Veritas Storage Foundation™ and High Availability Solutions Release Notes

Windows Server 2008 (x64), Windows Server 2008 R2 (x64)

6.0.1

March 2013

Symantec

Veritas Storage Foundation™ and High Availability Solutions Release Notes

Product version: 6.0.1

Document version: 6.0.1 Rev2

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

About Symantec Connect

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

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

Release Notes

This document includes the following topics:

Introduction

This document provides the release details of the following two products:

- Veritas Storage Foundation™ 6.0.1 for Windows (SFW)
- Veritas Storage Foundation™ HA 6.0.1 for Windows (SFW HA)

The information in the Release Notes supersedes the information provided in the product documents. You can download the latest version of this document from the Symantec SORT website.

https://sort.symantec.com

For the latest information on updates, patches, and software issues regarding this release, see the following Late Breaking News (LBN):

http://www.symantec.com/docs/TECH161556

Requirements

For information about the operating system, hardware, and other general requirements of Storage Foundation and High Availability solutions for Windows, see the *Veritas Storage Foundation™ and High Availability Solutions Installation and Upgrade Guide*.

For the latest information on supported hardware, see the Hardware Compatibility List (HCL) at:

http://www.symantec.com/docs/TECH152806

For the latest information on supported software, see the Software Compatibility List (SCL) at:

http://www.symantec.com/docs/TECH201485

About Symantec Operations Readiness Tools

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

SORT can help you do the following:

Prepare for your next installation or upgrade	■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture. ■ Analyze systems to determine if they are ready to install or upgrade Symantec products. ■ Download the latest patches, documentation, and high availability agents from a central repository. ■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems.

Manage risks	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository. ■ Identify and mitigate system and environmental risks. ■ Display descriptions and solutions for hundreds of Symantec error codes.
Improve efficiency	■ Find and download patches based on product version and platform. ■ List installed Symantec products and license keys. ■ Tune and optimize your environment.

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

To access SORT, go to:

https://sort.symantec.com

New features and changes for SFW and SFW HA

This section describes the new features and changes introduced in Veritas Storage Foundation for Windows (SFW) and Veritas Storage Foundation HA for Windows (SFW HA) 6.0.2.

Enhancements to the product installer

The following changes have been introduced in the product installer in this release:

- Product Updates
  In this release, a new feature is added that allows you to check for available product updates that can be downloaded and installed on the system. The product installer searches for the available product updates on the SORT website.
- Product Improvement Program
  This release has a new feature called Product Improvement Program. The focus of this feature is to provide data collection and upload capabilities for the product installations.
  When you install the product, you see an additional checkbox, **Participate in the Symantec Product Improvement Program by submitting system and usage information anonymously**, that asks you to allow the upload of installation data to Symantec. The collected information helps identify how

customers deploy and use the product. You also have the option to opt out of enabling this feature. If you choose to enable this feature, the collection runs during the installation only.

Veritas Storage Foundation

The following are the new features and changes for Veritas Storage Foundation in this release.

`Chkdsk` support for the VMDg resource in FOC

In a Microsoft Failover Cluster (FOC) configuration, you can now configure the Volume Manager Disk Group (VMDg) resource to clean the file system of all the volumes under it while the resource is being brought online. During this, it checks for dirty bit of the file system and, if it is set, then it runs the `chkdsk.exe` utility to fix errors in the file system.

For more details, refer to the *Veritas Storage Foundation Administrator's Guide*.

Performance tunables for storage migration

Performance tunables are provided to optimize the time taken for the storage migration operation.

The following performance tunables are introduced for storage migration:

- I/O size
- Task threads

Use these tunables to optimize the time required for the storage migration operation and improve storage migration performance. These tunables provide the flexibility to trade between storage migration performance and application performance. Usage of these tunables is limited to storage migration operations only.

The performance tunables are set at default values. However, the user can define these values to optimize the storage migration performance.

For more details, refer to the *Veritas Storage Foundation Administrator's Guide*.

Enhancements to storage migration

The SFW storage migration feature is enhanced to include support for the following:

- Change layout of target volumes at run time: The Storage Migration Wizard now provides options to change the volume layout of the target volumes selected for migration.
- Perform storage migration at an array enclosure level: The Storage Migration Wizard now allows you to perform migration for data on disks that belong to an array enclosure.
  The Storage Migration Wizard provides separate workflows that allow you to choose migration for volumes that are part of a SFW dynamic disk group, or volumes belonging to a Microsoft Hyper-V virtual machine, or disks from an array enclosure.

For more details, refer to the *Veritas Storage Foundation Administrator's Guide*.

Safeguarding the expand volume operation in SFW against limitations of NTFS

To prevent users from expanding volumes beyond the limitations of NTFS for NTFS cluster size, the following safeguard is implemented:

- SFW does a check for the NTFS cluster size of the volume while expanding it.
- The operation fails, if the user tries to expand the volume beyond a limit.
- The limit depends upon the cluster size of the specified volume during its creation.

When SFW receives an **expand volume** request, it validates the requested new size against the **MAX SIZE** supported for that volume.

For more details, refer to the *Veritas Storage Foundation Administrator's Guide*.

Working with disks that support thin provisioning

The timeout value for the reclaim I/O operations is set to 30 seconds by default. However depending on the configuration, the storage reclaim operation may sometimes take a significant amount of time to complete. To avoid the reclaim I/O failures due to timeout errors, you can change the reclaim I/O timeout value by adjusting the following registry tunable parameter:

`HKLM\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager\MaxReclaimIOTimeOut`

For more details, refer to the *Veritas Storage Foundation Administrator's Guide*.

Disaster Recovery (DR) fire drill enhancements

The following changes were made to the DR fire frill feature:

- The DR fire drill feature has been updated to enable complete support from the Veritas Operations Manager (VOM) console. For more information about

configuring and performing fire drills in VOM, see the *Veritas Operations Manager Add-on for Veritas Cluster Server Administration User's Guide*.

- A new agent named VVRSnap has been introduced to perform certain tasks in the VVR environment. The VVRSnap agent makes storage available for running fire drill in a disaster recovery environment. For more information, see the *Veritas Cluster Server Bundled Agents Reference Guide*.
- When preparing the fire drill configuration, the Fire Drill Wizard now performs the following actions:
  - Creates a fire drill service group with the VVRSnap resource in the VVR environment.
  - Sets a dependency between the fire drill service group and the application service group.

  For more information, see the SFW HA Solutions Guides.

For more details, refer to the *Veritas Storage Foundation Administrator's Guide*.

SSD Awareness in SFW operations

Solid-state devices, also known as flash drives or solid-state disks, used in storage are recognized by SFW. Using the Volume Manager, you can classify these devices.

Volume Manager space allocation operations are also media type aware and support both HDD and SSD devices.

An SSD plex is set as a preferred plex to augment the Read I/O performance of SFW.

Refer to the *SFW Administrator's Guide* for more information about automatic discovery, manual classification, SSD awareness and setting SSD as a preferred plex.

Veritas Cluster Server

The following new features and changes are introduced in Veritas Cluster Server in this release:

Added support to configure monitoring for Generic application and SQL Server 2008 in VMware virtual environment

For configuring application monitoring for generic applications and SQL Server 2008, in VMware virtual environment, VCS now provides a new application configuration wizard.

Use this wizard to configure application monitoring in a VMware environment. You can configure application monitoring either in a start/stop mode on a single system or a failover mode on multiple systems.

The wizard first configures a VCS cluster and then configures application monitoring.

vSphere integrated options to install VCS, monitor application status and manage licenses

In a VMware virtual environment, the Symantec High Availability solution now provides the vSphere integrated options to monitor the application status.

To configure application monitoring in VMware environment, install the Symantec High Availability Console. As part of the Console installation, the installer registers the Symantec High Availability plugin for VMware vCenter Server. This plugin enables integration of Symantec High Availability with VMware vSphere Client and adds the following options to the VMware vSphere Client:

- Menu to install the guest components
  A right-click menu is available to install the guest components when you select a cluster, a datacenter, or a virtual machine from the VMware vCenter Server inventory.
  The Symantec High Availability home page is also available as an vSphere Client extension under its Solutions and Applications pane.
  Use these options to install VCS as a part of the guest components installation.
- Symantec High Availability tab
  The Symantec High Availability tab is visible when you select a virtual machine from the VMware vCenter Server inventory. Use this tab to configure and control application availability on virtual machines that are managed from the VMware vCenter Server. You can perform these operations per virtual machine.
- Symantec High Availability dashboard
  The Symantec High Availability dashboard is visible when you select a cluster or a datacenter from theVMware vCenter Server inventory. Use the dashboard to administer the configured applications on virtual machines in a VMware datacenter. You can perform these operations at a VMware cluster or datacenter-level.
- Menu to manage licenses
  The Symantec High Availability home page is added as an vSphere Client extension under its Solutions and Applications pane.

Use this option to update the licenses by adding or removing specific license keys on the systems where you have installed the Symantec High Availability guest components.

Changes to MountV agent to secure mount point folders during the MountV offline operation

The MountV agent is enhanced to address the issue where the mount point folders are accessible after MountV takes the volumes offline. MountV agent can mount a volume as an NTFS folder. If this volume is unmounted, the NTFS folder remains accessible for writes by users and processes, which can create new files and folders in the NTFS folder. This can cause the MountV online operation to fail on the failover target node.

A new attribute, BlockMountPointAccess, is added to address this issue. This attribute defines whether the agent blocks access to the NTFS folder that is used as a folder mount point after the mount point is unmounted.

For example, if C:\temp is used as a folder mount for a volume, then after the mount point is unmounted, the agent blocks access to the folder C:\temp. The value true indicates that the folder is not accessible.

The default value false indicates that the folder is accessible.

In DR configuration, set this attribute to False if the volume is mounted on other volume which is a part of an application service group. Note: This attribute is applicable only for volumes mounted as folder mounts and not for drive letter mounts.

Changes to MountV agent to determine whether the volume being monitored by the agent belongs to the fire drill disk group

A new attribute, ForFireDrill, is added to the MountV agent to determine whether the volume being monitored by the agent belongs to the fire drill disk group.

The ForFireDrill attribute can take values of 1 and 0. The value 1 indicates that the volume belongs to the fire drill disk group. Default value of the attribute is 0.

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

Changes to MountV agent to achieve performance improvement and overall reduction in service group failover times

The MountV agent is enhanced to address the issue where the MountV resources take a long time to go offline if there are applications accessing the clustered mount points.

A new option, CLOSE_FORCE, is added to the ForceUnmount attribute of the MountV agent to address this issue.

If the value of the ForceUnmount attribute is set to CLOSE_FORCE, the agent does not try to lock the configured mount points and proceeds directly with the forceful unmount operation. This reduces the service group failover time.

Note: Forceful unmount may potentially cause a data corruption. When you use ForceUnmount with CLOSE_FORCE, then before you switch or take the MountV resources offline, verify that none of the applications are accessing the configured mount points.

The MountV agent now uses the native operating system APIs to obtain the status of the configured mount points. This results in faster detection.

Changes to VMDg agent attributes

The VxVMFailAction attribute is deprecated. You can use the VxSVCFailAction attribute instead. Similarly, the VxVMRestartAttempts attribute is deprecated. You can use the VxSVCRestartAttempts attribute instead.

Enhanced Disaster Recovery Configuration Wizard for Microsoft Hyper-V

You can now deploy Hyper-V disaster recovery for application virtual machines (VMs) in a Microsoft cluster using the Disaster Recovery Configuration Wizard for Microsoft Hyper-V. The wizard assists you to perform the following tasks:

- Export the application VMs configuration files on the primary site
- Configure the network settings file on the primary site and DR site
- Configure global clustering (GCO)
- Configure replication monitoring between the primary site and DR site

See the *Veritas Storage Foundation™ and Disaster Recovery Solutions for Microsoft Hyper-V™* for more information.

New VVRSnap agent to make storage available for running fire drill in a disaster recovery environment

The VVRSnap agent is designed to support a fire drill for a disaster recovery site that uses Veritas Volume Replicator (VVR) as the replication method. See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the agent.

New agents added to enable vMotion and VMware Distributed Resource Scheduler (DRS) in VCS clusters configured and deployed on virtual machines in VMware environment

When a VCS cluster with a shared disk is configured on virtual machines, VMware does not support VMware Distributed Resource Scheduler (DRS) and vMotion. Thus the vMotion and DRS capabilities are compromised. The solution to this issue would be to attach the disk(s) to a single virtual machine at a time in a VCS cluster. In case of a user initiated failover or a fault induced failover, these disks would failover (detach-attach) to the target virtual machine along with the Service Group.

The following agents are added to the VCS Bundled agents to enable vMotion and VMware Distributed Resource Scheduler (DRS) in VCS clusters configured and deployed on virtual machines in VMware environment:

- VMwareDisks agent
  VMwareDisks agent manages the attaching and detaching of the disks to the virtual machines. After the attach or detach operations, the tasks of initializing and managing the disks are performed by theVMNSDg agent (for SFW environment) and NativeDisks agent (for LDM environment).
- VMNSDg agent
  The VMNSDg (Volume Manager Non-Shared Diskgroup) agent manages dynamic disk groups and mounts created on local (non-shared) and non-SCSI storage, which can work without reservation.
  The agent imports, monitors, and deports a dynamic disk group configured using Storage Foundation for Windows. The agent makes the disk group highly available. The agent is represented by the VMNSDg resource type.
  In combination with the VMwareDisks resource, theVMNSDg agent is responsible for:
  - During import operation, ensure that the disks newly attached to the virtual machine are accessible to Windows.
  - During deport operation, prepare the disks for the detach operation from the virtual machine.

Note: The Fire Drill Wizard does not support the VMNSDg resource as of now.

- NativeDisks agent
  The NativeDisks agent manages the disks created on local (non-shared) and non-SCSI storage, which can work without reservation.
  The agent brings the disks online, monitors them, and takes them offline.
  The agent makes the disks highly available. In combination with the VMwareDisks resource, the NativeDisks agent is responsible for:
  - While bringing the disks online, ensure that the disks newly attached to the virtual machine are accessible to Windows.
  - While taking the disks offline, prepare the disks for the detach operation from the virtual machine.

The service group configuration wizards currently do not support configuration of these agents in the service group.

If you are using a non-shared storage configuration in a physical system environment, you should configure these agents manually using the Cluster Manager (Java Console) or using the VCS HA commands from the command line. VCS provides templates for configuring application service groups that use non-shared storage agent resources. The templates are located in the following directory:

```
%VCS_HOME%\Templates
```

Here, %VCS_HOME% is the default product installation directory for VCS, typically, C:\Program Files\Veritas\Cluster Server.

To configure these agents in a VMware virtual environment, you can use the new Symantec high availability configuration wizard. However, this wizard currently supports configuration for SQL Server 2008 and custom applications.

Veritas Volume Replicator

The following are the new features and changes for Veritas Volume Replicator in release 6.0.1.

Support for the Adaptive Compression feature

This release provides support for the new Adaptive Compression feature for Veritas Volume Replicator (VVR). The Adaptive Compression feature enables VVR to reduce excessive CPU consumption caused byVVRcompression. This feature tracks the CPU utilization, and when the CPU usage crosses a user-defined threshold, Adaptive Compression tries to reduce it by reducing the compression

rate gradually to zero. In some cases, if required, it even disables the VVR compression. Adaptive Compression is useful if you have enabled VVR compression, which is a CPU-exhaustive process. By enabling Adaptive Compression, VVR reduces the CPU utilization and improves the system performance. Adaptive Compression can be enabled on the Primary or Secondary node using the CLI. At any time, you can disable this feature without affecting VVR compression process.

For more information, refer to the *Veritas Volume Replicator Administrator's Guide*.

New features and changes for 6.0

This section describes the new features and changes introduced in Veritas Storage Foundation for Windows (SFW) and Veritas Storage Foundation HA for Windows (SFW HA) 6.0.

Enhancements to the product installer

The following changes have been introduced in the product installer, in this release.

- Changes to the product packaging
  The product packaging has been modified in this release. The new package now includes the following products in a single disc.
  - Storage Foundation for Windows (SFW)
  - Veritas Storage Foundation and High Availability Solutions for Windows (SFW HA)
  - Veritas Cluster Server for Windows
  - Dynamic Multi-Pathing (DMP) for Windows

  Also, these products can be individually downloaded from the following location:
  https://fileconnect.symantec.com
- New interface
  The user interface of the product installer wizard has been modified in this release. The new-look installer now provides improved usability.
- Faster install
  The product packaging and installation logic has been enhanced in this release. These changes have significantly reduced the overall product installation time. You will now experience faster installation with the new installer.

- Changes to the installation options
  The list of available installation options has been modified in this release. The installation wizard now does not allow you to select the VCS agents for installation. These agents are installed by default with the Server components. The Server components also include the client components by default.
- New license options to enter the license details
  The product installer now provides the following two options to enter your license details:
  - Keyless
  - User Entered Key

  Note: Evaluation license keys are now deprecated.

  A keyless license installs the embedded keys and allows you to use all the available product options. You can use the keyless license for 60 days. If you install the product using the keyless option, a message is logged everyday in the Event Viewer indicating that you must perform any one of the following tasks, within 60 days of product installation. Failing this, a non-compliance error is logged every four hours.
  - Add the system as a managed host to a Veritas Operations Manager (VOM) Management Server.
    For more details, refer to *Veritas Operations Manager Administrator's Guide*.
  - Add an appropriate and valid license key on this system using the Symantec product installer from Windows Add/Remove Programs.

  In case of a User Entered Key license, you must procure an appropriate license key from the Symantec license certificate and portal. The user entered license allows you to use the product options based on the license key you enter.
  https://licensing.symantec.com/
  The product installer enables you to switch from a keyless license to a user entered license and vice-a-versa. It thus helps you to overcome the issues faced while removing the left-over temporary keys.

Veritas Storage Foundation

The following are the new features and changes for Veritas Storage Foundation in release 6.0.

Summary of new features

The following is a summary list of new features and changes for SFW for release 6.0:

- Support for the online volume shrink feature.
- Veritas Volume Manager (vxvm) service replaced with Veritas Enterprise Services (vxsvc).
- The Veritas Object Bus (vxob) or vxvm has been replaced with Veritas Enterprise Services (vxsvc).
- SFW does not support Automated System Recovery on Windows Server 2008. SFW will use the vxcbr utility for troubleshooting issues.
- VxBridge support removed from 6.0 release onwards.
- Support for Vxcache removed from 6.0 release onwards.
- Enhancements to the I/O statistics feature of Veritas Storage Foundation for Windows, which allows performance tuning to improve overall disk and system performance.
  Earlier releases of SFW supported I/O statistics only of a device, i.e., for a disk seen from a particular path. For better storage management, administrators should have options to see the I/O statistics at various object levels. With 6.0 release, I/O statistics is enhanced to provide I/O statistics of all the paths of a disk, all paths of an array, and all paths managed by the DSM.
- Site-aware allocation feature is included with the Storage Foundation for Windows (SFW) 6.0 release. This feature enables applications and services to function properly at a site when other sites become inaccessible.

Pending timeout behavior change for Microsoft Failover Clusters

Due to VMDg resource timeout, the Application Group fails to come online. To avoid this, following change is made:

- Added check pointing mechanism to report to cluster nodes that VMDg resources are coming online. It renews the Pending Timeout for resources, which are waiting in queue for getting online.

Support for fast failover in clustered environments

Fast failover improves the failover time for the storage stack configured in service groups in a clustered environment. Fast failover includes several design changes and enhancements to the core SFW components. These changes provide significant performance improvement in the failover times taken by storage resources in

service group failovers. Faster service group failovers are particularly noticeable in clusters with large storage configurations, typically over 20 disk groups and over 150 volumes.

