop the resources:

```
# hastop -all
```

5 Restart the resources.

```
# hastart
```

Run the `hastart` command on all nodes in the cluster.

Creating CVMVolDg in a group

The following procedure creates CVMVolDg in a given group.

To create CVMVolDg

1 Make the configuration writeable:

```
# haconf -makerw
```

2 Add the CVMVolDg resource:

```
# hares -add name_of_resource CVMVolDg name_of_group
```

3 Add a diskgroup name to the resource:

```
# hares -modify name_of_resource CVMDiskGroup sdg1
```

4 Make the attribute local to the system:

```
# hares -local name_of_resource CVMActivation
```

5 Add the attribute to the resource.

```
# hares -modify name_of_resource CVMActivation  
\  
activation_value -sys nodename
```

Repeated this step on each of the nodes.

6 If you want to monitor volumes, enter the following command:

```
# hares -modify name_of_resource CVMVolume -add \  
name_of_volume
```

In a database environment, Symantec recommends you monitor volumes.

- 7 Modify the resource so that a failure of this resource does not bring down the entire group:

```
# hares -modify name_of_resource Critical 0
```

- 8 Enable the resources:

```
# hares -modify cvmvoldg1 Enabled 1
```

- 9 Make the configuration read-only:

```
# hares -modify cvmvoldg1 Enabled 1
```

- 10 Verify the configuration:

```
# hacf -verify
```

This should put the resource in the `main.cf` file.


Cluster Volume Manager in the control domain for providing high availability

The main advantage of clusters is protection against hardware failure. Should the primary node fail or otherwise become unavailable, applications can continue to run by transferring their execution to standby nodes in the cluster.

CVM can be deployed in the control domains of multiple physical hosts running LDoms, providing high availability of the control domain.

[Figure 5-1](#) illustrates a CVM configuration.

Figure 5-1 CVM configuration in an Solaris Logical Domain environment



If a control domain encounters a hardware or software failure causing the domain to shut down, all applications running in the guest LDomS on that host are also affected. These applications can be failed over and restarted inside guests running on another active node of the cluster.

Caution: As such applications running in the guests may resume or time out based on the individual application settings. The user must decide if the application must be restarted on another guest on the failed-over control domain. There is a potential data corruption scenario if the underlying shared volumes get accessed from both of the guests simultaneously.

Shared volumes and their snapshots can be used as a backing store for guest LDomS.

Note: The ability to take online snapshots is currently inhibited because the file system in the guest cannot coordinate with the VxVM drivers in the control domain.

Make sure that the volume whose snapshot is being taken is closed before the snapshot is taken.

The following example procedure shows how snapshots of shared volumes are administered in such an environment. In the example, `datavol1` is a shared volume being used by guest LDom `ldom1` and `c0d2s2` is the front end for this volume visible from `ldom1`.

To take a snapshot of `datavol1`

- 1 Unmount any VxFS file systems that exist on `c0d1s0`.
- 2 Stop and unbind `ldom1`:

```
primary# ldm stop ldom1
```

```
primary# ldm unbind ldom1
```

This ensures that all the file system metadata is flushed down to the backend volume, `datavol1`.

- 3 Create a snapshot of `datavol1`.

See the *Veritas Volume Manager Administrator's Guide* for information on creating and managing third-mirror break-off snapshots.

- 4 Once the snapshot operation is complete, rebind and restart `ldom1`.

```
primary# ldm bind ldom1
```

```
primary# ldm start ldom1
```

- 5 Once `ldom1` boots, remount the VxFS file system on `c0d1s0`.

Veritas Cluster Server: Configuring Logical Domains for high availability

This chapter includes the following topics:

- [About Veritas Cluster Server in a Logical Domain environment](#)
- [About Veritas Storage Foundation Cluster File System in a Logical Domain environment](#)
- [About Veritas Cluster Server configuration models in a Logical Domain environment](#)
- [Configuring Veritas Cluster Server to fail over a Logical Domain on a failure](#)
- [Configuring Veritas Cluster Server to fail over an application on a failure](#)
- [LDom migration in a VCS environment](#)

About Veritas Cluster Server in a Logical Domain environment

Veritas Cluster Server (VCS) 5.1 (or VCS 5.0 MP3 RP1 and later releases) provides high availability (HA) for a Sun Microsystems Logical Domain (LDom). You can configure VCS to monitor the complete LDom, its components, and the applications that run in LDom, or to monitor only the applications that run in LDom.