Changes to SFW include the following:

- A new mode, Deported Read-Only, is added to the existing dynamic disk group entity. In this mode, a disk group is imported in a Read-Only mode on all the passive nodes. The disk group configuration on the active node is automatically reflected on the passive nodes. During failover, instead of a complete disk group deport and Read/Write import, only a mode change occurs on the passive node.
- A new attribute, FastFailOver, is added to the VCS VMDg agent. This allows the agent to enable the configured disk groups to support fast failover. You need to enable this attribute to configure fast failover.
- The vxdg command is modified to display the new Deported Read-Only disk group state introduced for fast failover. The vxdg list and vxdg dginfo options now display the new Deported Read-Only disk group mode.
- The vxdg command has a new option, vxdg refreshff, which refreshes the state of disk groups on a system in order to support the new Deported Read-Only state. You must run this command if the disk group state on a passive node does not display the Read-Only state even though FastFailover is enabled for the VMDg resources.

Fast failover currently does not support the following:

- RAID-5 volumes
- SCSI-2
  The disk groups will not be able to use the fast failover feature if configured using SCSI-2.
- Active/Passive (A/P) arrays for DMP
  Only A/PF, A/A , and ALUA arrays are supported.

Refer to the *SFW Administrator's Guide* for more information about fast failover.

Aggregate I/O Statistics

Enhanced I/O statistics to provide I/O statistics of a Disk (I/O statistics of all the paths of a disk), Array (I/O statistics of all the Paths of an array), and DSM (I/O statistics of all the Paths managed by the DSM). Earlier releases of SFW 5.1 SP2 supported I/O statistics only of a device, i.e., a Disk seen from a particular path.

Site-aware allocation for campus clusters

Site-aware allocation feature is included with the Storage Foundation for Windows (SFW) 6.0 release. This feature enables applications and services to function properly at a site when other sites become inaccessible. It means that even during site disruption at least one complete plex of a volume is available at each site.

Storage Migration of SFW and Hyper-V Virtual machine data

SFW provides the ability to move volumes to new storage locations via the Storage Migration Wizard. Storage Migration feature facilitates moving multiple volumes to different set of disks while the volumes are still online. Volumes associated with a Hyper-V Virtual Machine (VM) or a SFW disk group can be moved in a single administrative operation while the volumes are online without stopping the applications or Hyper-V VMs.

For details refer to the *Storage Foundation for Windows Administrator's Guide*.

Live Migration support for SFW and Hyper-V Virtual Machine configuration

SFW and Hyper-V Virtual Machine (VM) configuration can be live migrated between nodes of a Microsoft Failover Cluster. Live Migration of a Hyper-V VM is achieved through the use of Windows Server 2008 R2 Failover Cluster feature. Live Migration significantly increases availability of the virtual machines during planned and unplanned downtime.

Live migration enables active virtual machines to be moved from one physical Hyper-V-based host server to another virtual machine within a Failover Cluster without any disruption or noticeable loss of service.

For details refer to *Veritas Storage Foundation and Disaster Recovery Solutions for Microsoft Hyper-V*.

Support for the online volume shrink feature

This release provides support for the online volume shrink feature for Veritas Storage Foundation for Windows (SFW). Using the online volume shrink feature, you can decrease or shrink the size of your dynamic volumes. This feature is helpful in reclaiming unused space to better utilize your resources. It calculates the amount of space that can be freed from the volume to create a new smaller volume size. The size of a volume after the volume shrink operation is approximately the difference of the current volume size and the amount of maximum reclaimable bytes.

For more information, refer to the *Veritas Storage Foundation™ Administrator's Guide*.

Extended attributes displayed for arrays and LUNs

SFW is now able to discover and display extended attribute information related to arrays and LUNs. This information has been added to the disk view and enclosure view in the VEA. Extended attributes can also be viewed using the `vxdisk diskinfo` command with the `-e` option.

For enclosures, the cabinet serial number is displayed as Cabinet ID.

The disk information shown for an array can include the following attributes, if available for that array type:

- Vendor ID
- Product ID
- Revision ID
- Cabinet Serial Number
- Array Volume ID
- Array LUN Type
- Array RAID Level
- Array Snapshot LUN
- Array Replication LUN
- Array Media Type
- Array Transport Protocol
- Array Port WWN
- Array Port Serial Number
- Array Controller ID
- Array Hardware Mirror

New enclosure naming convention based on Array Vendor ID (AVID)

The Veritas Device ID (VDID) library now supports naming devices based on Array Vendor ID (AVID). AVID naming provides meaningful names to storage teams and the names are consistent across servers. It allows identifying the physical array for a particular LUN.

Previously LUNs from the same array were collected under a separate enclosure named with the vendor ID of the array followed by number 0, 1, 2 etc. For example,

two EMC arrays would be named EMC0, EMC1. This naming convention does not identify the physical array.

Under the new naming convention, for supported arrays the enclosure name consists of VendorID_ProductID_CabinetSerialNo. For a generic enclosure the name has the format DISKS@*fully qualifiedhost name*. For example, the enclosure name EMC_CLAriiON_CK2000650002000 means disks under this enclosure are associated with LUNs which are from the EMC CLAriiON array having cabinet serial number CK2000650002000.

On upgrade to SFW/SFW HA 6.0 from a previous version, all existing enclosure names will be changed to follow the new enclosure naming convention. The cabinet serial number attribute will get added for each enclosure object.

Ability to map logical disk viewed in VEA GUI with actual LUN in an array

The VEA GUI can now display the vendor ID, product ID, cabinet serial number, and LUN serial number attributes associated with each LUN. This enables mapping the disk viewed in the VEA GUI with the actual LUN and the array associated with it.

For example, the information on display about a disk named Harddisk 8 can show that it is associated with LUN3, from an array with the vendor ID NETAPP, with the product ID LUN, cabinet serial number 311249411, and LUN serial number hpT3ZZwkSD

Configuring DSMs without storage connection

On Windows Server operating systems, DMP DSM will enable configuration of some DSM settings without a storage connection. When a storage array is later connected to a host under a MPIO management, the connected storage array will inherit the policy that has been previously set in a DSM.

Support for DSM settings without storage connection applies to operations that are general for storage devices without specific knowledge of a DMP path of a storage device.

The following are settings you can configure for a DSM without storage connection:

- Load balances: Active/Active, Active/Passive without choosing a primary path, Least Block, Balanced Path and Least Queue. The specific settings for Round Robin with Subset or Weighted Path load balances cannot be configured without a storage connection because the settings require the specific knowledge of a DMP path of a disk which does not exist yet.
- SCSI-3: Enabling or disabling the setting for SCSI-3 support.

The VEA DSM Configuration command enables viewing a list of installed DSMs on a system and applying load balance settings and SCSI-3 support settings to the selected DSMs. These capabilities are also available through the `vxdmpadm` utility in the command line interface (CLI). In addition, you can use the CLI `vxdmpadm getdsmattrib` command to view existing settings for the installed DSMs.

MPIO timer parameter settings

New DMP DSM CLI commands are provided for viewing and changing Microsoft MultiPath Input/Output (MPIO) timer parameter settings:

- `vxdmpadm getmpioparam`
- `vxdmpadm setmpioparam`

Warning: Symantec DSMs always use the default MPIO parameters. Attempting to change these MPIO parameters would affect the behavior of a DSM for I/O error, path failover, and DSM performance. Therefore, MPIO parameter settings should not be changed unless a customer has been advised by Microsoft to change the settings for debugging purposes. The `vxdmpadm setmpioparam` command provides a parameter to restore the default settings.

MPIO timer parameter settings are as follows (see MPIO timer on the Microsoft website for more details):

- PathVerifyEnabled: When enabled, MPIO will perform path verification for the amount of time specified in PathVerificationPeriod.
- PathVerificationPeriod: Specifies the amount of time MPIO will perform the path verification, if enabled.
- PDORemovePeriod: Specifies the amount of time an MPIO pseudo LUN, which represents a disk under MPIO control, will stay in memory after all paths of a disk have been removed from the system. It also specifies how long the pending I/O should fail after all paths have been removed from a disk.
- RetryCount: Specifies the number of times DSM will ask MPIO to retry the I/O when an I/O error occurs.
- RetryInterval: Specifies the amount of time MPIO should retry a failed I/O.

Veritas Cluster Server

The following new features and changes are introduced in Veritas Cluster Server in this release:

Changes to VMDg agent to support fast failover

To support the fast failover feature, a new attribute, FastFailOver, is added to the VCS Volume Manager Diskgroup (VMDg) agent. This attribute decides whether or not a disk group is enabled for fast failover.

The FastFailOver attribute can take values of 1 and 0. The value 1 indicates that the agent enables fast failover for the configured disk group. The default value 0 indicates that fast failover is disabled for the disk group.

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

Refer to the *SFW Administrator's Guide* for more information about fast failover.

Support for IPv6

Veritas Cluster Server adds support for the Internet Protocol version 6 (IPv6). The scope of support is as follows:

Note: Support is limited to mixed mode (IPv4 and IPv6) network configurations only; a pure IPv6 environment is currently not supported.

Types of addresses	The following types of IPv6 addresses are supported: ■ Unicast addresses Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	LLT over UDP is supported on both IPv4 and IPv6. You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.
VCS agents	■ VCS introduces a new IPv6 agent to support IPv6. The IPv6 agent monitors IPv6 addresses and the associated network prefix. Refer to the *VCS Bundled Agents Reference Guide* for more details about the IPv6 agent. ■ VCS agents that require an IP address attribute now support IPv6 address types described earlier.

Configuration wizards	■ All the VCS configuration wizards that configure or discover IP addresses now support IPv6 address types described earlier. This includes the Cluster Configuration Wizard (VCW), service group configuration wizards for applications such as Exchange and Oracle, and the solutions wizard such as the Disaster Recovery Configuration Wizard. ■ The wizards display options to choose the IP protocol during configuration. In case of IPv6 you are required to select the network prefix. The wizards use the prefix and automatically generate IPv6 addresses that are valid and unique on the network.
Other components	The VCS High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.

Embedded security services for secure cluster configuration

The Symantec Product Authentication Service (SPAS) is now embedded in to the product and has been renamed as VCS Authentication Service. This service is used to provide secure communication between the cluster nodes and the clients such as Cluster Manager (Java Console) and services. In a secure mode VCS uses digital certificates for authentication and uses SSL to encyrpt communication over the public network.

You no longer need to configure an additional system in your enterprise network to function as a root broker. During cluster configuration, VCS now configures all the cluster nodes as root brokers (RB) and Authentication Brokers (AB) and creates a copy of the certificates on all the cluster nodes. The VCS Cluster Configuration Wizard (VCW) provides an option to configure a secure cluster using the embedded security services.

Instantaneous fault detection using Intelligent Monitoring Framework (IMF)

VCS introduces Intelligent Monitoring Framework (IMF) that uses an event-driven design for monitoring VCS resources configured in a cluster. IMF is asynchronous and provides instantaneous resource state change notifications. This significantly improves the fault detection capability allowing VCS to take corrective actions faster. IMF works in addition to the poll-based monitoring.

The benefits of intelligent monitoring over poll-based monitoring are as follows:

- Instantaneous notification
  Faster notification of resource state changes result in improved service group failover times.

- Reduction in system resource utilization
  Reduced CPU utilization by VCS agent processes when number of resources being monitored is high. This provides significant performance benefits in terms of system resource utilization.
- Ability to monitor large number of resources
  With reduced CPU consumption, IMF enables VCS to effectively monitor a large number of resources.

See the *Veritas Cluster Server Administrator's Guide* for more information.

The following VCS agents support IMF-based monitoring:

- GenericService, ServiceMonitor
- IP, NIC
- VMDg, MountV, Mount
- Oracle, NetLsnr
- Process, RegRep
- SQLServer2005, SQLAgService2005, SQLOlapService2005, MSDTC
- SQLServer2008
- IIS
- ExchService2007, Exchange2010DB

Cluster ID can range from 0 to 65535

VCS LLT has been enhanced to support 16-bit cluster ID values. The cluster ID can now range from 0 to 65535 (was 0 to 255 earlier).

You can specify the cluster ID while configuring the cluster using the VCS Cluster Configuration Wizard (VCW) or using the silent configuration utility (vcwsilent).

VCS Simulator is not included with software

The VCS Simulator is no longer included as part of the software or the software disc. The functionality and features offered by the Simulator are supported; only the installer is no longer shipped with the software.

To download VCS Simulator, go to http://go.symantec.com/vcsm_download.

Veritas Volume Replicator

The following new features and changes are introduced in this release.

Support for the online volume shrink feature

This release provides support for the online volume shrink feature for Veritas Storage Foundation for Windows (SFW) and Veritas Volume Replicator (VVR). Using the online volume shrink feature, you can decrease or shrink the size of your data volumes. This feature is helpful in reclaiming unused space to better utilize your resources. It calculates the amount of space that can be freed from the volume to create a new smaller volume size. The size of a volume after the volume shrink operation is approximately the difference of the current volume size and the amount of maximum reclaimable bytes.

For more information, refer to the *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

Support for IPv6

In this release Veritas Volume Replicator (VVR) includes support for IPv6 addresses. You can specify IPv6 addresses for configuring replication.

Note the following:

- You must set the IP preference, whether VVR should use IPv4 or IPv6 addresses, before configuring replication.
  When you specify host names while configuring replication, VVR resolves the host names with the IP addresses associated with them. This setting determines which IP protocol VVR uses to resolve the host names.
  Use Veritas Enterprise Administrator (VEA) (Control Panel > VVR Configuration > IP Settings tab) to set the IP preference.
- The Replicated Data Set (RDS) wizard now allows you to specify IPv6 addresses associated with the primary and secondary host names.
- The VVR Security Service Configuration Wizard allows you to specify IPv6 addresses for hosts on which you wish to configure the VxSAS service.
- VVR commands that use an IP address, either as an input parameter or as an output, now support IPv6 addresses.
- VVR does not support replication in cases where the primary and secondary systems in an RDS use different IP addresses. For example, if the primary host uses an IPv4 address and the secondary host uses an IPv6 address, this configuration is not supported.
  In cases where the primary host uses only an IPv4 address, and the secondary host uses both IPv4 and IPv6 addresses, VVR automatically selects an IPv4 address for the secondary.

- VVR does not support replication for a IPv6-only system. An IPv6-only system is a system that implements only IPv6. It only has an IPv6 address in the name service database.

No longer supported

This section lists of features deprecated from this release.

Vxcache support deprecation

Vxcache feature is not supported for 6.0 release.

VVRDCOMBridge support deprecation

VVRDCOMBridge is no longer supported on 6.0.

Windows Server 2003 (x86 and x64) and Windows Server 2008 (x86)

SFW and SFW HA Server and Client components are no longer supported on Windows Server 2003 (x86 and x64) and Windows Server 2008 (x86).

Microsoft Exchange 2003 and Microsoft SQL Server 2000

SFW and SFW HA do not support Windows Server 2003 (x86 and x64) in this release. Hence Microsoft Exchange 2003 and Microsoft SQL Server 2000 are also no longer supported. The SFW HA agents for Exchange 2003 (ExchService, ExchProtocol) and SQL Server 2000 (SQLServer2000, MSSearch) are no longer available.

Microsoft Operations Manager (MOM) 2005

Microsoft Operations Manager (MOM) 2005 is no longer supported in this release.

The hasnap command

The hasnap command has been deprecated in this release.

The hasnap command can be used to take snapshots of the VCS configuration files on cluster nodes. The command can also be used to back up, restore, compare, and export the VCS configuration files.

Software limitations

The following software limitations apply to this release of the product. For the latest information on updates, patches, and software issues regarding this release, see the following TechNote:

http://www.symantec.com/docs/TECH161556

Cluster operations performed using the Symantec High Availability dashboard may fail

This issue occurs while performing the cluster operations in a multi-system VCS cluster that is configured in a VMware virtual environment. The cluster configuration is such that a single system belongs to one datacenter or ESX, and all the other systems belong to another datacenter or ESX. (2851434)

In such a cluster configuration, the operations initiated from the dashboard that is available from the datacenter or ESX to which the single cluster system belongs, may fail on the systems that belong to the other datacenter.

This situation arises if the following changes have occurred with the system:

- The system has lost its network connectivity
- The SSO configuration has become corrupt

This occurs because the operations are performed using the network details of the system that belongs to the datacenter or ESX from where they are initiated.

Volume Shadow Copy Service is not supported

The MountV agent is not supported on volumes with the copy-on-write feature of Volume Shadow Copy Service enabled.

VCS lock on shared volumes during Exchange recovery

VCS monitors the shared volume used for storing Exchange databases. During online, offline, or clean operations, VCS MountV resources exclusively lock the shared volume. This exclusive lock may conflict with recovery of an Exchange volume.

Workaround: Symantec recommends freezing the service group containing the MountV resources before recovering Exchange volumes.

To recover an Exchange volume that is monitored by VCS

1. In the VCS Java Console, identify the service group containing the MountV resources corresponding to the volume to be recovered.
2. Freeze the service group.
   - In the **Service Groups** tab of the configuration tree, right-click the service group name.
   - Choose **Freeze**, then choose **Temporary** or **Persistent** from the menu.
3. Recover the Exchange volume.
4. Unfreeze the service group.
   - In the **Service Groups** tab of the configuration tree, right-click the service group name.
   - Choose **Unfreeze**, then choose **Temporary** or **Persistent** from the menu.

If custom resources are configured in VCS to monitor a snapshotted volume, follow the procedure before snapping back to the original or the replica.

Note: If you cannot lock a volume for snapback, you can either force the operation or fail the operation and await administrator intervention.

Installation or upgrade

The following are installation or upgrade software limitation issues.

UUID files are always installed to the default installation path

During product installation, you can specify a different installation path than the default.

However, the installation process installs the UUID files in the following default path regardless of where the other binaries are installed:

```
C:\Program Files\Veritas\UUID\bin
```

License management

The following is a license management software limitation.

Silent installation does not support updating license keys after install

You can install SFW or SFW HA using either the product installer or the command line interface (for a silent installation).

Both installation methods enable you to specify license keys during product installation. The product installer also includes the functionality to update license keys after installation. However, the command line interface used in a silent installation does not support updating license keys after an installation.

To add license keys after a silent installation using the CLI, you use the `vxlicinst` utility located on the SFW HA product DVD:

To add license keys after silent installation using CLI

1. Insert the product DVD in a drive that you can access from the system on which you want to add the license.
2. Navigate to the vxlic_util directory on the product DVD:

   ```
   <DVD_ROOT_DIRECTORY>\Tools\storage_foundation_for_windows\vxlic_tools
   ```

3. Type the command as follows to specify the key to be added:

   ```
   vxlicinst -k <key>
   ```

You can also access the `vxlicinst` utility after an installation in the Volume Manager install directory.

The directory is: *%VMPATH%*.

Veritas Storage Foundation

This section covers limitations specific to Storage Foundation for Windows product functionality.

Only one disk gets removed from an MSFT compatible disk group even if multiple disks are selected to be removed

Only one disk gets removed from an MSFT compatible disk group even if multiple disk are selected to be removed. If such a disk group needs to be deleted, then remove disk operation has to be performed individually on all the disks in the disk group.(2581517)

Cannot create MSFT compatible disk group if the host name has multibyte characters

If a system or host name consists of multibyte characters, then it is observed that creating a Microsoft compatible disk group on such a system fails with the error message **Failed to migrate a basic disk to a VDS dynamic pack**.(2579634)

Fault detection is slower in case of Multipath I/O over Fibre Channel

If the storage is configured with Multipath I/O (MPIO) over Fibre Channel (FC), the storage framework takes more time to detect a storage failure and notify SFW. This results in a delay in the fault detection. (2245566)

Typically, fault detection takes 30 seconds or more without MPIO and up to a minute when MPIO is used over FC.

DSM ownership of LUNs

Do not use a DMP DSM together with a third-party DSM for the same array. Only one DSM at a time can claim the LUNs in an array. According to Microsoft Multipath I/O (MPIO) documentation, if multiple DSMs are installed, the Microsoft MPIO framework contacts each DSM to determine which is appropriate to handle a device. There is no particular order in which the MPIO framework contacts the DSMs. The first DSM to claim ownership of the device is associated with that device. Other DSMs cannot claim an already claimed device. Therefore, to ensure that the DMP DSM claims the LUNs of an array, no other DSM should be installed for that same array.