See [“About Veritas Cluster Server configuration models in a Logical Domain environment”](#) on page 99.

Table 6-1 lists the failure scenarios and the VCS failover options based on which you can plan your VCS configuration in an LDom environment.

Table 6-1 Veritas Cluster Server failover options for Logical Domain failure

Failure scenario	VCS failover	Typical VCS configuration
LDoms, their storage, or switches fail	VCS fails over the LDom from one node to the LDom on another node	VCS is installed in the control domain of each node. See “Veritas Cluster Server setup to fail over a Logical Domain on a failure” on page 100.
LDoms, their storage, or switches fail Or Applications that run in LDoms fail	VCS fails over the LDom from one node to the LDom on another node. The application starts on the same LDom after the LDom failover.	VCS is installed in the control domain of each node, and single node VCS is installed on each guest domain. See “Veritas Cluster Server setup to fail over a Logical Domain on a failure” on page 100.
Applications that run in LDoms fail Or LDom where the application is running fails	VCS fails over the application from one LDom to another.	VCS is installed in the guest domain of each node. See “Veritas Cluster Server setup to fail over an application on a failure” on page 103.

Veritas Cluster Server prerequisites

This document assumes a working knowledge of VCS.

Review the prerequisites in the following documents to help ensure a smooth VCS installation:

- *Veritas Cluster Server Release Notes*
Find this in the `cluster_server/docs` directory of the product disc.
- *Veritas Cluster Server Installation Guide*
Find this in the `cluster_server/docs` directory of the product disc.

Unless otherwise noted, all references to other documents refer to the Veritas Cluster Server documents version 5.1 for Solaris.

Veritas Cluster Server requirements

For installation requirements:

See [“System requirements”](#) on page 69.

For the configuration model where VCS is installed on the control domain:

- VCS requires shared storage that is visible across all the nodes in the cluster.
- Configure each LDom on a node.
- The LDom’s boot device and application data must reside on shared storage.

For the configuration model where VCS is installed on the guest domain:

- VCS requires the application data to reside on the shared storage.
- Each node can have more than one LDom.
- Each LDom can have its own boot device.

Veritas Cluster Server limitations

Depending on the configuration model, the following limitations apply to using VCS in an LDom environment.

Limitations for VCS configuration in the control domain:

- VCS does not support LDom migration.
See [“LDom migration in a VCS environment”](#) on page 112.
- When VCS is installed in control domains and managing guest domains, VCS does not support the use of alternate I/O domains for the guest domains. The use of alternate I/O domains may result in the loss of high availability.
- This release of VCS does not support attaching raw physical disks or slices to LDom. Such configurations may cause data corruption either during an LDom failover or if you try to manually bring up LDom on different systems. For details on supported storage configurations: See [“Configuration scenarios”](#) on page 105.
- Each LDom configured under VCS must have at least two VCPUs. With one VCPU, the control domain always registers 100% CPU utilization for the LDom. This is an LDom software issue.

Limitation for VCS configuration in the guest domain:

- If you want to configure I/O fencing in guest domain, then do not export physical devices to more than one guest domain on the same physical node. Otherwise, I/O fencing fences off the device whenever one of the guest domain dies. This situation causes the other guest domains also to lose access to the device.

Symantec recommends you to disable I/O fencing if you exported the same physical device to multiple guest domains.

Veritas Cluster Server known issues

The following section describes known issues using VCS in an LDom environment.

Shutting down the control domain may cause the guest domain to crash (1631762)

Set up	Two Sun SPARC Enterprise T5240 server physical boxes, each with a control domain and a guest domain. The guest domains in each of the physical boxes form a two node cluster. The nodes are named node 0 and node 1 in the following text.
Symptom	Gracefully shutting down the control domain of node 0 causes the guest domain of node 0 to crash.
Analysis	<p>Even though the guest domain can continue to function when the control domain is shut down, the heartbeats between node 0 and node 1 are lost as the control domain shuts down. As a result, the cluster forms two separate sub-clusters without the sub-clusters being able to see each others' heartbeats.</p> <p>I/O fencing resolves the split brain situation and determines that only one sub-cluster will continue to function while the other sub-cluster should panic. Therefore, the panic of node 0 is expected behavior.</p>

Resolution: None; this is expected behavior. However, Symantec recommends keeping the control domain highly available for the proper function of the SFCFS and SFRAC stack in the guest domains.

If you have set up a virtual private LLT heartbeats between the two guests (node 0 and node1), the guest will not crash.

About Veritas Storage Foundation Cluster File System in a Logical Domain environment

Veritas Storage Foundation Cluster File System (SFCFS) 5.1 allows clustered servers to mount and use a file system simultaneously as if all applications using the file system were running on the same server for a Sun Microsystems Logical Domain (LDom).

Veritas Storage Foundation Cluster File System limitations

Depending on the configuration model, the following limitations apply to using SFCFS in an LDom environment.

Limitations for SFCFS configuration in the guest domain:

- There is no support for replicating a shared disk group using VVR, when one or more guest domains share the disk group.
- If you want to configure I/O fencing in guest domain, then do not export physical devices to more than one guest domain on the same physical node. Otherwise, I/O fencing fences off the device whenever one of the guest domain dies. This situation causes the other guest domains also to lose access to the device.

Symantec recommends you to disable I/O fencing if you exported the same physical device to multiple guest domains.

About Veritas Cluster Server configuration models in a Logical Domain environment

When you configure VCS in an LDom environment, you need some specific information about the LDom, network, and the storage devices that the LDom requires to run.

You need to know the following information about your LDom:

- The LDom's name
- The names of the primary network interfaces for each node
- The virtual switch that the LDom uses
- The name and type of storage that the LDom uses

VCS configuration depends on whether you want VCS to fail over the LDom or the application on a failure:

- [Veritas Cluster Server setup to fail over a Logical Domain on a failure](#)
- [Veritas Cluster Server setup to fail over an application on a failure](#)

See [Table 6-1](#) for more information.

Veritas Cluster Server setup to fail over a Logical Domain on a failure

You can configure VCS to monitor the LDom, or to monitor both the LDom and the application in the guest domain.

- See [“Veritas Cluster Server installed in the control domain to monitor the Logical Domains”](#) on page 100.
- See [“Veritas Cluster Server installed in the control domain to monitor the applications in the guest domain”](#) on page 101.

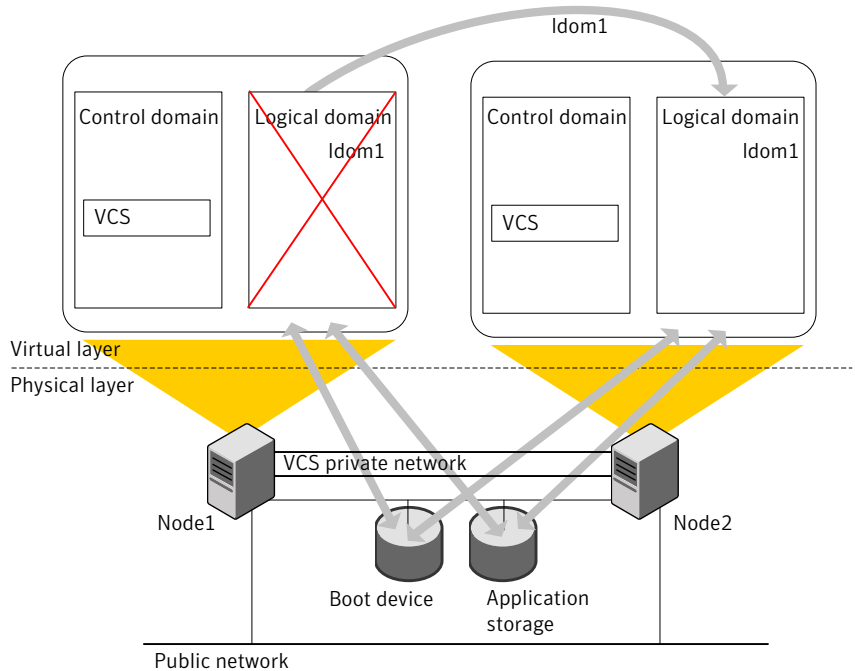
When you configure VCS in your LDom environment, VCS monitors the health of the LDom, and its supporting components. The LDom agent monitors the logical domain. If the agent detects that the LDom resource has failed, VCS moves the service group that contains the LDom resource to another node in the cluster.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for information on the LDom agent and other bundled agents.