SFW FlashSnap solution for EV does not support basic disks

The SFW FlashSnap solution for Enterprise Vault (EV) is not supported for EV databases that reside on basic disks.

The vxsnap command may fail with the following errors:

```
V-76-58657-2053: Failed to prepare some of the EV components.

See EVStatus.log for more details.

The operation to prepare the volumes for a snapshot has failed.

See the log file for details.
```

To use the FlashSnap solution for EV, ensure that the EV databases reside on dynamic disks.(2116246, 2122790)

Incorrect mapping of snapshot and source LUNs causes VxSVC to stops working

After mapping of snapshot LUNs to the host containing the source LUNs or mapping of the source LUNS to the host containing the snapshot LUNs, the following issues may occur:

- The `VxSVC` service goes into an invalid state and may stop working
- The host shows “unknown DG” for the snapshot LUNs disk group

To avoid these issues, do not connect the snapshot LUNs to the same host containing the source LUNs, or the source LUNs to the same host containing the snapshot LUNs. (2871055, 2794524)

It is recommended that you use Volume Shadow Copy Service (VSS) to take a snapshot and to import the snapshot LUN, preferably on a different node.

SFW does not support operations on disks with sector size greater than 512 bytes; VEA GUI displays incorrect size

No SFW operations are supported on disks with the sector size greater than 512 bytes. Similarly, VEA GUI displays incorrect size for such disks.

Limitations on 64-bit systems

This section describes limitations on 64-bit systems.

Limitations of SFW support for Dynamic Multi-pathing (DMP)

This section describes limitations of SFW support for Dynamic Multi-pathing (DMP).

Load balancing policies of third-party MPIO DSMs are not supported in SFW

Load balancing policies and path settings of third-party MPIO DSMs are not supported in SFW. This is because third-party MPIO DSMs may not implement a common method in the Microsoft MPIO framework for getting or setting load balancing policies. (820077)

Disconnected paths may not be reflected in VEA GUI with MPIO DSMs installed

Disconnecting paths from a host using MPIO DSMs may not be reflected in the VEA GUI. The VEA GUI is not automatically updated because of a communication problem between SFW and WMI. (326603)

Workaround: Perform a rescan operation to allow SFW to obtain information about the disconnected paths.

Other issues

The following are other issues:

Operations in SFW may not be reflected in DISKPART

If you perform an operation in DISKPART, it is reflected in the VEA GUI and the CLI. However, operations that are performed in SFW may not be automatically reflected in DISKPART. (100587, 101776)

Workaround: The workaround is to rescan in DISKPART to obtain these changes. The DISKPART utility does not support multiple disk groups, so it cannot reflect multiple disk groups that were created in SFW. DISKPART does indicate whether a disk is basic or dynamic

Disk signatures of system and its mirror may switch after ASR recovery

After an ASR recovery of a system with a mirrored system and boot disk, the disk signatures of the original system and boot disk and its mirror are sometimes switched.

The problem happens as a result of Microsoft's disk mapping algorithm. Under some conditions, the algorithm switches disk signatures. This is a known Microsoft issue. (100540)

Adding a storage group that contains many disks and volumes causes SFW and Microsoft Exchange System Manager to respond very slowly.

Adding or creating a storage group that has a dynamic disk group that contains many disks and volumes to an MSCS Exchange Virtual Server causes the VEA GUI and the Exchange System Manager GUI to respond very slowly. It seems that a greater number of disks and volumes increases the response time. This is a known Microsoft problem (SRX060621604113).(530035)

SFW does not support growing a LUN beyond 2 TB

Growing a dynamic disk that has the MBR partition style to a size of 2 TB or greater renders the disk unusable.(704839)

SFW cannot coexist with early Symantec Anti-virus software

Abnormal termination of SFW occurs when Symantec Anti-virus version 11.6.2 coexist on a system. (804143)

Workaround: Upgrade to Symantec Anti-virus version 11.6.8 or later.

SCSI reservation conflict occurs when setting up cluster disk groups

Setting up a cluster on Windows Server operating systems creates physical disk resources for all the basic disks on the shared bus. Later you create resources for the SFW cluster disk groups. Before doing so, you must remove any physical disk group resources for disks that are used in the cluster disk groups. Otherwise, a reservation conflict occurs.

Snapshot operation fails when the Veritas VSS Provider is restarted while the Volume Shadow Copy service is running and the VSS providers are already loaded

When the Volume Shadow Copy VSS service starts, it loads the Veritas VSS provider. If the Veritas VSS provider is restarted while the Volume Shadow Copy service is running and the VSS providers are already loaded, the snapshot operation fails with a VSS error (Event ID:12293).

When a node is added to a cluster, existing snapshot schedules are not replicated to the new node

When you create snapshot schedules in a clustered environment, schedule-related registry entries are created on all cluster nodes. When a failover occurs, the failover node can continue to run the schedules. However, if a new node is added to a cluster after the schedules are created, the schedules are not replicated to the new node. If the service group fails over to the node that was added, the scheduled snapshot tasks do not occur.

Workaround: Start the Quick Recovery Configuration Wizard from the Solutions Configuration Center (**Start**>**Run**>**scc**). Continue through the wizard until the **Synchronizing Schedules** panel shows that synchronization between cluster nodes is complete. Click **Finish** to exit the wizard.

Restore from Copy On Write (COW) snapshot of MSCS clustered shared volumes fails

On Windows Server operating systems, the restore operation using a COW snapshot of MSCS clustered shared volumes fails. This is a known Microsoft problem (KB945361). (1796788)

Dynamic Disk Groups are not imported after system reboot in a Hyper-V environment

In a Hyper-V environment, dynamic disk groups that reside on virtual disks that are attached to a SCSI controller are not imported automatically. This is a known Microsoft problem. (1406512)

Workaround: Configure the system to use the Veritas DG Delayed Import Service (VxDgDI) for these dynamic disk groups. Alternatively, you can manually import these disk groups after the system has completed the boot process.

Storage Agent cannot reconnect to VDS service when restarting Storage Agent

Stopping the VDS service while a VDS client is running on a system, results in a system error. Subsequently, stopping the Storage Agent and then restarting the Storage Agent, results in the Storage Agent not being able to reconnect to the VDS service.

All VDS clients, such as DISKPART, Storage Agent, or the Disk Management GUI, must be closed to avoid errors when stopping the VDS service and to enable the Storage Agent to be started again.(1794522)

Workaround: When the VDS service is stopped resulting in a system error, the vxvdsdyn.exe and vxvds.exe processes must be terminated. Also ensure that the vds.exe process has been terminated.

Use the following commands to stop these processes:

```
TASKKILL /F /IM vxvdsdyn.exe
TASKKILL /F /IM vxvds.exe
TASKKILL /F /IM vds.exe
```

At this point, restarting the Storage Agent restarts the VDS service automatically.

SFW does not support transportable snapshots on Windows Server

SFW does not support transportable snapshots on Windows Server operating systems

Windows Disk Management console does not display basic disk converted from SFW dynamic disk

A basic disk that was converted from an SFW dynamic disk does not appear in the Windows Disk Management console or in the results of the `DISKPART list disk` command. (930388)

Workaround: The disk can be displayed in the Windows Disk Management console by performing a refresh or a rescan disks operation. In addition, the disk can be displayed in the results of the `DISKPART list disk` command by performing a DISKPART rescan operation first.

SharePoint components must have unique names

When creating SharePoint components, ensure that the names of the components are unique. Performing operations on components with names that are not unique may cause unpredictable results.(1851186)

DCM or DRL log on thin provisioned disk causes all disks for volume to be treated as thin provisioned disks

Having a volume on a disk that is not a thin provisioned disk and then adding a DCM or DRL log that resides on a thin provisioned disk to the volume, causes the volume to be enabled for thin provision disk operations. Performing thin provision disk operations in this situation causes the operations to fail.(1601143)

After import/deport operations on SFW dynamic disk group, DISKPART command or Microsoft Disk Management console do not display all volumes

The Microsoft Disk Management console and DISKPART CLI command may not display all volumes after repeated import/deport operations are performed on an SFW dynamic disk group.

Symantec recommends that using SFW CLI commands instead of the Microsoft DISKPART command for scripts to monitor the status of volumes.

Restored Enterprise Vault components may appear inconsistent with other Enterprise Vault components

Selected Enterprise Vault components that were restored may appear to be inconsistent with Enterprise Vault components that were not restored. Inconsistencies may appear as dangling Saveset entries in a VaultStore database or index, or a Saveset component with missing Saveset entries in a database or index.

Symantec recommends that the user verify the restored component and components dependent on the restored component.

Use the EVSVR.exe tool (available with the Enterprise Vault installation) for the verification operation.(1671337)(1780009)

Note: Any discrepancies that are discovered can be repaired with the EVSVR.exe tool that is available with Enterprise Vault 8.0 SP2.

Enterprise Vault restore operation may fail for some components

The restore operation fails for an Enterprise Vault component when an open handle exists for a volume on which the component resides. (1788920)

Workaround: Specify the **Force** option in the Enterprise Vault restore wizard or CLI command to allow the operation to proceed successfully.

VDS limits the length of a volume name

On Windows Server 2008, Virtual Disk Service (VDS) limits the length of volume names. As a result, VDS commands may fail on SFW dynamic volumes that have names that are too long.

This is a known Microsoft issue and has been resolved with KB958556. KB958556 is automatically installed with this release. (1598800)

Note: You must activate Windows before installing or upgrading to SFW 5.1 SP1, otherwise the installation of KB958556 fails.

Shrink volume operation may increase provisioned size of volume

Performing a shrink volume operation on a volume that resides on a thin provisioned disk may result in an increase of the provisioned size of the volume.(1935664)

Reclaim operations on a volume residing on a Hitachi array may not give optimal results

Reclaim operations on a striped volume that resides on thin provisioned disks in Hitachi arrays may not give optimal results. This is due to the size of the allocation unit of the arrays. (1922235)

Veritas Cluster Server

This section describes the software limitations for Veritas Cluster Server.

NBU restore changes the disk path and UUID due to which VMwareDisks resource reports an unknown state

When you restore a VMware virtual machine using NetBackup (NBU), it changes the path and UUID of disks because of which the VMwareDisks agent resource goes into an unknown state as it has the old path and UUID configured in its “DiskPaths” attribute. As a workaround, you need to manually provide the new disk path in the “DiskPaths” attribute of the affected VMwareDisks resource and delete the incorrect UUID (and the colon after it) from the attribute. (2913645)

SQL service group configuration wizards fail to discover SQL databases that contain certain characters

The VCS SQL service group configuration wizards fail to discover SQL databases if the database name contains any of the following characters:

- " (inverted commas)
- , (comma)
- [ ] (square brackets)

MountV agent does not detect file system change or corruption

Even if IMF is enabled in the cluster, the VCS MountV resource cannot detect corruption or a change in the file system format. The MountV resource or the service group does not fault or fail over in the cluster. The agent is able to detect a fault only after the application writes begin to fail on the configured volumes. (2245295)

If the MountV agent attribute AutoFSClean is set to true and you take the resource offline and then bring it online again, the agent attempts to open a read-only handle to the volume. If it is unable to do so, it attempts to clean the file system using the Windows command Chkdsk /x. If the file system clean does not resolve the issue, the resource faults. The MountV agent logs contain a "File system is not clean" message to indicate this issue.

Windows Safe Mode boot options not supported

The Windows Safe Mode boot options are not supported. VCS services and wizards fail to run if Windows is running in Safe Mode.(1234512)

Security issue when using Java-GUI and default cluster admin credentials

While configuring the cluster using the VCS Cluster Configuration Wizard (VCW) if you do not choose the secure mode (Use Single Sign-on option) on the **Configure Security Service Option** panel, VCW creates a user with user name as admin and password as password. The user credentials are auto-populated in the respective fields, by default. This user has administrative privileges to the cluster.

Symantec recommends that you create a different user instead of accepting the default values.(1188218)

VCW does not support configuring broadcasting for UDP

VCW does not provide options to configure broadcasting information for UDP. You can configure broadcasting for UDP by manually editing the `llttab` file. Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

Cluster Manager (Java Console)

The following are Cluster Manager (Java Console) software limitations.

Latest version of Java Console for VCS is required

Cluster Manager (Java Console) from previous VCS versions cannot be used to manage VCS 6.0.1 clusters. Symantec recommends always using the latest version of Cluster Manager.

Running Java Console on a non-cluster system is recommended

Symantec recommends not running Cluster Manager (Java Console) for an extended period on a system in the cluster.

All servers in a cluster must run the same operating system

All servers in a cluster must run the same operating system. You cannot mix the following Windows operating systems within a cluster:

- Windows Server 2008 (x64) and Windows Server 2008 R2
- Windows Server 2008 (full install) systems and Windows Server 2008 Server Core

Service group dependency limitations

The following are Service group dependency software limitations.

No failover for some instances of parent group

In service groups in which the group dependency is configured as parallel parent/failover child, online global, remote soft or firm, the parent group may not online on all nodes after a child group faults.

System names must not include periods

The name of a system specified in the VCS configuration file, `main.cf`, must not be in the fully qualified form; that is, the name must not include periods. The name in main.cf must be consistent with the name used in the `llthosts.txt` file.

Incorrect updates to path and name of types.cf with spaces

The path of the `types.cf` file, as referenced in the `main.cf`, updates incorrectly if the path contains spaces. For example, C:\Program Files\, would update incorrectly. Running a combination of the `hacf` commands `hacf -cmdtocf` and `hacf -cftocmd` truncates the path of the `types.cf` file and updates the `main.cf` file with the truncated path.

Lock by third-party monitoring tools on shared volumes

Some third-party monitoring tools (such as Compaq Insight Manager) hold an exclusive lock or have an open file handle on the shared volumes they monitor. This lock may prevent VCS from offlining a service group that includes the volume as a resource. VCS requires a lock on resource in a service group when taking the group offline.

Workaround: Symantec recommends adding a custom resource as the topmost parent for an affected service group. Use the custom resource to manage onlining, monitoring, and offlining of the third-party monitoring tool.

Schedule backups on online nodes

If you are scheduling backups in a VCS cluster, schedule them on the node on which the service group is online. If the Exchange virtual server fails over to another node, you must set up the backup schedule again on the new node.

Undefined behavior when using VCS wizards for modifying incorrectly configured service groups

If you use the VCS wizards to modify service groups that are incorrectly configured through the VCS Cluster Manager (Java Console), the wizards fail to modify the service groups. This may also result in undefined behaviors in the wizards.(253007)

MirrorView agent resource faults when agent is killed

If all of the parent resources of the MirrorView Agent are offline when the MirrorView Agent is killed, or has crashed, then the resource will fault once the MirrorView Agent has automatically restarted. This behavior only occurs if all of the parent resources of the MirrorView agent are offline before the MirrorView Agent being killed, or crashing. (508066)

Cluster address for global cluster requires resolved virtual IP

The virtual IP address must have a DNS entry if virtual IP is used for heartbeat agents.

Systems in a cluster must have same system locale setting

VCS does not support clustering of systems with different system locales. All systems in a cluster must be set to the same locale.

Virtual fire drill not supported in Windows environments

The virtual fire drill feature available from the VCS command line and the Cluster Manager (Java console) is not supported in Windows environments.

Cluster Manager consoles do not update GlobalCounter

To avoid updating Cluster Manager views with unnecessary frequency, the Java and Web Console do not increment the GlobalCounter attribute of the cluster.

Symantec Product Authentication Service does not support node renaming

Symantec Product Authentication Service (earlier known as Veritas Security Services) does not support renaming nodes.

WAN cards are not supported

The VCS Configuration Wizard (VCW) does not proceed with network card discovery if it detects a WAN card.

Enterprise Vault Task Controller and Storage services fail to start after running the Enterprise Vault Configuration Wizard if the MSMQ service fails to start

The Enterprise Vault (EV) Task Controller Service and Storage Service are dependent on the Message Queuing Service. If MSMQ is not configured correctly, the MSMQ service may fail to start. In that case, the EV Task Controller and Storage services fail to start after you finish running the Enterprise Vault Configuration Wizard, and you will not be able to bring the EV service group online.

Before running the Enterprise Vault Cluster Setup Wizard to configure the EV service group, ensure that the MSMQ service can be started normally.

Note: The MSMQ service may fail to start after you move the MSMQ storage directory from a default system drive to new directory on a non-system drive if you create the new directory manually. Instead, use the following steps to generate the new directory automatically and move the storage to it. Select **Computer Management > Services and Applications > Message Queuing > Properties**. In the **Storage** tab, browse to the non-system drive location and specify a directory name to be created. Click **Apply** and then **OK** to close the Properties window.

Note: The MSMQ Service may also fail to start due to a problem with the incoming sequence checkpoint files (MSMQ Event 2053). For a description of that problem and the workaround for it, refer to the following Technote:

http://www.symantec.com/docs/TECH87839

Veritas Volume Replicator

This section covers limitations specific to Veritas Volume Replicator (VVR).

Resize Volume and Autogrow not supported in Synchronous mode

The Resize Volume and Autogrow operations are not supported when replication is done in Synchronous mode. While Synchronous replication is paused to resize volumes, writes necessary to grow the file system cannot occur. (103613)

Workaround: To resize the volume, temporarily change the mode of replication to Asynchronous or Synchronous Override. After you finish resizing the volume, you can switch replication back to the Synchronous mode.

Expand volume not supported if RVG is in DCM logging mode

VVR does not support the Expand Volume operation if the Replicated Volume Group (RVG) is in DCM-logging mode.

Fast failover is not supported if the RLINK is in hard synchronous mode

In synchronous mode of replication, if fast failover is set, then the RVG cannot be stopped and started when a disk group fails over to another node. If the RLINK is in hard synchronous mode, it may not be connected when the volume arrives, and the I/Os may fail. In such case, the Event Viewer displays NTFS errors and file system reports the volume as RAW. Therefore, fast failover is not supported if the RLINK is in hard synchronous mode. (2711205)

Solutions Configuration Center

This section describes the Solutions Configuration Center software limitations.

Quick Recovery wizard displays only one XML file path for all databases, even if different file paths have been configured earlier

When running the Quick Recovery wizard, the XML file path you specify applies to all the databases selected in that run of the wizard. If you schedule databases in separate runs of the wizard, you could specify a different XML file path for each database. However, if you later run the wizard to modify the snapshot schedule and select more than one database, the Quick Recovery wizard displays the XML file path for the first database only.

Workaround: If you want to view the XML file path of each database, run the wizard again and specify one database at a time to modify.

Disaster Recovery, Fire Drill, and Quick Recovery wizards fail to load unless the user has administrative privileges on the system

Disaster Recovery, Fire Drill, and Quick Recovery wizards require that the user have administrative privileges on the system where they are launched. If a user with lesser privileges, such as user privileges, tries to launch the wizards, the wizards will fail to load, with the message "Failed to initialize logging framework".

Discovery of SFW disk group and volume information sometimes fails when running Solutions wizards

Discovery of Storage Foundation for Windows disk group and volume information may fail when running a Solutions wizard. This issue applies to the Fire Drill Wizard, Quick Recovery Configuration Wizard, or the Disaster Recovery Configuration Wizard.(1802119)

To workaround this known discovery failure issue

1. Make sure that the Storage Agent service is running on the target system.
2. From the VEA console, click **Actions** > **Rescan** to perform a rescan.
3. Restart the wizard.

DR Wizard does not create or validate service group resources if a service group with the same name already exists on the secondary site

If a service group with the same name as the one selected in the Disaster Recovery Wizard already exists on the secondary site, the Disaster Recovery Wizard does not validate the configuration or add missing resources.

Workaround: Remove the service group with the same name that exists on the secondary site. Then run the wizard again so that it can clone the service group that exists on the primary site.