Veritas Cluster Server installed in the control domain to monitor the Logical Domains

[Figure 6-1](#) illustrates a typical setup where VCS installed in the control domain provides high availability to LDom and its resources.

Figure 6-1 Typical setup for Logical Domain high availability with VCS control domains



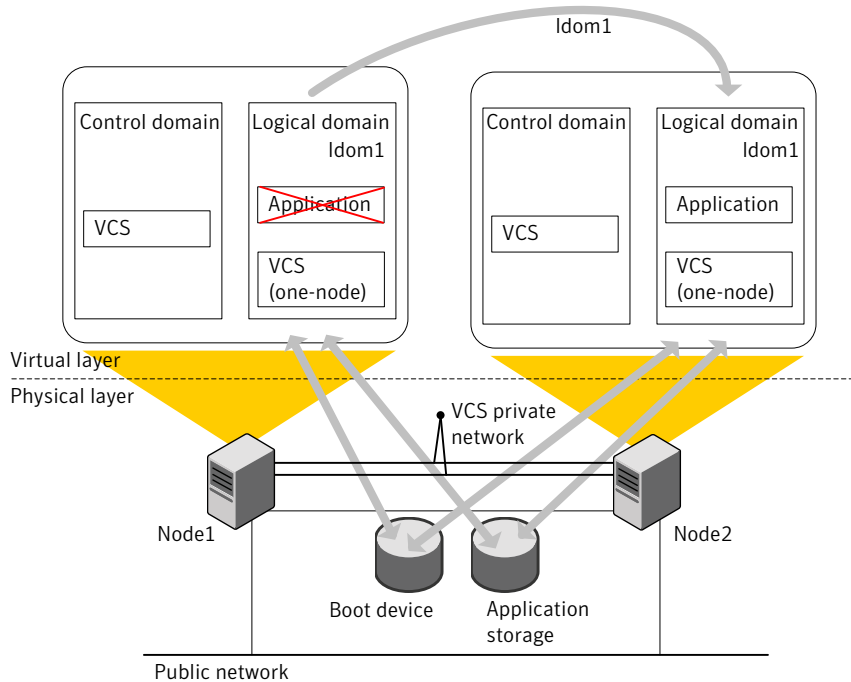
A typical two-node VCS configuration for LDom high availability has the following software and hardware infrastructure:

- Sun LDom software is installed on each system Node1 and Node2.
- Shared storage is attached to each system.
- An LDom ldom1 exists on both the nodes with a shared boot device.
- VCS is installed in the control domains of each node.

Veritas Cluster Server installed in the control domain to monitor the applications in the guest domain

Figure 6-2 illustrates a typical setup where VCS installed in the control domain provides high availability to applications that run in the guest domains.

Figure 6-2 Typical setup for application high availability with VCS in control domains



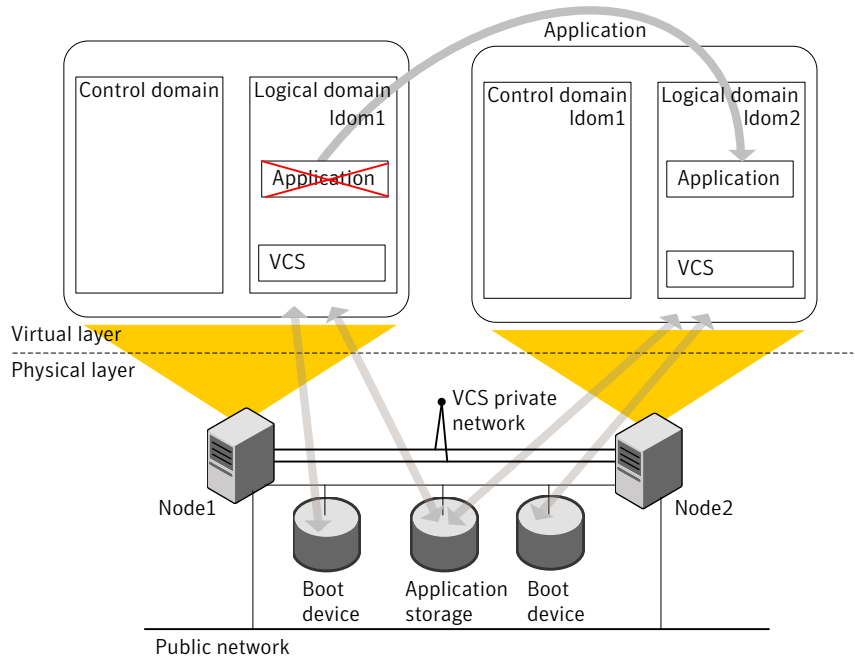
A typical two-node VCS configuration that fails over the LDomS to keep the applications that run in LDomS highly available has the following infrastructure:

- Sun LDom software is installed on each system Node1 and Node2.
- Shared storage is attached to each system.
- An LDom ldom1 with same configuration details exists on both the nodes with a shared boot device.
- Each LDom has an operating system installed.
- VCS is installed in the control domains of each node.
- Each guest domain has single-node VCS installed. VCS kernel components are not required.
- VCS service group exists for the application that VCS must manage.
- VCS RemoteGroup service group with an online global firm dependency to the LDom service group is created to monitor the Application service group.

Veritas Cluster Server setup to fail over an application on a failure

Figure 6-3 illustrates a typical VCS setup to provide high availability for applications that run in guest domains.

Figure 6-3 Typical setup for applications high availability with Veritas Cluster Server in control domains



A typical two-node VCS configuration that fails over the applications to keep the applications that run in LDomS highly available has the following infrastructure:

A typical two-node configuration where VCS keeps the applications that run in LDomS highly available has the following software and hardware infrastructure:

- Sun LDom software is installed on each system Node1 and Node2.
- Shared storage is attached to each system.
- LDomS are created on both the nodes that may have local boot devices.
- VCS is installed in the guest domains of each node.

Configuring Veritas Cluster Server to fail over a Logical Domain on a failure

You can configure VCS to keep the LDom highly available. In addition to monitoring the LDom, you can also configure VCS to monitor the applications that run in the LDom.

You must perform additional steps for VCS in the control domain to manage applications in the guest domains. After you install VCS in the control domain, you must create separate service groups for the RemoteGroup resource and LDom resource with online global firm dependency.

Note: If you create the RemoteGroup resource as part of the LDom service group, then the RemoteGroup resource state remains as UNKNOWN if the LDom is down. So, VCS does not probe the service group and cannot bring the LDom online. The online global firm dependency between the service groups allows VCS to fail over a faulted child LDom service group independent of the state of the parent RemoteGroup service group.

Perform the following tasks to configure VCS to fail over an LDom on an LDom failure:

- Review the configuration scenarios
See [“Configuration scenarios”](#) on page 105.
- Configure logical domains
See [“Configuring logical domain”](#) on page 107.
- Install VCS on control domain
See [“Installing Veritas Cluster Server inside the control domain”](#) on page 108.
- Create VCS service group for LDom
See [“Creating the Veritas Cluster Server service groups for Logical Domain”](#) on page 108.