Solutions wizard support in a 64-bit VMware environment

In a 64-bit VMware virtual machine environment, the Disaster Recovery, Quick Recovery, and Fire Drill wizards are supported on VMware ESX 3.5 and above. No support is provided for VMware Workstation version.

Known issues

The following known issues exist in this release of the product.

For the latest information on updates, patches, and software issues regarding this release, see the following TechNote:

http://www.symantec.com/docs/TECH161556

Issues relating to installing, upgrading and licensing

This section provides information on the issues you may face during the product installation, upgrade or during managing the licenses.

Side-by-side error may appear in the Windows Event Viewer

While installing SFW or SFW HA, the installation progress may get hung and the Windows Event Viewer of that system displays a Side-by-Side error.

This issue occurs on the systems where the Microsoft VC redistributable package or the Dot Net installation is corrupted. (2406978)

Workaround: You must repair the VC redistributable package or the Dot Net installation.

Installation may fail with "Unspecified error" on a remote system

The SFW or SFW HA installation may fail on a remote system with "Unspecified error".

This issue occurs if the vxInstaller service does not start on the remote node to begin the installation. (2429526)

Workaround: Run the installation locally on the system where the installation has failed.

Installation may fail with a fatal error for VCS msi

The product installation wizard may fail to install SFW HA with a fatal error for installing the VCS msi. This error occurs on the Installation panel of the product installation wizard.

During the installation the product installer accesses the user profile folder and the SID path for the logged on user. While logging in to the system, if the user profile does not load properly or if the logged on user profile is corrupt, the product installer fails to perform the required installation task. This causes the installation to fail with a fatal error. (2515584)

Workaround: Reboot the system and run the installation again. If the problem persists, contact your system administrator.

Delayed installation on certain systems

You may experience a slower installation on certain systems.

This issue occurs, if you have configured any software restriction policies on the system. During the installation the restriction policies increases the package verification time and thus the over all installation time is increased. (2516062)

Installation may fail with the "Windows Installer Service could not be accessed" error

This issue occurs if the Windows Installer Service is not accessible during the installation. Since the service is not accessible, the installer fails to proceed with the installation. (2497344)

Workaround: The Windows Installer Service is a native component of an operating system. Typically, the Installer Service inaccessible issue occurs, if the service is damaged or unregistered and thus repairing the operating system installation serves as a workaround.

For more details on the workaround, refer to the following Microsoft knowledge base articles.

http://support.microsoft.com/kb/315353

http://support.microsoft.com/kb/315346

"Run Configuration Checker" link available on the CD browser only downloads the Configuration Checker

The "Run Configuration Checker" link available on the CD Browser, enables you to only download the Configuration Checker. To launch the Configuration Checker, you must navigate to the directory path and double-click the setup.exe (2143564)

The installation may fail with "The system cannot find the file specified" error

This issue occurs if the vxinstaller service is in a failed state during the product installation. (2560071)

Workaround: Delete the vxinstaller service and then run the installation wizard again.

Log on to remote nodes before installation

Installation on a remote node may fail if the user does not first log on to the remote node. This situation occurs when using a domain account and the installer to install on a remote machine that has just joined the domain. If the user does not log on to the remote node before installing, the node will be rejected and fail the validation phase of the installation. For remote nodes that join the domain, there is a security requirement that the user must log on to the node at least once before the node can be accessed remotely. (106013)

Uninstallation may fail to remove certain folders

After a successful uninstallation, the product installer may fail to remove the following folders:

- VERITAS Object Bus
- Veritas Shared
- Veritas Volume Manager

These folders contain application logs. The reinstallation of the product will not be affected if these folders are not deleted. (2591541, 2654871)

Workaround: You can safely delete these folders manually.

Error while uninstalling SFW HA after uninstalling VOM

This issue may occur if you uninstall SFW HA after the successful uninstallation of VOM. (2921462)

The following error message is displayed:

```
Windows cannot find 'bin\xprtlc.exe'. Make sure you typed the name
correctly, and then try again.
```

Workaround: You can ignore this error and click OK to proceed with the uninstallation.

Internationalization

The following known issues may be observed when running Storage Foundation for Windows or Storage Foundation HA for Windows in locales other than U.S. English.

Only US-ASCII characters are supported

File paths and names of servers, clusters, disk groups, volumes, databases, directories and files that include non-ASCII characters are not supported by SFW or SFW HA.

You may not be able to view the snapshot history for volumes that include non-ASCII characters. (862762, 860579, 860186, 2426567, 2581502)

Workaround: Only use US-ASCII characters when naming servers, clusters, disk groups, volumes, databases, directories, files and file paths.

Language preference in Veritas Enterprise Administrator (VEA) must be set to English (United States) or Japanese (Japan)

You can set the display language preference for the Veritas Enterprise Administrator (VEA) console by selecting Tools > Preferences. However, after selecting languages other than English (United States) or Japanese (Japan), displayed characters will be corrupted and unreadable even if you have the local language's character set installed in your system and the system's default language is set for your local language. The Japanese (Japan) displays properly only if the SFW Japanese language pack is installed. In Japanese, SFW or SFW HA displays most screens, buttons, and descriptions in Japanese. (849597)

Workaround: Select only English (United States) or Japanese (Japan) as the display language.

Known issues relating to WinLogo certification

For details on the known issues pertaining to WinLogo certification, refer to the following technote:

http://www.symantec.com/docs/DOC4749

General issues

This section provides information about general issues.

Troubleshooting errors

Unexplained errors in the Quick Recovery configuration wizard or Disaster Recovery configuration wizard may be resolved by stopping and then starting the Plugin Host service. Note that Restart does not resolve the issue. (766137)

VMDg resources fault when one of the storage paths is disconnected

This issue occurs when the IBM DSM (Array: IBM DS5020 A/P-C) is installed and configured in an SFW HA environment.

When you disconnect one of the storage paths, VMDg and MountV resources fault and the VCS service groups begin to fail over. This occurs only when SFW is set to use SCSI-3 commands. (2600019)

On Windows operating systems, non-administrator user cannot log on to VEA GUI if UAC is enabled

If User Access Control (UAC) is enabled on Windows Server operating systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.

Veritas Storage Foundation

This section provides information on known Storage Foundation issues.

For volumes under VMNSDg resource, capacity monitoring and automatic volume growth policies do not get available to all cluster nodes

For a volume under a VMNSDg (Volume Manager Non-Shared Diskgroup) resource in a VCS clustered environment, this issue occurs while configuring capacity monitoring or automatic volume growth. During any of the two operations, if you want to make their policies available to another cluster node after a failover, it does not work. However, the policies work on the nodes where they are created. (2932262)

Workaround: To resolve this issue, you need to manually create the same policies on the other cluster nodes as well.

For a dynamic disk group, application component snapshot schedules are not replicated to other nodes in a cluster if created using VSS Snapshot Scheduler Wizard

In a clustered environment, this issue occurs when you create snapshot schedules of an application component for a dynamic disk group using the VSS Snapshot Scheduler Wizard in VEA GUI. In this case, the snapshot schedules do not get replicated to the other nodes in the cluster. However, this issue is not present in case of volume snapshot schedules. (2928909)

Workaround: To resolve this issue, use the Quick Recovery Configuration Wizard to create application component snapshot schedules for dynamic disk groups.

The installer fails to add or remove the product options after upgrading SFW Basic 5.1 to SFW Basic 6.0.1

This issue may occur if you have upgraded from SFW Basic 5.1 to SFW Basic 6.0.1. In this case, if you use the Add Remove program to add or remove the product options, the operation fails. (2925279)

Workaround: After completing the upgrade, delete the following registry key manually:

```
HKLM\SOFTWARE\Wow6432Node\Veritas\VPI\Components\vrts.comp.vm.opt.mpio.dsm\vrts.soln.sfw.server
```

Retry performing the operation using the Add Remove program.

In FoC, if VxSVC is attached to Windows Debugger, it may stop responding when you try to bring offline a service group with VMDg resources

In a Microsoft Failover Cluster (FoC) configuration, the Veritas Enterprise Administrator Service (VxSVC) may stop responding when you try to bring offline a service group with VMDg resources. This happens if VxSVC is attached to Windows Debugger (WinDbg) and there are multiple VMDg resources in a service group. (2807048)

Workaround: There is no workaround for this issue.

After installation of SFW or SFW HA, mirrored and RAID-5 volumes and disk groups cannot be created from LDM

This issue occurs while creating a mirrored or RAID-5 volume or a disk group from Logical Disk Management (LDM) after installing Storage Foundation for Windows (SFW) or Storage Foundation for Windows and High Availability (SFW HA). Because of the presence of Veritas VDS Dynamic Provider (`vxvdsdyn`), the options for creating mirrored and RAID-5 volumes are disabled. Similarly, the

disk group fails to be created on the disks that are removed from the Available Disks list in Microsoft Failover Cluster (FoC). (3030226, 2170857)

Workaround: To resolve this issue, use the Veritas Enterprise Administrator (VEA) GUI to create a mirrored or RAID-5 volume or a disk group, instead of the LDM GUI.

Storage reclamation commands do not work when SFW is run inside Hyper-V virtual machines

This issue is observed on Hyper-V virtual machines where disks that support thin provisioning and reclamation are presented in a pass-through mode. SFW storage reclamation commands run inside a virtual machine appear to succeed, but the provisioned size of the LUNs remains unchanged. Hyper-V filters certain SCSI commands sent from the guest operating systems to the pass-through disks. Refer to the Microsoft Hyper-V documentation here:

http://technet.microsoft.com/en-us/library/dd183729(WS.10).aspx

This issue occurs because SFW uses one of the filtered SCSI commands to request thin storage reclamation. (2611988)

Workaround: In Windows Server operating systems, Hyper-V allows disabling the filtering of SCSI commands. This allows the full SCSI command set to be sent to the pass-through disks mapped to the virtual machine.

Note: Hyper-V does not support disabling filtering of SCSI commands on Windows Server operating systems.

To disable SCSI command filtering, modify the virtual machine configuration and set the **AllowFullSCSICommandSet** property to **True**. Use the Virtualization WMI provider or edit the virtual machine configuration xml file manually. Refer to the Microsoft Hyper-V documentation for more details.

Alternatively, you can also use the following PowerShell script to disable SCSI command filtering for a virtual machine:

```
$HyperVGuest = $args[0]

$VMManagementService = gwmi Msvm_VirtualSystemManagementService
-namespace "root\virtualization"

foreach ($Vm in gwmi Msvm_ComputerSystem
-namespace "root\virtualization" -Filter "elementName='$HyperVGuest'")

{
$SettingData = gwmi -Namespace "root\virtualization"
```

```
-Query "Associators of {$Vm}
Where ResultClass=Msvm_VirtualSystemGlobalSettingData
AssocClass=Msvm_ElementSettingData"
$SettingData.AllowFullSCSICommandSet = $true
$VMManagementService.ModifyVirtualSystem
($Vm,$SettingData.PSBase.GetText(1)) | out-null
}
```

Save this script to a file and run it from the PowerShell command line on the Windows Server Hyper-V host system. The name of the virtual machine must be passed as an argument.

For example, if you save this script to a file named disablescsifiltering.ps1, run this script from the PowerShell command prompt as follows:

```
C:\>.\disablescsifiltering.ps1 virtualmachine_name
```

This script sets the AllowFullSCSICommandSet property value to True.

Note: Before you run this script, you may have to set the PowerShell execution policy to allow execution of unsigned scripts on the local system. Refer to the Windows PowerShell documentation for more information.

First failover attempt might fault for a NativeDisks configuration

The NativeDisks resource might fail to come online on the failover node after first failover. (2857803)

Workaround: Clear the fault and re-attempt the failover.

Failed volume shrink after a successful file system shrink leaves the file system shrunk

During a shrink operation, if a file system shrink is successful but volume shrink fails, then it leaves the file system in a shrunk state while the volume remains of the original size. In such cases, as the sizes of the file system and volume differs, you cannot use the empty region created by the shrink operation on the volume. (2367659)

Workaround: To resolve this issue, you can grow the file system size by using the `vxvol growfs` command. The command grows the file system size approximately equal to volume size. For information about using the command, refer to the *Veritas Storage Foundation™ Administrator's Guide*.

Unknown disk group may be seen after deleting a disk group

This issue occurs while performing the Destroy Dynamic Disk Group operation. In some cases, while performing this operation, an unknown disk group object is displayed. This may happen if the VxVDS service crashes. The unknown disk group consists of disks that originally belonged to the deleted disk group. (2573763)

Workaround: For more information and assistance to resolve this issue, please contact the Symantec Technical Support team.

Wrong information related to disk information is displayed in the Veritas Enterprise Administrator (VEA) console. Single disk is displayed as two disks (harddisk and a missing disk)

When certain operations like create disk group and mirroring is performed on a disk, then it is observed that wrong information is displayed in the **Disk View** on the VEA console. Single disk is displayed as hard disk and a missing disk. (2296423)

Workaround: Perform **vxassist refresh** from CLI or do a refresh from VEA console.

SFW Configuration Utility for Hyper-V Live Migration support wizard shows the hosts as Configured even if any service fails to be configured properly

While configuring cluster nodes using the SFW Configuration Utility for Hyper-V Live Migration support wizard, sometimes a dialog box is displayed with the message "**Please refer to logs for more details**." Additionally, on verifying the cluster node state it is observed that the node state is shown as **Configured** even when the service fails or the subsequent cluster configuration is an invalid configuration. (2571990)

Workaround: Unconfigure and reconfigure the cluster nodes using the SFW Configuration Utility for Hyper-V Live Migration support wizard through the Solutions Configuration Center (SCC).

Refer to *Veritas Storage Foundation™ and Disaster Recovery Solutions for Microsoft Hyper-V*for details.

Though disk partition gets created successfully, a failure message is displayed in the Event Viewer

Even though the disk partition operation is completed successfully, a failure message is displayed in the Event Viewer "**VDS create partition request failed. Investigate further based on operation return status**" along with the message "**Created a new partition on disk Harddisk**." (2293995)

Even when the disk group state is in a Read/Write state, subsequent I/Os keep happening on a failed FoC cluster node

On a Microsoft Failover Cluster (FoC) cluster node, if the service or application group fails over from one node to another node unexpectedly and the VMDGg resource is marked as Failed on the failed node, it is observed that I/Os keep happening on the failed cluster node even when the disk group is in a **Read/Write** state. (2579667)

Workaround: It is recommended to verify the disk group state on the failed cluster node by performing the below-mentioned steps:

To verify the disk group state

1. Run the `vxdg list` command from the command prompt to see the **Access** state of the disk group on the cluster node.
2. If the **Access** state of the disk group is in a **Deported None** state, then no further action is required. However, if the **Access** state of the disk group is in a **Imported Read/Write** state, then it needs to be deported forcefully from the failed node to avoid data corruption.
3. Use the `vxdg deport` command to deport the disk group forcefully from the failed node.

   ```
   vxdg -g<DynamicDiskGroupName> -f deport
   ```

System shutdown or crash of one cluster node and subsequent reboot of other nodes resulting in the SFW messaging for Live Migration support to fail (2509422)

If a cluster node crashes or shuts down abruptly, then it is noticed that on subsequent reboot of the other remaining cluster nodes, the SFW Configuration Utility for Hyper-V Live Migration Support shows the crashed node as **Invalid**Configuration.

In such cases, the following is observed:

- The SFW messaging for Live Migration support will not work between the remaining nodes
- TheVMDg **LiveMigrationSupport** attribute cannot be set to **True** for any new VMDg resource

To resolve this issue, it is recommended to first Unconfigure and then Configure the remaining cluster nodes using the SFW Configuration Utility for Hyper-V Live Migration Support through the Solutions Configuration Center (SCC).

Changing the FastFailover attribute for a VMDg resource from FALSE to TRUE throws an error message

Changing the **FastFailover** attribute for a VMDg resource from **False** to **True** throws an error message. However, the VMDg resource Properties window displays the attribute value as **True**. (2522947)

Workaround: Perform the following to resolve this issue:

- Configure the SFW Hyper-V Live Migration Support using the SFW Configuration Utility for Hyper-V Live Migration Support Wizard through the Solutions Configuration Center (SCC).
- Reset the VMDg resource **FastFailover** attribute to **True**.

Refer to *Veritas Storage Foundation and Disaster Recovery Solutions for Microsoft Hyper-V* for details.

A remote partition is assumed to be on the local node due to Enterprise Valult (EV) DNS alias check

A remote partition is assumed to be on the local node when a DNS alias check is performed for the Enterprise Vault (EV)server. FlashSnap operations fail on such remote partitions. (2572106)

Workaround: Perform or schedule the FlashSnap operations on such partitions from the local host on which they are created.

After performing a restore operation on a COW snapshot, the "Allocated size" shadow storage field value is not getting updated on the VEA console

When restore operation is performed on a COW snapshot, it is observed that the **Allocated size** field value of shadow storage is not getting updated on the Veritas Enterprise Administrator (VEA) console. After performing the `vxassist refresh` operation, the field values are updated and the correct values are displayed on the Veritas Enterprise Administrator (VEA) console. (2275780)

Workaround: Perform the `vxassist refresh` CLI command operation.

Scheduler service does not retain its credentials after upgrading from SFW and SFW HA 5.x to 6.0

After upgrading SFW and SFW HA 5.x to 6.0, the following services do not retain their credentials if they are configured on a Domain in the Active Directory:

- Veritas Scheduler Service

- Veritas Volume Replicator Security Service

You must reconfigure the above-mentioned services in the domain user account after upgrade as user credentials are not preserved after upgrade. (2529295)

Workaround: Manually reset the service credentials after upgrade for all services, if it was set to a different account other than the local system prior to upgrade.

Shrink volume operation may increase provisioned size of volume

Performing a shrink volume operation on a volume that resides on a thin provisioned disk may result in an increase of the provisioned size of the volume. (1935664)

Workaround: There is no workaround for this issue.

Computer crashes if you perform certain operations while the volume shrink is in progress

This issue occurs in Windows Server operating systems because of a bug in the NTFS file system. (2366140, 2400288, 2406659)

If you perform any of the following operations while the online volume shrink operation is in progress, then the computer crashes:

- Deport a disk group
- Fail over a disk group
- Use the Automatic Volume Growth feature
- Remove disks forcefully from the computer or bring them offline

Workaround: This is a known Microsoft issue and there is no workaround for it. To avoid this problem, do not perform any of the mentioned operations when the online volume shrink operation is in progress.

Enterprise Vault (EV) snapshots are displayed in the VEA console log as successful even when the snapshots are skipped and no snapshot is created

When Enterprise Vault component's snapshot is created from the VEA console or using Vxsnap CLI, then the VEA console log displays the snapshot as being created successfully even when no snapshot is created and the snapshot operation fails. This issue is seen when there is a problem with remote communication. (2142378)

Workaround: Make sure that the Scheduler service is configured with Domain admin credentials and restart the Symantec Plug-in Host service if remote communication is not working and delete any previously prepared snapshot mirrors. Perform a **Prepare** operation again and then take a snapshot of the Enterprise Vault components.

BSOD 0xCA during install of DDI-1 when array connected over iSCSI with MPIO

Bluescreening of setup nodes while installing Device Driver Integration (DDI)-1 for Storage Foundation for Windows 5.1 Service Pack 1. (2082546)

Workaround: Reconfigure MS iSCSI initiator without the MPIO option enabled and then install the Device Driver Integration (DDI-1) package.

On a Windows Vista x86 node, the VEA fails to launch with a "MSVCR71.dll not found" error

This issue occurs when the Veritas Enterprise Administrator (VEA) cannot find MSVCR71.dll under C:\Program Files\Veritas\Veritas Object Bus\bin on the node. (2059877, 2132462)

Workaround: To launch the VEA successfully, copy MSVCR71.dll from C:\Program Files\Veritas\Veritas Object Bus\eat\bin to C:\Program Files\Veritas\Veritas Object Bus\bin.