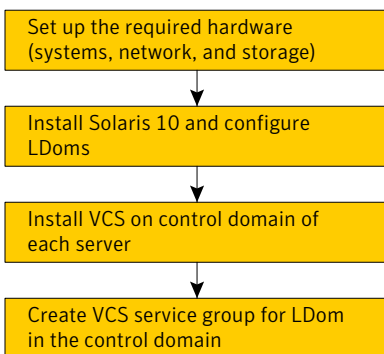
Perform the following additional tasks to configure VCS to fail over an LDom on an application failure:

- Install single-node VCS on guest domain
See [“Installing single-node Veritas Cluster Server inside the guest domain”](#) on page 109.
- Configure VCS in control domain to monitor the application in guest domain
See [“Configuring Veritas Cluster Server to monitor the application in the guest domain”](#) on page 109.

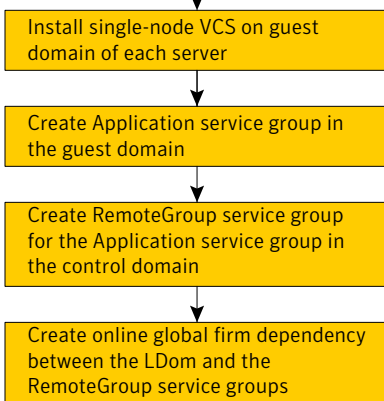
Figure 6-4 depicts the workflow to configure VCS to manage the failure of an LDom or the failure of an application that runs in an LDom.

Figure 6-4 Workflow to configure VCS to fail over a Logical Domain on a failure

For VCS to monitor the LDom:




For VCS to monitor the application in guest domain:



Configuration scenarios

Figure 6-5 shows the basic dependencies for an LDom resource.

Figure 6-5 A Logical Domain resource depends on storage and network resources



Network configuration

Use the NIC agent to monitor the primary network interface, whether it is virtual or physical. Use the interface that appears using the `ifconfig` command.

[Figure 6-6](#) is an example of an LDom service group. The LDom resource requires both network (NIC) and storage (Volume and DiskGroup) resources.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the NIC agent.

Storage configurations

Depending on your storage configuration, use a combination of the Volume, DiskGroup, and Mount agents to monitor storage for LDom.


Note: VCS in a control domain supports only volumes or flat files in volumes that are managed by VxVM for LDom storage.

Veritas Volume Manager exposed volumes

Veritas Volume Manager (VxVM) exposed volumes is the recommended storage solution for VCS in a control domain. Use the Volume and DiskGroup agents to monitor a VxVM volume. VCS with VxVM provides superior protection for your highly available applications.

[Figure 6-6](#) shows an LDom resource that depends on a Volume and DiskGroup resource.

Figure 6-6 The Logical Domain resource can depend on many resources, or just the NIC, Volume, and DiskGroup resources depending on the environment




For more information about the Volume and DiskGroup agents, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

Image files

Use the Mount, Volume, and DiskGroup agents to monitor an image file.

[Figure 6-7](#) shows how the Mount agent works with different storage resources.

Figure 6-7 The mount resource in conjunction with different storage resources



See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the Mount agent.

Configuring logical domain

You must perform the following steps to configure logical domain.

To configure logical domain

- 1 Make sure that the network and storage setup meet the VCS requirements. See “[Veritas Cluster Server requirements](#)” on page 97.
- 2 Make sure that you have installed the required Solaris operating system in the logical domain.
- 3 Create an LDom (ldom1) on each system with an identical configuration and boot device.

Installing Veritas Cluster Server inside the control domain

You must install VCS in the control domain of each system.

To install Veritas Cluster Server in a Logical Domain environment

- ◆ Install and configure VCS in the primary control domain of each system.

The process for installing VCS in the control domain is very similar to the regular installation of VCS. However, you must specify the name of the control domain for the name of the host where you want to install VCS.

See the *Veritas Cluster Server Installation Guide* for VCS installation and configuration instructions.

Creating the Veritas Cluster Server service groups for Logical Domain

You can also create and manage service groups using the Veritas Cluster Server Management Server, the Cluster Manager (Java Console), or through the command line.

See the *Veritas Cluster Server User's Guide* for complete information about using and managing service groups, either through CLI or GUI.

Configuring Veritas Cluster Server for application monitoring

You must perform the following procedures to configure domain to monitor the application in the guest domain:

- See “[Installing single-node Veritas Cluster Server inside the guest domain](#)” on page 109.
- See “[Configuring Veritas Cluster Server to monitor the application in the guest domain](#)” on page 109.

Installing single-node Veritas Cluster Server inside the guest domain

Perform the following steps to install one-node Veritas Cluster Server (VCS) inside the guest domain:

To install and configure one-node Veritas Cluster Server inside the logical domain

- 1 Install one-node VCS (no kernel components required) in the guest domain.
See the *Veritas Cluster Server Installation Guide* to perform a single-node VCS installation in the logical domains.
- 2 Start the VCS engine.

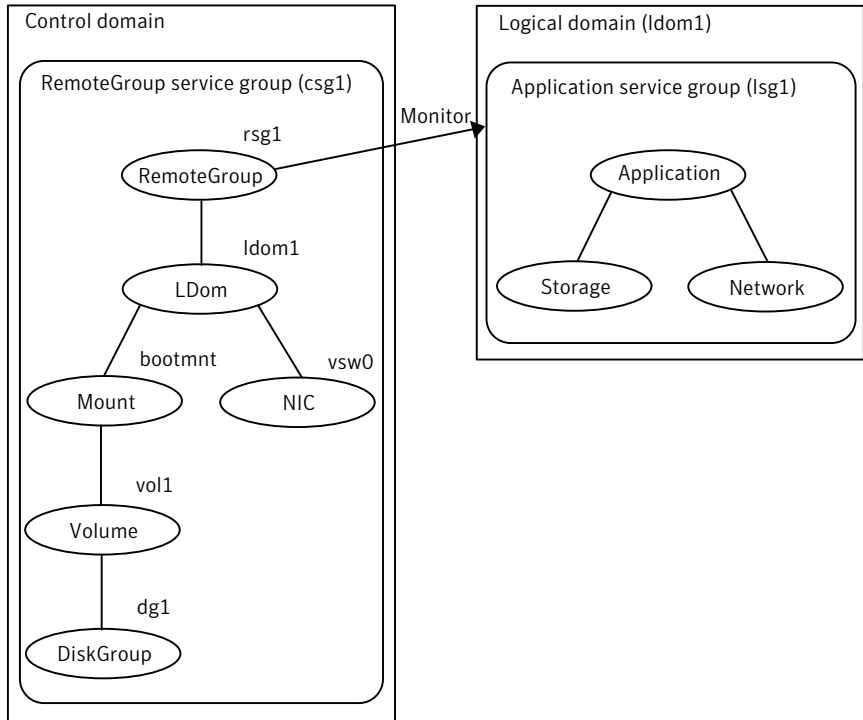
Configuring Veritas Cluster Server to monitor the application in the guest domain

Perform the following steps to configure Veritas Cluster Server (VCS) in the control domain to monitor the application in the guest domain.

To configure Veritas Cluster Server to monitor the application in the guest domain

- 1 Configure a VCS service group (lsg1) for the application.
The ManualOps attribute of the service group must remain set to true, the default value.
- 2 Add a VCS user (lsg1 -admin) with the minimum privilege as the group operator of the VCS service group (lsg1).
- 3 Configure a RemoteGroup service group (rsg1) in the control domain to monitor the VCS service group (lsg1) that was configured in 1.
- 4 Set the value of the following RemoteGroup resource attributes as follows:
 - RestartLimit attribute of the resource or type to 1 or higher
 - OfflineWaitLimit attribute of the resource or type to 1
 - ToleranceLimit attribute of the resource or type to 1
- 5 Create the dependencies between the groups and resources as shown in [Figure 6-8](#).

Figure 6-8 Group and resource dependency diagram



Note: RemoteGroup and Application service groups are required only when you want to configure VCS to monitor the application in the guest domain.