On a clustered setup, split-brain might cause the disks to go into a fail state

In a cluster environment, Split-brain might cause disks and volumes to go into a failing state.

During internal testing it is observed with HP MSAP2000 array that after split-brain, disks are going into a failing state. This is due to some of the SCSI reservation commands taking longer time than expected. (2076136)

Workaround: Reactivate the volumes and disks from the VEA console manually and then online the ServiceGroup in case of a VCS setup and online the ApplicationGroup for Microsoft Failover Cluster (FoC) from the clustered GUI console.

To avoid this in future, increase the value of registration time in the registry. Create a DWORD key with name `RegistrationTimer`' and value should be calculated as below: “If SCSI-3 is enabled from SFW {3 seconds for each disk in the disk group; number of disks in the disk group (DG) => sleeping reservation

for the going to be online DG¦". Value should be specified in Milliseconds. Default value of this key is 7000 (7Secs).

Takeover and Failback operation on Sun Controllers cause disk loss

DSMs report IO errors in case of a takeover and failback operation leaving the volume degraded. This happens both for a standalone setup & clustered setup. (2084811)

FoC disk resource may fail to come online on failover node in case of Node Crash or Storage Disconnect if DMP DSMs are installed

In case of an active node crash or a majority disk loss, Microsoft Failover Cluster (FoC) disk resources may fail to come online on the failover node if DMP DSMs are installed. (2920762)

Workaround: Set the "Clear SCSI reservation" policy on the failover node, and then bring the FoC cluster resource online.

The Veritas Enterprise Administrator (VEA) console cannot remove the Logical Disk Management (LDM) missing disk to basic ones

This issue is intermittently produced and there is no workaround for this issue. (1788281)

After breaking a Logical Disk Manager (LDM) mirror volume through the LDM GUI, LDM shows 2 volumes with the same drive letter

Volume state is properly reflected in Diskpart. Issue is seen only in disk management (diskmgmt) MMC. (1671066)

Workaround:

To reflect the proper volume state in the diskmgmt console

1. If disk management console is open, then close it.
2. Run the command `net stop vxsvc`
3. Stop vds by running `net stop vds`
4. Restart the vxsvc service by running `net start vxsvc`.
5. Now, restart the disk management console .

Unable to failover between cluster nodes. Very slow volume arrival

Slow disk import because of large number of COW snapshots. Significant amount of time is spent in comparing disks. (2104970)

Workaround: Disable the COW processing on disk group Import by creating the following DWORD registry key

`SOFTWARE\\VERITAS\\Vxsvc\\CurrentVersion\\VolumeManager\\DisableCOWOnImport=1`

Blue screening noticed after clussvc crash while creating a VSS snapshot

While taking the VSS backups, the clussvc crashes and system blue screening is noticed. (2118628)

Workaround: Install Microsoft KB at: http://support.microsoft.com/kb/950267

VDS errors noticed in the event viewer log

If a user performs a rescan or reboot operation on the passive node when storage is being mounted on the active node, then the following event is noticed in the event viewer **"Unexpected failure. Error code: AA@02000018"**. Note that this issue does not cause any harm or unexpected behavior. (2123491)

Removing the VMDg resources from a service group to available storage from the Failover Clustering MMC results in an "Assertion Failed" error

When the Volume Manager Disk Group (VMDg) resource is removed from a service group, then the delete operation throws an **"Assertion Failed"** error. Clicking on **Retry** may cause the entire Failover Clustering MMC to hang up. (2080714)

Workaround: Fix for the above issue is to upgrade from Windows Server 2008 SP1 to Windows Server 2008 SP2 as this issue is specific to Windows Server 2008 SP1 only.

Restore with `-a` option for component-based snapshot fails for Exchange mailboxes on a VCS setup

Create a component-based normal snapshot on, say for example on, `C:\ vxsnap -x witha.xml create writer="microsoft exchange writer" component=Final`. Now run the `vxssnap restore` command with the `-f` and `-a` option. Restore operation fails because it tries to restore mounted databases when used with the `-a` option. (1873821)

Workaround: Do either of the following to avoid the restore operation from failing:

- Dismount the mailbox before the restore operation.
  or
- To avoid this problem fully restart the storage agent service once after the service group is online after the initial setup.

An extra GUI Refresh is required to ensure that changes made to the volumes on a cluster disk group having the Volume Manager Disk Group (VMDg) resource gets reflected in the Failover Cluster Manager Console

Workaround: To reflect the changes made to the cluster disk group, go to the Cluster name, right-click on it and perform a Refresh. The changes get reflected in the Failover Manager console.

DR wizard cannot create an RVG that contains more than 32 volumes

For VVR replication, the DR wizard cannot create a Replicated Volume Group (RVG) that contains more than 32 volumes. If you select more than 32 volumes while running the DR wizard, the create RVG task fails when you reach the Implementation panel. (2010918)

To configure DR when any RVG contains more than 32 volumes, use the following steps:

1. Run the DR wizard until the Application Installation panel is displayed and then exit the wizard.
2. Complete the application installation on the secondary nodes.
3. Using the Veritas Enterprise Administrator (VEA) console, create a replicated data set (RDS) with a Primary and Secondary RVG with only 32 volumes. Do not select more than 32 volumes in the create RVG operation.

   For more information on creating an RDS, see the *Veritas Volume Replicator Administrator's Guide*.
4. Once the RVG is created, right click on the RDS name and select **Add Volume**. Using the Add Volume wizard, add the remaining volumes that are part of the RVG.
5. Finish running the DR wizard to complete the service group cloning, replication configuration, and global cluster option (GCO) configuration.
6. In the VEA, change the IP address in the replication settings to match what you entered for the replication IP in the DR wizard, as follows:

- Open the VEA and connect it to a system on the primary site where the primary RVG is configured. From the same VEA GUI, connect to a system on the secondary site where the secondary RVG is configured.
- Go to Replication Network View.
- Right click on the secondary RVG and select **Change Replication Settings**.
- Change the primary side IP address and secondary side IP address to the same values which you provided in the DR wizard on the Replication Attribute Settings panel.

Allow restore using the vxsnap restore command if mailbox is removed or missing

If a mailbox database component is missing or deleted from an Exchange 2010 configuration, then use the `vxsnap restore` command to recover the missing or deleted database component. (2013769)

To restore a missing database, perform the following steps:

1. Create mailbox with same name as that of the missing mailbox database and mention the same database file and log path.

   Make sure to uncheck the **Mount this database** checkbox.

2. Set the database properties by right-clicking the database and enabling the checkbox **This database can be overwritten by a restore** option in the Exchange Management Console.

3. Now try to restore the missing database using the vxsnap restore command.

   ```
   vxsnap -x <filename> [-f] [-b] [-r] [-a] restore
   restoreType=<PIT|POF> writer=<writername>
   [subComponent=<subComponentName>] [RSG=<Yes|No>]
   ```

   Note that the subcomponent and RSG=Yes|No is not valid for Exchange 2010.

Scheduled VSS snapshots of an Exchange mailbox database configured under a VCS cluster setup starts with some delay of around two to three minutes

Prepare and Create VSS snapshot operations starts with some delay and takes sometime to get launched, mostly 2 to 3 minutes. (2021279)

Event viewer shows error message "Could not impersonate Veritas Scheduler Service login user ....." when VSS restore and snapback operations are performed

In case of an Exchange 2010 VCS cluster setup, if the Scheduler Service is stopped then the user may get the following error message **"Could not impersonate Veritas Scheduler Service login user. Make sure this service is started and configured with a domain user account"** in the VEA console and Application event log. (2028835)

Workaround: If a Fileshare resource is configured for a VCS cluster setup, ensure that the Veritas Scheduler Service is running and configured with an appropriate user account. If snapshot metadata files are stored on a local volume, then this error can be ignored.

For a cluster setup, configure the Veritas Scheduler Services with a domain user account

In case of a clustered (VCS or DAG/FoC) setup with more than one node, on each node of the cluster you must configure the Veritas Scheduler Services with a domain user account that has administrative privileges.

Snapshot metadata files are not deleted after VSS Snapback and PIT Restore operation

For an Exchange 2010 VCS cluster setup that has a fileshare resource configured to store snapshot metadata files, it is noticed that after performing a VSS Point in Time (PIT) restore or snapback operation the snapshot metadata files are not deleted even though the restore or snapback operation completes successfully. These snapshot files are displayed in the VSS restore and snapback wizards. (2030283)

Workaround: Manually delete the older snapshot metadata files or if there is requirement to use the same file name, then use the -o option with the vxsnap utility command.

For example if \\FSAP1\share\mb1.xml is the snapshot set file name specified during Create snapshot operation, then files with the names mb1.xml$, mb1.xmlwmd, and mb1.xml should be deleted. They are present at the location \\FSAP1\share.

If an inaccessible path is mentioned in the vxsnap create CLI, the snapshot gets created and the CLI fails

If a wrong path or path which is not accessible is specified during the VSS create snapshot operation, then the operation fails after actual snapshot of the volumes and while generating a snapshot metadata file. (2030292)

Workaround: Perform manual snapback of the volumes which are part of a component and take a VSS snapshot again by specifying a valid and accessible path.

If snapshot set files are stored on a Fileshare path, then they are visible and accessible by all nodes in the VCS cluster

If snapshot metadata files are stored on a fileshare path, they are visible and accessible by all nodes in a VCS cluster. Hence, VSS restore and reattach operations should be performed only on a node where the component is online.

A remote partition is assumed to be on the local node due to Enterprise Valult (EV) DNS alias check

A remote partition is assumed to be on the local node when a DNS alias check is performed for the Enterprise Vault (EV)server. FlashSnap operations fail on such remote partitions. (2572106)

Workaround: Perform or schedule the FlashSnap operations on such partitions from the local host on which they are created.

Storage Management issues

The following are Storage Management issues.

Mirrored volume in Microsoft Disk Management Disk Group does not resynchronize

On Windows Server 2008, a mirrored volume in a Microsoft Disk Management Disk Group does not resynchronize when a failed mirror is reattached. (1150292)

Workaround: Reactivate the disk and resynchronize the volume using Microsoft Disk Management.

Expand volume operation not supported for certain types of volumes created by Microsoft Disk Management

The resize operation to expand a volume created by Microsoft Disk Management is not supported for mirror, stripe, or RAID-5 volumes. Also, extending a volume to more than one disk in a single operation is not supported. A volume can only be extended on one other disk during a resize operation. However, the resize

operation can be repeated so that the volume can be extended to more than one disk. (1128016)

Snapshot and Restore issues

The following are Snapshot and Restore issues.

Vxsnap restore CLI command fails when specifying a full path name for a volume

Specifying a full path name for a volume in the `vxsnap restore` CLI command fails with an error message, **"The volume is not present in the snapshot."** (1897541)

Workaround: Specify either the drive letter or the drive path of the volume in the `vxsnap restore` command instead of specifying the full path name of the volume.

Restoring COW snapshots causes earlier COW snapshots to be deleted

On Windows Server operating systems, when restoring all earlier COW snapshots in reverse chronological order (restoring the latest snapshot to the earliest snapshot) causes earlier COW snapshots to be deleted. These COW snapshots are deleted after the second COW snapshot is restored. (1864268)

Workaround: This is a known Microsoft problem. Refer to Microsoft KB975803 for more information.

http://support.microsoft.com/default.aspx?scid=kb;EN-US;975803

COW restore wizard does not update selected volumes

The COW restore wizard requires that the snapshot set (XML file) be specified for the restore operation. The specification of the snapshot set allows the wizard to display the volumes associated with the snapshot set.

When you specify the snapshot set, continue to view the volumes to restore, and then go back to specify a different snapshot set, the volumes associated with the new snapshot set are not displayed in the **Select Volumes** screen of the wizard. The volumes that are displayed are the volumes associated with the first snapshot set. (1881148)

Workaround: Cancel the COW restore wizard and launch it again specifying the appropriate snapshot set.

Snapshot operation requires additional time

On Windows Server operating systems, creating a new snapshot volume by performing a snapshot operation (mirror break) on a volume that already has a COW snapshot volume, and then performing an operation on this snapshot volume

(e.g. assigning a drive letter, restore, or a snapshot operation that assigns a drive letter) requires additional time to complete.

Subsequent operations on the snapshot volume do not require additional time. (1872810)

Incorrect message displayed when wrong target is specified in vxsnap diffarea command

Issuing the `vxsnap diffarea -c` CLI command with the wrong value for the target parameter results in the display of an incorrect error message in the VEA console and in the Windows Event Viewer. The incorrect message that is displayed is "Failed to remove shadow storage area". The correct message that should be displayed is "Failed to change shadow storage area".

However, the correct message is displayed in the CLI command window. (1879829)

Restore operation specifying missing volume for SQL component fails

The operation to restore an SQL component specifying a missing volume fails when the operation has completed and the drive letter of the restored volume is changed to the drive letter of the original volume. (1876307)

Workaround: Change the drive letter of the snapshot volume to the drive letter of the original volume before starting the restore operation.

Snapshot operation of remote Sharepoint database fails when it resides on local SharePoint server

After configuring a remote database with a separate machine name and IP address on the local SharePoint server, taking a snapshot of the database fails.

This situation creates a call-back loop and returns the error condition, snapshot operation already in progress. (1847861)

Snapshot of Microsoft Hyper-V virtual machine results in deported disk group on Hyper-V guest

Creating a dynamic disk group with SCSI disks on a Hyper-V guest machine and then taking a snapshot of the Hyper-V guest with the Hyper-V host causes the disk group to be deported. (1859745)

Enterprise Vault restore operation fails for remote components

The restore operation fails for an Enterprise Vault component, when a part of the component resides on the local server and a part resides on a remote server. An open handle may exist on a volume where one of the parts reside causing the operation to fail. (1729872)

Workaround: Specify the Force option in the Enterprise Vault restore wizard or CLI command to allow the operation to proceed successfully.

Persistent shadow copies are not supported for FAT and FAT32 volumes

A shadow copy is persistent when it is not deleted after a backup operation. A persistent shadow copy is only supported for NTFS volumes. They are not supported for FAT or FAT32 volumes. (1779879)

This is a known Microsoft problem. Refer to Microsoft technical support for more information about this problem.

http://msdn.microsoft.com/en-us/library/aa384613(VS.85).aspx

Copy On Write (COW) snapshots are automatically deleted after shrink volume operation

On Windows Server operating systems, the operation to shrink a volume that contains a shadow storage area causes VSS to delete any shadow copies (COW snapshots) that reside on the volume. (1863910)

Shadow storage settings for a Copy On Write (COW) snapshot persist after shrinking target volume

On Windows Server operating systems, the shadow storage (DiffArea) setting for the size of the target volume does not change after shrinking the size of the target volume to less than minimum size. The DiffArea settings for the size of the target volume reflect the DiffArea size of the target volume before the shrink operation. (1592758)

Copy On Write (COW) shadow storage settings for a volume persist on newly created volume after breaking its snapshot mirror

On Windows Server operating systems, the shadow storage (DiffArea) settings for a volume are applied to the newly created volume after breaking the snapshot mirror. These shadow storage settings can be displayed with the `vxsnap refresh` CLI command. (1678813)

Conflict occurs when VSS snapshot schedules or VSS snapshots have identical snapshot set names

An XML file is created when a VSS snapshot is taken. This XML file contains database and snapshot volume metadata. If two snapshot schedules, or a snapshot schedule and a VSS snapshot, are created with the identical snapshot set name and directory path, the schedule that is launched later overwrites the XML file that was created by the schedule or VSS snapshot operation that was launched earlier.

Since the earlier XML file does not exist, subsequent VSS reattach/VSS restore operations for that schedule or snapshot fails. (1303549)

Workaround: Ensure that snapshot set names are unique in a given directory path to avoid conflict with other VSS snapshot schedules or VSS snapshots.

Memory leak occurs during snapshot and snapback operations

A memory leak occurs during snapshot and snapback operations when the Microsoft Virtual Disk Service (VDS) is called. VDS causes the memory leak.

This is a known Microsoft problem. Refer to Microsoft technical support for more information about this problem. (1234278)

Microsoft Outlook 2007 Client (caching mode enabled) does not display restore messages after VSS Exchange restore operation completes

After a VSS Exchange restore operation completes, restore messages are not displayed in the Outlook 2007 Client when caching is enabled.

For more information about this issue, refer to Microsoft Outlook 2007 technical support. (1287199)

Volume information not displayed correctly in VSS Restore wizard

If a subcomponent of Microsoft Exchange is configured to use more than one volume, then the last page of the VSS Restore wizard does not display the list of volumes correctly. This is only a display issue and does not affect the restore operation. (1179162)

Vxsnap restore operation fails with "Pre-Restore failed by Writer" error

SFW dismounts the Exchange 2007 stores before beginning the vxsnap restore operation. If it fails to dismount the stores, the restore operation fails with a "Pre-Restore failed by Writer" error.

This occurs when the Exchange Storage group is not offline/dismounted or when databases have not been set to overwrite by restore. (1253095)

Workaround: Make sure to dismount the stores, manually set them to overwrite, and repeat the vxsnap restore operation.

VSS Writers cannot be refreshed or contacted

VSS Writers cannot be refreshed or contacted as in the following:

- `Vxsnap refresh` CLI operation fails because VSS fails to gather data from the VSS Writers
- Windows Event Viewer encounters a VSS error, "An internal inconsistency was detected in trying to contact shadow copy service writer." (Event ID 12302)

These are known Microsoft problems. (1275029)

Workaround: Refer to Microsoft KB940184 for steps to correct the issue.

Memory leaks occur during VSS snapshot operations

During snapshot operations using VSS, memory leaks occur.

This is a known Microsoft problem (KB 933653). (859058, 1239666)

Time-out errors may occur in Volume Shadow Copy Service (VSS) writers and result in snapshots that are not VSS compliant

In some circumstances, you may receive VSS errors showing that the volume shadow copy freeze timed out. As a result the snapshots that were created are not VSS compliant and the snapshot XML file used by the VSS-based wizards and `vxsnap` commands is not generated. Therefore, you cannot use any of the `vxsnap` commands or VSS-based wizards to restore or to reattach the snapshot. If the snapshot volumes have been scheduled for automatic updates with the Quick Recovery Configuration Wizard or VSS Snapshot Scheduler Wizard, the updates cannot occur. (633219)

For a detailed description of the problem, see:

http://support.microsoft.com/kb/915331

Workaround: If a snapshot fails with this error, you can use volume-based commands to manually snapback individual snapshot volumes. You can use the `vxassist snapback` command or the Snap Back command from the Volumes node in the Veritas Enterprise Administrator console. Once the volumes are reattached and resynchronization is complete, you can create a new snapshot manually or scheduled snapshots can resume.

In addition, Microsoft supplies a hotfix that you can install to resolve this issue. For additional information, see Microsoft Knowledge Base 915331:

The backup process may fail and a time-out error may occur in Volume Shadow Copy Service writers

http://support.microsoft.com/default.aspx?scid=kb%3Ben-us%3B915331

The vxsnapsql restore CLI command may fail when restoring an SQL database

On an SFW HA system that is configured with VCS, VVR, and GCO options, using the `vxsnapsql restore` CLI command to restore a SQL database may fail with the following error message: (895239)

```
Recovering production volumes from Snapshot Backup set ...
Can not reattach a mirror to a volume that is in use by another
```

```
application. Please close applications, consoles, Explorer windows,
or third-party system management tools accessing the volume and
then retry the operation. The SQL command failed after it was
initiated. The operation failed.
```

Workaround: The workaround for this problem is to first offline all the SQL server and MountV resources for the volume which contains the SQL database and Logs on VCS and then to bring them back online.

The `vxsnapsql restore CLI` command works correctly after performing this procedure.

VSS objects may not display correctly in VEA

On systems running both SFW and Microsoft Exchange, VSS objects may not be displayed in VEA after a reboot. Also, VSS objects may not display correctly as a result of changes to storage groups or databases in Exchange. (307402)

Workaround: Select **Refresh** from the **Action** menu of the VEA menu bar (or use the `vxsnap refresh` CLI command). Refreshing VEA displays these VSS objects.

VSS Snapshot of a volume fails after restarting the VSS provider service

The Veritas VSS Provider Service contacts the Microsoft VSS service to complete the snapshot operation. Restarting the Veritas VSS Provider Service disables the contact to the Microsoft VSS service. (352700)

Workaround: Restart Microsoft VSS service after restarting the Veritas VSS Provider Service.

Restoring SQL databases mounted on the same volume

When you restore a Microsoft SQL database that resides on a volume that contains another SQL database, the vxsnapsql utility restores both databases. (258315)

Workaround: Avoid this situation by configuring each SQL database on its own separate dynamic volume.

Mirror attach operation hangs and does not complete

The mirror reattach operation may not finish and hangs at 99% complete. Although the operation appears not to finish, the volume is healthy and it is accessible. (406420)

Workaround: The workaround is to issue a rescan to signal the completion of the operation.

Snapshot operation fails if components with the same name exist in different Exchange virtual servers

If multiple Exchange virtual servers are online on the same server, snapshot operations may fail. This can occur when using the `vxsnap start` and `vxsnap create` commands or the Quick Recovery Configuration Wizard. (508893, 1104325)

Workaround: Use the **VSS Snapshot Wizard** (VEA GUI) to take a snapshot in an environment with two virtual Exchange servers, when both have a storage group with the same name. To use the Quick Recovery Configuration Wizard, rename the storage groups with the same name.

To rename the service group, in the Exchange Management Console, right-click the storage group that you want to rename and click **Properties**. In General Properties, change the name in the editable box and click **Apply**.

CLI command, vxsnap prepare, does not create snapshot mirrors in a stripe layout

When using the `vxsnap prepare` command, specifying the layout type as stripe should create snapshot mirrors in a stripe layout. However, if the number of columns is not also specified in the `vxsnap prepare` command, then snapshot mirrors with a concatenated layout are created. (839241)

After taking a snapshot of a volume, the resize option of the snapshot is disabled

After performing a snapshot operation on a volume, the volume might be designated as read-only, which means the Resize Volume option is disabled. (Right-click the volume in tree view and in the menu, Resize Volume... is disabled). (866310)

Workaround: In the volume properties page, deselect the **Read Only** check box. When you right-click the volume in tree view, **Resize Volume** > **Expand** is now enabled.

If the snapshot plex and original plex are of different sizes, the snapback fails

When a snapshot volume and the original volume are of different sizes, the snapback fails. (867677)

Workaround: Make the snapshot volume read-write manually, increase the size of the snapshot volume to match the size of the corresponding original volume, and then reattach.

Snapshot scheduling issues

The following are Snapshot scheduling issues.

Snapshot schedule fails as result of reattach operation error

On Windows Server operating systems, a snapshot schedule fails when the reattach operation fails during a snapshot procedure on mounted volumes. A **"volumes are in use, cannot reattach"** error occurs for the reattach operation. Subsequent snapshot schedules fail with the same error. The reattach operation fails as a result of a known Microsoft volume lock problem (SRX080317601931). (1280848)

Workaround: Snapshotted volumes that do not have assigned drive letters do not encounter this error. When creating snapshot schedules, select the "no driveletter" for the snapshotted volumes.

Next run date information of snapshot schedule does not get updated automatically

When selecting a snapshot schedule object in the VEA GUI , information about the next run date is displayed.

If the next run date changes, such as after a scheduled run, the new next run date information is not automatically updated in the VEA GUI. (930269)

Workaround: Reselecting the snapshot schedule in the VEA GUI updates the display of the next run date information.

Changes related to Daylight Savings

Time Impact of the United States Energy Policy Act of 2005. Beginning Spring 2007, daylight saving time (DST) start and end dates have been changed. DST dates in the United States:

- Start three weeks earlier (2:00 A.M. on the second Sunday in March)
- End one week later (2:00 A.M. on the first Sunday in November).

To address this change, Microsoft provides patches for Windows servers.

These patches are available at:

http://support.microsoft.com/kb/928388

The SFW Snapshot Scheduler relies on the Windows system clock and does not function correctly without the application of the Microsoft DST patch.

VEA GUI may not display correct snapshot schedule information after Veritas Scheduler Service configuration update

In a cluster environment, the Veritas Scheduler Service needs to be configured on each node with domain administrator privileges. This configuration change requires that the scheduler service be restarted on each node to enable the new settings. This is done to ensure that the schedule information is reflected on all

the nodes in the cluster in case of failover. However, the VEA GUI may not show the correct schedule information after the service is restarted. (1260683)

Workaround: To ensure that the VEA GUI displays the correct schedule information, the Storage Agent Service also needs to be restarted after the Scheduler Service is restarted. In this way, the Storage Agent Service is able to receive any changes in the schedule information from the Veritas Scheduler Service. Alternatively, to get the correct schedule information, you must perform a VSS refresh command with the VEA GUI or a `vxsnap refresh` CLI command every time you want to display the correct schedule information.

Scheduled snapshots affected by transition to Daylight Savings Time

The transition from Standard Time to Daylight Savings Time (DST) and the transition from Daylight Savings Time to Standard Time affects the Snapshot Scheduler. (929625)

- On the first day of DST, any snapshots scheduled during 2:00 A.M.- 2:59 A.M. are taken during 3:00 A.M.- 3:59 A.M. DST.
- On the last day of DST, any snapshots scheduled during 1:00 A.M. - 1:59 A.M. are taken 1:00 A.M. - 1:59 A.M. Standard Time.
- If during 1:00 A.M. - 1:59 A.M. on the last day of DST the Veritas Scheduler Service is started/restarted or a VSS refresh occurs, some snapshots scheduled for this period are not taken. For example, if a VSS refresh occurs at 1:30 A.M. on the last day of DST, then any snapshots scheduled during 1:00 A.M. - 1:29 A.M.are not taken.

In a cluster environment, the scheduled snapshot configuration succeeds on the active node but fails on another cluster node

In a VCS cluster environment, in some cases configuring a snapshot schedule fails on one or more of the cluster nodes and the Quick RecoveryWizard or VSS Snapshot Scheduler Wizard displays an error message to that effect. In that case, the schedule succeeds on the active node but in the case of a failover, scheduled snapshots do not occur. (800772)

Workaround: Start the Quick Recovery Configuration Wizard from the Solutions Configuration Center (**Start**>**Run**>**scc**). Continue through the wizard until the **Synchronizing Schedules** panel shows that synchronization between cluster nodes is complete. Click **Finish** to exit the wizard.

After a failover occurs, a snapshot operation scheduled within two minutes of the failover does not occur

When a failover occurs and the disk group is imported on the active node, the scheduler waits for two minutes. Then the schedule-related information is

refreshed. If a snapshot operation, such as a mirror preparation or a snapshot, is scheduled within those two minutes, it does not occur at that time. The schedule starts working with the next scheduled snapshot operation. If the mirror preparation operation was skipped, it is performed at the time of the next scheduled snapshot. (798628)

Unable to create or delete schedules on an FoC cluster node while another cluster node is shutting down

If you are creating or deleting a snapshot schedule on a Microsoft Failover Cluster (FoC) cluster node while another node in the cluster is shutting down, the schedule creation or deletion fails. You can no longer create or delete schedules on the original node until the Veritas Storage Agent (`vxvm` service) is restarted on the original node. However, any existing schedules continue to run, and you can create or delete schedules from other nodes in the cluster. (894830)

Workaround: Restart the Veritas Storage Agent (`vxvm` service) on the node on which you attempted to create or delete the schedule.

Quick Recovery Wizard schedules are not executed if service group fails over to secondary zone in a replicated data cluster

In a replicated data cluster configured with primary and secondary zones, Quick Recovery snapshot schedules are not executed if the service group fails over from the primary zone to the secondary zone. (1209197)

On Windows Server, a scheduled snapshot operation may fail due to mounted volumes being locked by the OS

A Windows Server operating system issue causes the operating system to intermittently lock mounted volumes. This can result in a failure in a scheduled snapshot operation, if the user specified mount points or mount paths for the snapshot volumes or manually mounted the snapshot volumes after a snapshot operation completed. If the operating system locks mounted volumes, when the scheduler tries to do the next scheduled operation, it fails with the error "volumes are in use". The error can be found in the .sts file corresponding to the schedule. (1205743)

Workaround: Check if any programs or processes are holding a lock on the storage groups and take the necessary steps to release the lock on the relevant volumes. Remove the mount for the volume before the next scheduled snapshot.

Quick Recovery Configuration Wizard issues

The following are Quick Recovery Configuration Wizard issues.

Quick Recovery Wizard allows identical names to be assigned to snapshot sets for different databases

The Quick Recovery Configuration Wizard lets you edit the snapshot set names and XML file names. If you select multiple databases during one run of the wizard, the wizard validates the names you assign to ensure that they are unique across all databases and snapshot sets. However, if you specify different databases during different runs of the wizard, the wizard is unable to validate that the names assigned during the later run are different from the names assigned earlier. If you later run the wizard to modify both databases at the same time, the wizard recognizes the names are the same and will not proceed further. (1090276)

Workaround: Select both databases in a single run of the wizard when configuring for the first time, so that the wizard can validate the names, or ensure that you specify unique names. If you have already assigned the same names by running the wizard multiple times for multiple databases, select the databases on different runs in modify mode as well.

VEA Console issues

The following are VEA Console issues.

VEA GUI sometimes does not show all the EV components

In some cases, this issue may occur while viewing the Enterprise Vault (EV) components in the Veritas Enterprise Administrator (VEA) GUI. The VEA GUI does not display some of the EV components because the GUI is not refreshed properly. However, this does not have any impact on the functionality of the product. (2846344)

Workaround: To resolve this issue, refresh the VEA GUI using the Refresh command.

VEA GUI incorrectly shows yellow caution symbol on the disk icon

This issue occurs when the snapshot or snapback operations are performed multiple times in quick succession, either by creating schedules using the Quick Recovery Configuration Wizard or performed manually using the `vxsnap` command. The Veritas Enterprise Administrator (VEA) GUI incorrectly shows the yellow caution symbol on a disk's icon because VEA GUI has not updated the recent changes. However, this does not have any impact on the functionality of SFW. (2879200)

Workaround: Perform the Refresh command to resolve this issue.

Reclaim storage space operation may not update progress in GUI

Performing a reclaim operation may not allow the GUI to automatically update the progress of the operation. In this situation, the progress of the operation does not change. (1955322)

Workaround: Perform a rescan operation to allow SFW to obtain the progress about the operation and to refresh the GUI.

VEA GUI fails to log on to iSCSI target

On a Windows Server operating systems, the operation to log onto an iSCSI target fails when selecting the initiator adapter and the source portal (using the "Advanced settings" option). The failure of the operation is not obvious. However the connection object displayed in the VEA GUI for the logon session shows an invalid IP address of 0.0.0.0. (1287942)

Workaround: When it is necessary to specify the initiator adapter and source portal during logon of an iSCSI target, you can use the Microsoft iSCSI Initiator Applet to successfully perform the operation.

VEA GUI incorrectly displays a new iSCSI disk as online

On Windows Server operating systems, when a new iSCSI disk is presented, the VEA GUI incorrectly displays the disk as being online. (1362395)

Workaround: Perform a rescan to correctly display the new iSCSI disk as being offline.

VEA does not display properly when Windows color scheme is set to High Contrast Black

Launching the VEA GUI and then changing the color scheme in the Appearance settings of Windows to High Contrast Black causes the VEA GUI not to display properly. (1225988)

Workaround: To enable the VEA GUI to display properly, close the VEA GUI and launch it again.

VEA displays objects incorrectly after Online/Offline disk operations

On Windows Server operating systems, after performing online/offline disk operations on disks that belong to the Microsoft Disk Management Disk Group, the VEA GUI may display objects related to this disk group incorrectly. Missing disk or duplicated volume objects may be displayed in the VEA GUI. Generally, performing a rescan operation corrects this issue. However, a rescan may not be effective and may possibly cause the Veritas Storage Agent Service to terminate abnormally. This situation may also occur when the dynamic disk in the Microsoft

Disk Management Disk Group is disabled and then enabled with the Device Manager. (1196813, 1200302, 1202847, 1204590, 1205352)

Workaround: To have the VEA GUI display objects related to the Microsoft Disk Management Disk Group correctly, restart the Storage Agent Service. However, after the Storage Agent has restarted, performing some operations on the disk group (such as write signature or create simple volume) using SFW may fail. In this situation, perform a rescan operation after the Storage Agent Service has restarted.

Disks displayed in Unknown disk group after system reboot

If all disks in a dynamic disk group are brought online after a server is booted, the disks are incorrectly displayed in the Unknown disk group. (1138080)

Workaround: Perform a rescan to display the disk group correctly.

Device type displayed for a disk may not be accurate

The device type displayed for a disk in the VEA may not be accurate.

When the device type is displayed as FIBRE for a disk, the device type may be a different type, such as SCSI. SFW obtains the device type value from a Microsoft API. This issue has been sent to Microsoft for investigation. (291887)

Internationalization issues

The following are Internationalization issues.

VEA GUI cannot show double-byte characters correctly on (English) Windows operating system

VEA GUI relies on the font setting of the Windows operating system to be enabled to display double-byte characters after enabling East Asian Languages in the **Windows Regional and Language Options** dialog box. The default font setting for the (English) Windows operating system cannot display double-byte characters. (1238207)

Workaround: The following procedures enable the display of double-byte characters.

To enable display in Windows XP

1. Right click on the desktop.
2. Select **Properties** > **Appearances** tab.
3. In the window that appears, click **Advanced**.
4. Select "Message Box" from the Item drop-down list.

5 Select a font from the Font drop-down list that supports double-byte characters. (For example: "MS Mincho".) .

6 Click **OK** to complete the setting.

To enable display in Windows Server 2008 or Windows Vista

1 Right click on the desktop.

2 Select **Personalize**.

3 Select Windows Color and Appearance.

4 In the window that appears, click **Advanced**.

5 Select "Message Box" from the Item drop-down list.

6 Select a font from the Font drop-down list that supports double-byte characters. (For example: "MS Mincho".)

7 Click **OK** to complete the setting.

VEA can't connect to the remote VEA server on non-English platforms

When connecting to the remote VEA server on non-English platforms, you might see a VEA error that says **"Request to server has timed out"**. (804330, 861289)

Workaround: Set up the target server's subnet in the DNS Reverse Lookup Zone. For example, if the remote VEA server is 10.198.91.111, set the target server's subnet to 10.198.91.* in the DNS Reverse Lookup Zone.

Note that setting the DNS Reverse Lookup Zone Configuration is a network requirement for VEA and VVR. When setting up your network, verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported. Make sure that a reverse lookup zone exists in the DNS.

Dynamic Multi-pathing (DMP) issues

The following are Dynamic Multi-pathing (DMP) issues.

Bug check may occur when adding DMP DSM option

After installing SFW, adding the DMP DSM option, with Windows Add or Remove Programs, may result in bug check 0xD1. This issue has been reported to Microsoft (SRZ080421000462). (1251851)

Changes made to a multipathing policy of a LUN using the Microsoft Disk Management console, do not appear on the VEA GUI

DMP DSMs do not manage the load balance settings made with the Microsoft Disk Management console. So changes made to a multipathing policy using the Microsoft Disk Management console do not appear on the VEA GUI.

Changing the load balance settings for DMP DSMs must be done using the SFW VEA GUI or CLI. (1859745)

VEA or CLI operations for DMP DSMs fail without providing error message if WMI service is disabled

The Windows Management Instrumentation (WMI) service is required for using the DMP DSM feature. If you disable the WMI service, the wizards or commands for DMP DSM operations that require the WMI service will fail. The message window displays only an error code without a message explaining the cause of the failure. (2590359)

Microsoft Systems Center Operations Manager 2007 (OpsMgr 2007) issues

The following are Microsoft Systems Center Operations Manager 2007 (OpsMgr 2007) issues.

Performance Graph for MPIO does not display data

When monitoring the performance activity of a volume using MPIO Path Performance counters, the performance data is displayed as the number of reads, number of writes, bytes read, and bytes written. The counter can increment to a point where the graph appears to level off or becomes negative. This condition continues until the affected counter rolls over to zero, at which time an accurate graph is displayed. (914312)

When deleting the last RVG or moving an RVG, the VVR state view is not updated

The **VVR State** view > **Detail** view is not updated in OpsMgr 2007 when the last RVG is deleted or when an RVG is moved. This is due to OpsMgr 2007 being unable to recognize the empty collection of item sent by discovery workflows. (1051217, 1051220)

Other issues

The following are other issues.

SFW volume operations fail on Windows Server

On Windows Server operating systems, SFW operations on a volume may fail with a message stating that the volume is already in use. This is caused by Windows not releasing an open handle. This is a known Microsoft problem. (KB952790) (1093454)

Sharing property of folders not persistent after system reboot

On Windows Server operating systems, folders that reside on a volume in a dynamic disk group and were set up as shared folders are no longer shared after a system reboot. The following is the workaround procedure for this issue. (1856737)

Note: Perform the following before system reboot.

To work around the sharing property issue

1. Enable the dynamic disk group that contains the shared folders for latestart.

   For example use the CLI command:

   ```
   vxdg -gDiskGroup1 latestart on
   ```

2. In regedit, navigate to
   `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\lanmanserver`
3. Right-click the lanmanserver node and select **New** > **Multi-String Value** to enter a new REG_MULTI_SZ entry.
4. Name the Multi-String Value as DependOnService and enter the service name for the Veritas DG Delayed Import Server name in the Data field. (The default name for this service is VxDgDI.)
5. Reboot the system.

Microsoft Disk Management console displays an error when a basic disk is encapsulated

On a Windows Server operating systems with the Microsoft Disk Management console launched, adding a basic disk that contains a primary partition/extended partition with a logical drive to an SFW dynamic disk group using the VEA GUI, may cause a pop-up error message on the Microsoft Disk Management console. The pop-up error message on the Microsoft Disk Management console is not meaningful and can be ignored. (1601134)

Results of a disk group split query on disks that contain a shadow storage area may not report the complete set of disks

When performing a disk group split query command on a set of disks that contain a shadow storage area of volumes on disks having mirrored volumes, the resulting report may not be comprehensive. In this case, the report does not indicate the complete set of disks for split closure. (1797049)

Extending a simple volume in Microsoft Disk Management Disk Group fails

On Windows Server operating systems, while extending a simple volume in a Microsoft Disk Management disk group, the operation fails with the error message, **"Fail to grow volume"**. This issue also affects the automatic volume growth operation when resizing a volume in a Microsoft Disk Management disk group. This is a known Microsoft problem (KB975680) (1596070, 1834611).

Unable to add a shared folder in Microsoft Failover Cluster environment

After creating a cluster disk group and VMDG resource (dynamic volume) in a Microsoft Failover Cluster environment on Windows Server operating systems, the dynamic volume cannot be found when trying to add a shared folder with the "Provision a shared folder Wizard". This is a problem with Windows Server and has been reported to Microsoft. (1233387)

Workaround: Moving the VMDG resource to each of the nodes in the cluster environment, one node at a time, allows adding the dynamic volume as a shared folder with the "Provision a shared folder Wizard" in the last node.

After the VMDG resource reaches the last node, the "Provision a shared folder Wizard" works correctly when the VMDG resource is moved back to any of the other nodes.

Volume automatically assigned a drive letter after dynamic disk group deport/import operations

On Windows Server operating systems, the operations of deporting and then importing a dynamic disk group that contains a volume that does not have an assigned drive letter results in the assignment of a drive letter to the volume. In addition, the drive letters of other volumes in the dynamic disk group may change. This is a Windows Server problem and has been reported to Microsoft. (1282891)

Workaround: Manually remove the automatically assigned drive letter of the volume after importing the dynamic disk group. Also adjust the drive letters of the other volumes in the dynamic disk group as needed.