RemoteGroup resource definition

The resource definition for the RemoteGroup resource is as follows:

```
RemoteGroup rsg1 (  
    GroupName = lsg1  
    IpAddress = <IP address of ldom1>  
    ControlMode = OnOff  
    Username = lsg1-admin  
    Password = <lsg1-admin's password>  
)
```

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on the RemoteGroup agent.

Verifying a Logical Domain service group failover

Verify the configuration in different situations.

Using a switch command

Switch the LDom to another node in the cluster to make sure the service group fails over. If all the resources are properly configured, the service group shuts down on the first node and comes up on the second node.

Other verification scenarios

In all of these verification scenarios, you are stopping or moving an LDom, or stopping a resource for that LDom. VCS should detect the failure, or the movement, and either failover the effected LDom or take no action.

The following list presents some quick testing scenarios;

- From outside of VCS control, stop the LDom. VCS should fail the LDom over to the other node.
- Boot the LDom through VCS by entering a `hagrp -online` command. move the LDom to another node by shutting it down through VCS on the node where the LDom is running. boot the LDom outside of VCS control on the other node- the service group comes online on that node.

Configuring Veritas Cluster Server to fail over an application on a failure

You must install and configure Veritas Cluster Server (VCS) in the guest domains of each system to enable VCS to manage applications in the guest domains.

To configure Veritas Cluster Server to manage applications in the guest domains

- 1 Install and configure VCS in the guest domains of each system.
See the *Veritas Cluster Server Installation Guide* for installation and configuration instructions.
- 2 Create two virtual NICs using private virtual switches for private interconnects.
You can configure virtual switches with no physical network interfaces if you want the failover across LDomS in the same control domain.
- 3 Configure VCS service group for the application that you want to monitor.
 - Configure Mount and Disk resources to monitor the storage.

- Configure NIC resources to monitor the network.
- Configure application resources using the application-specific agent.

See the *Veritas Cluster Server User's Guide* for more details on configuring applications and resources in VCS.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for details on the storage and networking bundled agents.

LDom migration in a VCS environment

VCS supports a warm LDom migration when VCS is installed in a logical domain. A warm migration of an LDom can take around some time. While the LDom is migrating, the LDom's software freezes the LDom and the application that runs inside the LDom.

You do not have to start and stop LLT and GAB. In a warm migration, LLT and GAB restart themselves gracefully.

To perform a warm migration for an LDom when VCS is installed in logical domains

- 1 Stop VCS engine using `hastop -local -force` on the LDom system that you plan to migrate.

Perform this step so that GAB does not have to kill the VCS engine process when the migration is complete. GAB wants all clients to reconfigure and restart when the configuration is not in sync with other members in the cluster.

- 2 If CVM is configured inside the LDom, perform this step.

Set LLT `peerinact` to 0 on all nodes in the cluster so that while the LDom is in migrating, the system is not thrown out of the cluster by the other members in the cluster. If the CVM stack is unconfigured, the applications can stop.

- 3 If fencing is configured in single instance mode inside the LDom, perform this step.

Unconfigure and unload the `vxfen` module in the LDom.

Perform this step so that GAB does not panic the node when the LDom migration is complete.

- 4 Migrate the logical domain from the control domain using the `ldm` interface. Wait for migration to complete.

- 5 Perform this step if you performed step 3

Load and configure `vxfen` module in the LDom.

6 Perform this step if you performed step 2

Reset LLT peerinact to the original value on all nodes in the cluster.

7 Use the `hastart` command to start VCS engine inside the LDom.

VCS does not support LDom migration when VCS is installed in the control domain. If you plan to use the LDom migration feature in a VCS environment, however, follow these guidelines.

To migrate an LDOM when VCS is installed in the control domain**1** Ensure that you have created an XML configuration file and have configured the `CfgFile` attribute of the LDom agent.

Note: If you do not complete this step, VCS reports the status of the resource on the source node as `UNKNOWN` after the migration is complete.

2 Before initiating migration, freeze the service group that contains LDom resources.**3** Migrate the logical domain.**4** After the migration is complete, VCS takes about five minutes to detect the online status of the service group. You may probe the resources on the target node to verify status.**5** Unfreeze the service group.