SFW cannot merge recovered disk back to RAID5 volume

For a Microsoft Disk Management RAID5 volume on Windows Server operating systems, a recovered disk is displayed by SFW as a RAID 5 volume, however the volume has a degraded status. SFW is not enabled to perform a reactivate operation on the volume to change the volume to a healthy status. (1150262)

Workaround: Use Microsoft Disk Management to reactivate the disk or the RAIDS 5 volume to resynchronize the plexes in the RAID5 volume and change the volume to a healthy status.

Request for format volume occurs when importing dynamic disk group

During the import of a dynamic disk group, or other operation that involves mounting a volume, that has an unformatted volume with a drive letter or assigned mount point, a pop-up window appears that requests formatting the volume. Avoid completing the formatting operation if there is any existing data on the volume. (1109663)

Logging on to SFW as a member of the Windows Administrator group requires additional credentials

On Windows Server operating systems, by design, logging on to SFW as a member of the Windows Administrator group should allow access to SFW without additional credentials. However, only the Administrator userid is allowed access to SFW in this way. Other members of the Administrator group are not allowed access unless additional credentials are given. (1233589)

Workaround: Other members of the Administrator group should provide their Windows userid and password when prompted to gain access to SFW.

Removing a cluster disk group causes an error message

On Windows Server operating systems, removing a cluster disk group causes an error message describing that an Assertion Failed. However, in this case, the remove cluster disk group operation has completed successfully and this message can be ignored. Click the **Ignore** button at the bottom of the message window to continue. This is a known Microsoft problem (KB953402) (1233398).

Certain operations on a dynamic volume cause a warning

On Windows Server operating systems, operations on a dynamic volume (such as change drive letter, delete, or shrink) result in a warning message stating that the volume is currently in use. This is a known Microsoft volume lock problem (SRX080317601931) (1093454).

Workaround: If no applications are utilizing the volume, complete the operation by responding to the warning message to perform the operation with force.

Avoid encapsulating a disk that contains a system-critical basic volume

On Windows Server operating systems, if a disk contains a system-critical basic volume (as determined by VSS), then the disk should not encapsulated by SFW. The disk needs to be managed by Microsoft Logical Disk Manager (LDM) so that in a recovery situation it can be recovered by ASR. Encapsulating the disk would not allow recovery by ASR. (1180702)

Sharing property of folders in clustering environment is not persistent

In a clustering environment on Windows Server operating systems, the sharing property of folders is not persistent when first the cluster disk group is deported and the system is rebooted, and then the cluster disk group is imported back to the system. Also, the sharing property is not persistent when the cluster disk group is deported to another node. In addition, the file share property of a volume is not persistent when it arrives after system boot up. (1195732)

Access violation error occurs when performing simultaneous rescan operations

On a Windows Server operating system, performing two rescan operations simultaneously, one in SFW and one in Microsoft Disk Management, results in an access violation error. (1219999)

Workaround: Ensure that the Veritas Storage Agent service and the Virtual Disk Service (VDS) are both stopped; and then restart both services.

Fileshare cannot be created using Failover Cluster Management

In a clustering environment on Windows Server operating systems, a file share cannot be created by using Microsoft Failover Cluster Management. An alternate way to create the file share is to use Windows Explorer. In addition, connecting to the file share using the virtual IP address is not possible.

To connect to the fileshare, use the virtual name of the fileshare, not the virtual IP address. (1159620, 1195732)

Installation of SFW or SFW HA into a non-default installation directory results in the creation of duplicate files and directories

If you choose to specify an installation directory instead of accepting the default directory, duplicate files, and directories are created. This does not affect the function of the product.(861852)

Entries under Task Tab may not be displayed with the correct name

Tasks displayed under the Task tab of the VEA GUI console may appear as an entry labeled as "NoName". These labels are not harmful and refer to a task that is running. (797332)

Attempting to add a gatekeeper device to a dynamic disk group can cause problems with subsequent operations on that disk group until the storage agent is restarted

If your storage array has a gatekeeper device (disk), do not add this disk to a dynamic disk group. The operation to include this disk in a dynamic disk group

fails, and subsequent operations on the disk group, such as snapshot operations, fail until the storage agent is restarted. (864031)

Workaround: Remove any gatekeeper devices from the dynamic disk group and restart the Veritas Storage Agent (vxvm service).

Installing SFW in a non-default path causes an abnormal termination

An abnormal termination occurs when installing SFW in a location that is not the default installation path. This is due to a problem with Microsoft Virtual Disk Service (VDS). This is a known Microsoft problem (SRX061018602975).(829850)

ASR fails to restore a disk group that has a missing disk

When a disk group is missing a disk or a volume, you should not perform an ASR backup and restore procedure, as that action is not supported.(844084)

Use only U.S. ASCII characters in the SFW or SFW HA installation directory name

Using non-ASCII characters in the SFW or SFW HA installation directory may result in the creation of duplicate directories and files. (858913)

Workaround: No workaround. Use only U.S. ASCII characters in directory names.

Unable to create an FoC Volume Manager Disk Group resource type on the active node

In a two node Microsoft Failover Cluster (FoC) cluster, you cannot create an FoC Volume Manager Disk Group resource type on the active node after SFW has been uninstalled on the standby node.

This issue occurs when the Volume Manager Disk Group (VMDg) resource type does not already exist in the cluster before uninstalling SFW on the standby node. (301263)

Workaround: The workaround is to run ClusReg.cmd on the active node after uninstalling SFW on the standby node and before trying to create the VMDg resource.

ClusReg.cmd is located in the VM5INF folder and is in the path where SFW has been installed. For example, if SFW has been installed on a 64-bit server using the default path, then VM5INF is located at `C:\Program Files(x86)\VERITAS\VERITAS Volume Manager 4.3\VM5INF`

Veritas Cluster Server

This section provides information on known Veritas Cluster Server issues.

VMware virtual environment specific issues

This section provides information on the known issues specific to the VMware virtual environment. Review these details if you plan to configure application monitoring in a VMware virtual environment.

Storage agent issues and limitations

The following limitations and configuration issues apply for non-shared storage configured using NativeDisks, VMNSDg, and VMwareDisks agents in a VMware virtual environment:

- In case the VMwareDisks agent resource is configured manually, care should be taken not to add the operating system disk in the configuration. The VMwareDisks agent does not block this operation. This might lead to a system crash during failover. (2843813)
- Non-shared disks partitioned using GUID Partition Table (GPT) are not supported. Currently only Master Boot Record (MBR) partition is supported. (2861160)
- VMwareDisks agent does not support disks attached to the virtual machine using IDE controllers. The agent resource reports an unknown if IDE type of disks are configured. (2844255)
- Application may fail to start or report an unknown state if VMware vMotion and application failover is triggered simultaneously.
  If VMware vMotion is triggered for the failover target system at the same time as the application is failed over or switched over to the same failover target system, the application successfully stops on the current system, but may fail to start on the target system or report an unknown state. (2861106, 2874316)
  This issue occurs because as a part of the switch over operation the data disk are successfully detached from the current virtual machine but they cannot be attached to the failover target system since the VMware vMotion is in progress for the target system.
  The application agent tries to start the application on the target system for the configured number of attempts. During this operation the application may report an unknown state and eventually start if the time taken for the VMware vMotion to complete does not exceed the time taken for restarting the application for the configured number of attempts. However, if the time taken for VMware vMotion exceeds the application online retry limit, then the application fails to start on the target system.
  Workaround: Ensure that you do not switch the application to a system for which the VMware vMotion is in progress. However, if you happen to do so and the application fails to start on the target system, you must manually start

the application, using the “Start Application” operation from the Symantec High Availability tab.

- VMware snapshot operations may fail if VMwareDisks agent is configured for a physical RDM type of disk. Currently only virtual RDM disks are supported. (2856068)

  This is a limitation from VMware.

- In case VMware HA is disabled and the ESX itself faults. VCS moves the service group to the target failover system on another ESX host. VMwareDisks agent registers the faulted virtual machine on the new ESX host. When you try to power on the faulted system, you may see the following message in the vSphere Client:

  ```
  This virtual machine might have been moved or copied. In order to configure certain management and networking features, VMware ESX needs to know if this virtual machine was moved or copied. If you don't know, answer "I copied it".
  ```

  You must select “I moved it” (instead of the default “I copied it”) on this message prompt. (2853873)

- The VMwareDisks agent updates its attributes, saves the configuration, and makes it read-only during the monitor cycle in the following scenarios:

  - Successfully performing Storage vMotion.
  - Not specifying UUID in the DiskPaths attribute of the resource.

  Before performing the following tasks, discard any changes to the configuration that you do not wish to retain:

  - Performing Storage vMotion.
  - Adding, enabling, or probing a resource that does not have a UUID specified in its DiskPaths attribute. (2865463)

Wizard, Dashboard, and Symantec High Availability tab specific issues

Symantec High Availability tab does not provide an option to remove a node from a VCS cluster [2944997]

The **Symantec High Availability** tab does not support removal of a system from a VCS cluster.

Workaround: You can use VCS commands to remove the system from the VCS cluster. You can also unconfigure the entire VCS cluster from the **Symantec High Availability** tab.

SSO configuration fails if the system name contains non-English locale characters

If you install the Symantec High Availability guest components on a system that has non-English locale characters in its name, then the SSO configuration between such a system and the Symantec High Availability Console host fails. (2910613)

Workaround: Ensure that the system name does not contain any non-English locale characters.

The Symantec High Availability Configuration Wizard gives an error for invalid user account details if the system password contains double quotes (")

This issue occurs while configuring application monitoring using the Symantec High Availability Configuration Wizard.

On the Configuration Inputs panel, if the specified the user account password for the selected systems contains double quotes ("), then the wizard fails to proceed. Even though the user account details entered are correct, it displays an invalid user account details error.(2937186)

Workaround: Ensure that the user account password for the systems which you want add to the VCS cluster systems list do not include the double quotes.

The vCenter Server tasks shown in the vSphere Client may fail to display the correct status if the Symantec High Availability Guest Components installation fails

This issue occurs in case of VMware vCenter Server version 5.1.

If you choose to install the Symantec High Availability Guest Components using the vSphere Client menu and if the installation fails, then the vCenter Server task shown in the vSphere Client gets held-up at 10%. It however, does not display that the installation has failed. (2824900)

Workaround:

Refer to the installer logs at the following locations on the Symantec High Availability Console server. Rectify the cause and re-run the wizard.

```
% ALLUSERSPROFILE %\Symantec\ApplicationHA\Logs\ApplicationHA.log

% ALLUSERSPROFILE %\Symantec\ApplicationHA\Logs\VII\
```

You can also refer to logs on the system at the following location:

```
%ALLUSERSPROFILE %\Veritas\
```

Alternatively, choose to install the guest components using the product installer or the CLI. For details refer to the product-specific installation and upgrade guide.

The Symantec High Availability view does not display any sign for the concurrency violation

In a failover type of VCS cluster configuration the application must be online on one cluster system at any point of time. However, if the application is online on more than one cluster system, then the VCS cluster is in a state of concurrency violation.

The Symantec High Availability view shows that the application is online on more than one cluster system. However, it does not indicate the state as a concurrency violation. On the contrary the concurrency violation is indicated using a red icon in the VCS Java GUI, VOM and the CLI output. (2924826)

The Symantec High Availability installer may fail to block the installation of unrelated license keys

This issue occurs when you initiate to manage the licenses from the Symantec High Availability home view, that is available under the Solutions and Applications menu in the vCenter Server.

The Symantec High Availability installer may fail to validate if the entered license key is applicable to the selected product and may proceed to install the key even if it is applicable to a different product.(2924831)

Even though the installer shows that the validation is successful and installs the license key, you may face unknown issues later.

Workaround: Manage the licenses using the Windows Add or Remove Programs.

For more details refer to the product installation and upgrade guide.

The Symantec High Availability Guest Components Installer does not block the installation of ApplicationHA 6.0 over VCS 6.0.1

If you initiate the installation of ApplicationHA 6.0 through the vSphere Client menu on a system where VCS 6.0.1 is installed, then the installer does not block the installation.(2922421)

However, VCS 6.0.1 and ApplicationHA 6.0 are incompatible products. Do not install ApplicationHA 6.0 on the systems where VCS 6.0.1 is installed.

Symantec High Availability Dashboard fails to differentiate the VCS clusters

While configuring application monitoring using the Symantec High Availability wizard, the wizard generates a unique ID for the VCS cluster after verifying its uniqueness in the network. If the network contains multiple clusters, the wizard verifies the generated ID with the IDs assigned to all the accessible clusters. The wizard does not verify the ID with the clusters that are not accessible during the verification. (2908536)

After the cluster configuration is complete, if any of the inaccessible cluster starts running and its cluster ID matches to the one assigned to this new cluster, then the Symantec High Availability Dashboard fails to differentiate the VCS clusters and displays them as a single cluster.

Workaround: Edit the VCS cluster ID or the cluster name. For more details on modifying the cluster ID or name, refer to the *VCS Administrator's Guide*.

Alternatively, you may consider to unconfigure the VCS cluster and then reconfigure it again. However, note that unconfiguring the VCS cluster requires you to unconfigure your application monitoring configuration.

VCS cluster may display “stale admin wait” state if the virtual computer name and the VCS cluster name contains non-English locale characters

This issue occurs after configuring application monitoring using the Symantec High Availability Configuration wizard.(2906207)

While configuring application monitoring using the Symantec High Availability Configuration wizard, if you specify non-English characters for any of the following, the wizard successfully configures the VCS cluster and completes the application monitoring configuration. However, after the configuration workflow is complete the VCS cluster fails to start and displays the “stale admin wait” state, in the Symantec High Availability tab and the Symantec High Availability Dashboard.

- Virtual name, on the Virtual Network Settings panel
- VCS cluster name, on the Edit Cluster Details panel

Workaround: Edit the virtual name or the VCS cluster ID/name. For more details on modifying the virtual name or the cluster ID/name, refer to the *VCS Administrator's Guide*.

Alternatively, you may consider to unconfigure the VCS cluster and then reconfigure it again. However, note that unconfiguring the VCS cluster requires you to unconfigure your application monitoring configuration.

Issues faced while configuring application monitoring for a Windows service having non-English locale characters in its name

While configuring application monitoring for a Generic Service, if the service that you select on the Windows Service Selection panel has non-English locale characters in its name, you may face the following issues: (2906275)

- The wizard fails to display the service name correctly
- The wizard successfully configures the VCS cluster and completes the application monitoring configuration. However, the resources configured for the service display “unknown” state

- During the Add Failover System operation, the wizard fails to validate the system that you want add to the VCS cluster or as a Failover target

Workaround: Using the VCS Java Console, modify the "ServiceName" attribute of the GenericService agent.

By default the attribute value is set to "Service Display Name". You must change this to "Service Key Name".

The Symantec High Availability configuration wizard fails to configure the VCS cluster if UAC is enabled

This issue occurs while configuring application monitoring in VMware virtual environment.

The Symantec High Availability configuration wizard fails to configure the VCS cluster if the selected cluster systems have User Access Control (UAC) enabled and the user has logged on to the systems using a non-default administrator user account. (2867609, 2908548)

Workaround:

Perform the following steps:

1. Exit the wizard and disable UAC on the systems where you want to configure application monitoring.
2. Reboot the systems and run the Symantec High Availability configuration wizard again.

Alternatively, configure the VCS cluster using the Veritas Cluster Server configuration wizard and then use the Symantec High Availability Configuration Wizard to configure application monitoring.

Generic issues

This section provides information on some general known issues found while configuring application monitoring in a VMware virtual environment.

VMwareDisks resource is unable to go offline if the disks configured are thin provisioned (TP) disks

This issue occurs if VMwareDisks agent is configured to monitor thin provisioned (TP) disks and VMware snapshots are taken.

When the service group is taken offline or failed over, the VMwareDisks resource is unable to go offline. The VMwareDisks agent is not able to perform the disk detach operation as the size of the TP snapshot disk is different than the base disk size.

The following message is displayed in the agent log:

VCS ERROR V-16-10061-22523
VMwareDisks:<AppResourceName>-SG-VMwareDisks:offline:Failed to detach disks from VM on ESX <ESXIPaddress> with error 'Invalid configuration for device '0'.' (2801599)

Workaround: There is no workaround for this issue at this time.

Guest virtual machines fail to detect the network connectivity loss

In a VCS cluster that is configured in a VMware environment, if the ESX host loses its network connectivity, the guest virtual machines residing on the ESX host fail to detect the network loss. The configured virtual IP address remain online even though the underlying network has disconnected. (2901327)

In case of a failover, the application successfully starts on a virtual machine that resides on another ESX host and the configured virtual IP address is accessible over the network. However, when you attempt to failback the application to the original virtual machine, the application status shows "online" but the configured virtual IP address is not accessible.

Workaround: Using the VCS Java Console, configure the "PingHostList" attribute for the VCS NIC agent. For more details, refer to the Veritas Cluster Server Bundled Agents Reference Guide.

VMware vMotion fails to move a virtual machine back to an ESX host where the application was online

In a VCS cluster configured using VMware non-shared disk, if a virtual machine (VM1) on which the application is online is moved to another ESX host (for example, ESX 1), then the storage disk also relocates along with VM1. (2896662)

Now, if the application is failed over to a virtual machine (VM2) that resides on an alternate target ESX host (for example, ESX 2), then the storage disk relocates to VM2. The application is now online on VM2. VMware vMotion now fails to move VM2 back to ESX 1, because of the earlier data logs.

Workaround:

Perform the following steps, to resolve this issue:

1. Fail over the application to VM1
2. Move VM2 to ESX1
3. Fail back the application to VM2

VCS commands may fail if the snapshot of a system on which the application is configured is reverted

In a VCS cluster, the cluster configuration and application monitoring configuration details are replicated on all the cluster systems.

When the snapshot of a cluster system is reverted, that system reverts back to an earlier state while the remaining cluster systems retain the current state. Because of this mismatch, the communication between the cluster systems fail and thus the VCS commands used for the cluster operations fail. (2884317)

Workaround: Perform the following steps, using the command line:

1. Stop the VCS cluster using the following command:

```
hastop -all -force
```

2. Run the following commands sequentially on each cluster system:

```
net stop vcscomm

net stop gab

net stop llt
```

3. Restart the VCS cluster using the following command:

```
hastart -all
```

VDS error reported while bringing the NativeDisks and Mount resources online after a failover

This error may occur in a VCS configuration of VMwareDisks, NativeDisks, and Mount resources. (2886291)

While bringing the NativeDisks and Mount resources online after a failover, the following VDS error message may be reported in the Windows Event Viewer:

Unexpected failure. Error code: D@01010004

Workaround: This is an information message and can be safely ignored.

Hyper-V DR attribute settings should be changed in the MonitorVM resource if a monitored VM is migrated to a new volume

In a Hyper-V DR environment, if the storage of a virtual machine is migrated to a new volume, then its configuration path changes. This causes the MonitorVM resource in the VCS configuration to go to the unknown state. Hence the VM is not monitored for disaster recovery.

Workaround: Modify the VMNames attribute in the MonitorVM resource and set it to the new configuration path.

Event-viewer error message contents are not displayed for VMwareDisks error messages

Workaround: You can search for the EventID in the VMwareDisks log to get the corresponding error message. (2760545)

'Connection Refused' error while using Remote Cluster Configuration Wizard or Global Group Configuration Wizard from Java console

This error may occur in the following scenarios:

- If you try to delete a remote secure cluster using the Remote Cluster Configuration Wizard from Java console
- If you try to configure a service group as a global service group using the Global Group Cluster Configuration Wizard from Java console

The following error is displayed:

Following clusters had problems while connection:Connection refused.

This error occurs because the VCS Java console requires you to re-enter user credentials, even though it should ideally try the logged-in user first. (2740392, 2859468)

Workaround: Re-enter the user credentials.

PrintSpooler resource faults and service group fails over while bringing more than 10 PrintShare resources of newly added printers online. Any further offline of PrintShare resources fails with the PrintShares remaining in 'Unable to go offline' state

These issues may occur if you add more than 10 printers on the active node and then configure PrintShare in VCS and try to bring the service group online on same node. (2924014)

Issue 1: The windows spooler service crashes. As a result, the service group faults and failovers to the failover node.

Issue 2: If issue 1 occurs, then the service group comes online on the failover node. If you now failover to the active node, the failover succeeds. But any more attempts to take the service group offline on the active node will fail.

Workaround: Before bringing any newly added PrintShare online, set the NumThreads attribute to 1. This will increase the time required to bring the resources online. To avoid this, ensure that you revert back this change after the newly added PrintShares come online.

Restart the ‘Server’ service on the node where spooler crashed, after the service group fails over to the failover node. You will be asked to restart the Veritas Storage Foundation Messaging Service in case of a SFW installation. Click Yes.

VCS engine HAD may not accept client connection requests even after the cluster is configured successfully

This issue may occur after you run the VCS Cluster Configuration Wizard (VCW) to configure a cluster to use single sign-on authentication. Even though VCW indicates that the cluster is configured successfully, the VCS high availability engine (HAD) fails to accept client connection requests. This may happen if VCW fails to configure the VCS authentication service on one or more cluster nodes. (2609395)

You may observe one or more of the following:

- If you try to launch any of the VCS service group configuration wizards, you will see the following error:

```
The VCS engine (HAD) is not running on the cluster nodes.
Failed to get the required cluster information.

The wizard will quit.

Error V-16-13-160
```

- If you run the hasys –display command to check the status of HAD in the cluster, you will see the following error on the command prompt:

```
VCS ERROR V-16-1-53007 Error returned from engine:
HAD on this node not accepting clients.
```

- If you try to connect to the cluster using the Cluster Manager (Java Console), you will see the following error:

```
VCS ERROR V-16-10-106
Could not connect to a live system in the cluster localhost:14141.
Please check the application event log for more details.
Closing all windows.
```

- If you run VCW again to reconfigure the cluster, you will see the following error on the Edit Cluster Options panel:

```
Failed to connect to the cluster.
Error reason: Failed to open socket connection to port 14141 on
host <node_name> (1)
```

Workaround: In the following steps we manually modify the cluster that was configured to use single sign-on authentication to use VCS authentication instead and then reconfigure the cluster using VCS Cluster Configuration Wizard (VCW).

Perform the following steps:

1. Stop the VCS high availability engine (HAD) on all the cluster nodes.

   On each cluster node, type the following on the command prompt:

   ```
   net stop had
   ```

2. Perform the remaining steps on one of the cluster nodes.

   Navigate to `%vcs_home%\conf\config` and locate and delete the **.secure** file from that directory.

   Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.

3. From `%vcs_home%\conf\config` directory, locate the configuration file **main.cf** and open it in a text editor.

4. In main.cf, search for the text "**SecureClus=1**" and delete that line altogether.

5. Save the file and close the text editor.

6. Start the VCS engine (HAD) locally on the node where you performed the earlier steps.

   Type the following on the command prompt:

   ```
   hastart
   ```

7. Set the cluster configuration to read/write mode.

   Type the following on the command prompt:

   ```
   haconf -makerw
   ```

8. Add a user to the cluster and assign it with cluster administrator privileges.

   Type the following command:

   ```
   hauser -add <username> -priv Administrator
   ```

9. Enter the password for the user, when prompted.

10 Save and make the cluster configuration read-only.

Type the following on the command prompt:

```
haconf -dump -makero
```

11 Start the VCS high availability engine (HAD) on the remaining cluster nodes.

Type the following on the command prompt:

```
hastart -all
```

Use the cluster user you added in earlier steps to connect to the cluster using Cluster Manager (Java Console). If required, run the VCS Cluster Configuration Wizard (VCW) and reconfigure the cluster to use single sign-on authentication.

VCS engine HAD fails to start on upgrading to SFW HA 6.0.1

This issue may occur if you successfully upgrade to SFW HA 6.0.1, roll back to the previous release, and then upgrade to SFW HA 6.0.1 again. (2894197)

In this scenario, the VCS High Availability Engine (HAD) fails to start and the following message is displayed in the Windows Event Viewer:

```
Failed to create process:
"C:\ProgramFiles\Veritas\VRTSPerl\Bin\perl.exe"
```

Workaround: Repair the installation of SFW HA 6.0.1.

VCS BlackBerry resource BB_MDS fails to probe after service group configuration

After configuring a service group using the BlackBerry service group template from the VCS Cluster Manager (Java Console), the VCS GenericService resource, BB_MDS, for the BlackBerry Mobile Data Service (MDS) Connection Service fails to probe. (2604445)

Workaround: The GenericService resource fails to probe because the VCS BlackBerry service group template uses an incorrect service name for the configured BlackBerry service.

Perform the following steps from the Cluster Manager (Java Console):

1 Double-click **BB_MDS** resource and then in the properties view, click the edit icon to edit the **ServiceName** attribute.

2 For the ServiceName attribute, change the service name from **BlackBerry Mobile Data Service** to **BlackBerry MDS Connection Service**.

3 Click **OK** and then close the attributes properties view.

4 Click **File > Save Configuration** to save the changes to the VCS configuration file.

5 Probe the BB_MDS resource again.

SQL Server 2005 resources do not fault even if detail monitoring fails

This issue may occur when detail monitoring is configured for SQL Server 2005 and IMF is enabled for SQL Server 2005 agents.

If detail monitoring fails (either the database becomes unavailable or there is a failure in the detail monitoring script), the SQL 2005 service resource (SQLServer2005) faults and VCS may then fail over the service group if the agent's FaultOnDetailMonitorScriptFailure attribute is set to 1.

It is observed that when IMF is enabled for SQL 2005 agents, the SQL Server 2005 resource in the service group does not fault. This occurs if the SQL Server 2005 Agent service is set to auto restart (SQL Management Studio > SQL Server Agent Properties).

This occurs because when detail monitoring fails, the SQL 2005 service agent invokes the clean function and as a result the SQL database engine service goes for a restart. As the SQL Agent service depends on the database service, the database service first stops the SQL Agent service as part of its own restart process. As soon as the SQL agent service is stopped, SQL Server immediately restarts the SQL Agent service as defined in the SQL Server Agent properties. As a result, the clean function initiated by the SQL 2005 agent does not complete and the SQL server 2005 resource does not fault. (2572158)

Workaround: To use IMF-based monitoring for SQL Server 2005, disable SQL Server 2005 Agent service auto-restart in the service properties.

Perform the following steps to disable SQL Server Agent service auto-restart:

1 Connect to SQL Server instance using SQL Server Management Studio.

2 From the Object Explorer pane on the left, right-click **SQL Server Agent** and select **Properties**.

3 On the SQL Server Agent Properties dialog box, clear the **Auto restart SQL Server Agent if it stops unexpectedly** check box.

4 Click **OK** to save the changes.

SQL Server 2008 service resource does not fault even if detail monitoring fails

This issue may occur when detail monitoring is configured for SQL Server 2008 and IMF is enabled for SQL Server 2008 agents.

If detail monitoring fails (either the database becomes unavailable or there is a failure in the detail monitoring script), the SQL 2008 service resource (SQLServer2008) faults and VCS may then fail over the service group if the agent's FaultOnDMFailure attribute is set to 1.

It is observed that when IMF is enabled for SQL 2008 agents, the SQL 2008 service resource does not fault. Instead, the SQL 2008 Agent service resource (SQLServerAgent) faults.

This occurs because when detail monitoring fails, the SQL 2008 service agent invokes the clean function and as a result the SQL database engine service goes for a restart. As the SQL Agent service depends on the database service, the database service first stops the SQL Agent service as part of its own restart process. IMF instantly detects that the SQL Agent service has stopped and as a result the SQL Agent resource (SQLServerAgent) in the service group faults. As the SQL Agent resource has faulted, VCS initiates a fail over of the service group. The SQL 2008 service resource receives this VCS initiated offline and therefore does not fault in response to the original detail monitoring failure event.

This is applicable to SQL Server 2008 and SQL Server 2008 R2. (2535806)

Workaround: There is no known workaround for this issue at this time.

If a disk group is deleted from the active node, the change fails to update on the passive nodes

This issue may occur if fast failover is enabled for the disk groups in the cluster.

When you delete a disk group on the active node, the disk group configuration change does not get updated on the passive nodes where the disk group is imported in Deported Read-Only mode. (2236774)

Workaround: From the VEA Console, perform a storage agent rescan on all the passive nodes.

Drive letters assigned to volumes not configured with VCS are not removed during failover

This issue may occur if fast failover is enabled for the disk groups in the cluster.

SFW does not remove drive letters assigned to volumes that are not configured under VCS. As a result, after service group failover the drive letters for such

volumes are visible from Windows Explorer and format dialog boxes may be seen for such those volumes. (2246763)

Disk group import is not clean when a disconnected node becomes available in the cluster

This issue may occur if fast failover is enabled for the disk groups in the cluster.

In a multi-node cluster when you disconnect the node where the service groups are online (active node); the service group resources including the disk groups begin to fail over to the passive node. Now if the active node is connected to the network again, SFW begins to import the disk groups on the node in a Deported Read-Only mode. However, in some cases, the disk group import either takes a long time, or is not in a clean state. The disk groups status shows as Deported None and one or more disks may be missing from the configuration.(2289119)

This issue is observed only when the storage is connected using Fibre Channel (FC).

Workaround: Run the vxassist rescan command on the active node to remove the missing disks.

After a complete storage disconnect on the active node, the service groups failover to one of the passive nodes. However on storage reconnect, the failback to the active node does not happen until all the disk removal and arrival events are processed. Symantec recommends that in case of a large storage configuration, you reboot the node after reconnecting the storage.

Invalid arguments error seen when vxdg dginfo command is run on a passive node

This issue may occur if fast failover is enabled for the disk groups in the cluster.

Invalid arguments error may be seen when you run the vxdg dginfo command on a passive node where the disk group Read-Only refresh operation is in progress. (2274302)

Workaround: Wait for some time for the refresh operation to complete and then run the vxdg dginfo command again.

CPU spikes observed in VMDg and MountV agent processes

This issue may occur if VCS IMF and SFW fast failover is enabled for the disk groups in the cluster.

Intermittent CPU spikes are observed for the VCS Volume Manager Diskgroup (VMDg) and MountV agent process (VCSAgDriver.exe). (2299143)

Workaround: Disable the VMDg detail monitoring by setting the DetailMonitorFreq attribute value to 0.

One or more VMNSDg resources may fail to come online during failover of a large service group

In a large application service group configuration with many VMwareDisks and VMNSDg resources, one or more VMNSDg resources may fail to come online with the following error: (2924009)

```
VMNSDg:<resource_name>:online:Online diskgroup : The Diskgroup is
not present.
```

Workaround: To avoid this issue, perform one of the following procedures:

- Bring the corresponding VMwareDisks resource in the service group online and run the following command:

  ```
  vxassist rescan
  ```

  Then bring the other resources of the service group online.
- Set the OnlineRetryLimit for VMNSDg resource to greater than 1.

PlugPlayManager error displayed in System Event Viewer while bringing VMwareDisks resources offline

This issue is observed on Windows Server 2008 (x64). During an application failover, while bringing the VMwareDisks resources offline, the following error message is displayed in the Event Viewer: (2926238)

```
The device <device_name> disappeared from the system without first
being prepared for removal.
```

Workaround: VCS prepares the disks for the detach operation. Hence this message can be safely ignored.

Several issues while you configure VCS on systems where Symantec Endpoint Protection (SEP) version 12.1 is installed

The following issues may occur while you install and configure VCS on systems where Symantec EndPoint Protection (SEP) version 12.1 is installed. (2439737, 2487369, 2574748, 2530343)

- The VCS Cluster Configuration Wizard (VCW) may fail to connect to the systems while configuring the cluster.
  The following error may be displayed:

  ```
  WMI Connection failed. Error=800706BA
  ```

- If LLT is configured over UDP on an IPv6 network, the status of the VCS High Availability Engine (HAD) on all the remote nodes in the cluster remains in the REMOTE_BUILD state.
- If you set up Disaster Recovery using the Global Cluster Option (GCO) in an IPv6 environment, the status of the remote cluster (cluster at the secondary site) shows as “initing”.

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the Installation and Upgrade Guide for list of ports and services used by SFW HA.
- For the VCW issue, add a custom rule to the SEP firewall policy and define the properties as follows:
  - Rule name: Type **VCS TCP 135** as the name.
  - Action: Select **Allow this traffic**.
  - Protocol: Select **TCP** from the drop-down list.
  - Remote Ports: Type **135** in the field.
- For IPv6 networks, configure the SEP firewall policy to allow IPv6 traffic on the cluster nodes.
  Edit the Firewall Rules table and edit the following settings:
  - **Block IPv6**: Change the Action field value to “Allow”.
  - **Block IPv6 over IPv4 (Teredo)**: Change the Action field value to “Allow”.
  - **Block IPv6 over IPv4 (ISATAP)**: Change the Action field value to “Allow”.
    Refer to the SEP documentation for detailed instructions on how to edit the firewall policy.
- For the LLT over UDP issue on IPv6 network (REMOTE_BUILD issue), perform these steps. Note that the steps require you to stop the Veritas High Availability Engine (HAD).
  - Stop the VCS HAD in the cluster.
    From one of the cluster nodes, type the following on the command prompt:
    ```
    hastop -all -force
    ```
  - Perform the following steps on all the cluster nodes, one node at a time:
    - Stop the LLT service on the cluster node.
      Type the following on the command prompt:
      ```
      net stop llt
      ```

- Navigate to `%vcs_root%\comms\llt` and open the `lltab.txt` file in a text editor.
  Here, vcs_root typically resolves to `C:\Program Files\Veritas`.
- Modify all the link entries in the llttab.txt file as follows:
  Change this entry: `link <link#> udp6 - udp6 <udpport#> - <IPv6address> -`
  to this: `link <link#> udp6 - udp6 <udpport#>` **`1380`** `<IPv6address> -`
  Note that "1380" is added to the link entry. This defines the MTU (packet size) that LLT uses for its communication.
  For example, a sample link entry is as follows: `link Link1 udp6 - udp6 50000 1380 2001:db8:0:11:5c56:6867:e398:6152 -`
- Save and close the llttab.txt file.

- Start the VCS HAD in the cluster.
  From one of the cluster nodes, type the following on the command prompt:
  `hastart -all`

Cluster node may become unresponsive if you try to modify network properties of adapters assigned to the VCS private network

This issue may occur if after configuring the cluster you try to modify the network properties of the adapters assigned to the VCS private network. After you make changes in the adapter properties dialog box and click OK, the properties dialog may hang and in some cases the cluster node itself may become sluggish or unresponsive. (2408878)

Workaround: To resolve this issue, you must either terminate the network properties dialog from Windows Task Manager or restart the VCS LLT service.

Recommendation: If you want to modify the network properties of the adapters assigned to the VCS private network, Symantec recommends that you perform the following steps in the given order.

To modify the private network adapter properties

1. Stop the Veritas High Availability Engine (HAD) on one of the passive nodes. A passive node is a node where are no service groups online.

   Type the following on the command prompt:

   ```
   hastop -local -force
   ```

2. Stop the VCS LLT service on the node.

   Type the following on the command prompt:

   ```
   net stop llt
   ```

3. Modify the network properties of the network adapters on the node, as desired.
4. Start the Veritas High Availability Engine (HAD) on the node.

   Type the following on the command prompt:

   ```
   hastart
   ```

5. Repeat step 1 to step 4 on all the other passive nodes in the cluster.
6. Switch the online service groups from the active node to another node in the cluster. The active node is the node where the service groups are online.
7. Repeat step 1 to step 4 on the active node.
8. If you have assigned a new IP address to any of the network adapters used in the VCS private network, you must reconfigure the private network in the cluster using the VCS Cluster Configuration Wizard (VCW).

   This step is required only if you have configured LLT over UDP in the cluster.

VCS services do not start on systems where Symantec Endpoint Protection 11.0 MR5 version is installed

This issue occurs on Windows Server 2008 systems. (1710556)

The VCS High Availability Engine (HAD) may fail to start if you install and configure VCS on systems where Symantec EndPoint Protection (SEP) version 11.0 MR5 is already installed.

The following error may be displayed:

```
Failed to start the cluster. Error=FFFFFFFF. Failed to start
services on all the nodes.
```

Workaround: Disable Symantec EndPoint Protection (SEP) on all the systems where you have installed VCS and then configure the VCS cluster.

After completing the cluster configuration tasks, enable SEP on all the systems.

File Share Configuration Wizard may create dangling VMDg resources

This issue occurs if you use folder mounts for creating file share service groups. The VCS File Share Configuration Wizard successfully creates file share service groups, but may fail to configure dependency for one or more VMDg resources. The VMDg resources are configured correctly but may not be part of the overall service group resource hierarchy. (2097155)

Workaround: You may have to manually configure the dependency for such dangling VMDg resources. VMDg resources are typically configured as child of MountV resources in VCS service groups.

Exchange 2007 database fails to mount after performing a database restore operation using Backup Exec 12.5

This issue occurs if you are using Backup Exec (BE) 12.5 on Windows Server 2008 systems where Exchange 2007 is configured with VCS.

When the **Mount database after restore** option is enabled in the Restore Properties Settings for Exchange, after performing a database restore operation using Backup Exec (BE) 12.5 the Exchange 2007 database fails to mount on the cluster node. (1674378)

Storage agent resources may go into an unknown state after upgrading Windows Server 2008 SP1 systems

If you upgrade the Windows Server 2008 SP1 operating system after upgrading SFW HA cluster nodes, the storage agent resources (VMDg and MountV) in the cluster may go into an UNKNOWN state and the storage agent service (vxvm) may fail to start.

This issue occurs if you upgrade Windows Server 2008 SP1 to Windows Server 2008 SP2 or Windows Server 2008 R2. (1786188, 1835035, 1835031)

Workaround: Complete the following steps depending on the upgrade case:

To perform Windows Server 2008 SP1 to Windows Server 2008 SP2 upgrades:

1. Stop the Veritas High Availability Engine (HAD) on all the cluster nodes.

   Type the following at the command prompt:

   ```
   hastop -all -force
   ```

2. From the command prompt, run %vmpath%\FixVDSKey.bat on all the cluster nodes.

   %vmpath% is the default installation directory for Veritas Volume Manager, typically, `C:\Program Files\Veritas\Veritas Volume Manager 5.1`. (`C:\Program Files (x86)\Veritas\Veritas Volume Manager 5.1` on 64-bit systems)

3. Reboot all the cluster nodes.

4. Verify that the Veritas High Availability Engine (HAD) is running on all the cluster nodes.

   Type the following at the command prompt:

   ```
   hasys -state
   ```

   The status should display as RUNNING.

   If HAD is not running, start it.

   Type the following at the command prompt:

   ```
   hastart
   ```

To perform Windows Server 2008 SP1 to Windows Server 2008 R2 upgrades

1. From the Windows Control Panel, launch **Add or Remove Programs**.
2. Click **Service Pack 2 for SFW 5.1, SFWHA 5.1, and VCS 5.1 for Windows** and then click **Change**.
3. On the **Symantec Product Installer Selection** panel, click **Repair**, and then click **Next**.
4. On the **Validation** panel, click **Next**.
5. On the **Summary** panel, click **Repair**.
6. After the installer completes, click **Next** and then click **Finish**.
7. Reboot the systems.

While taking volume-based snapshots using the VEA console, the Exchange 2007 service group fails over

While taking volume-based snapshots from the Veritas Enterprise Administrator (VEA) console on a node where an Exchange 2007 service group is online, the VSS Exchange writer crashes causing the MSExchangeIS resource to fault. This results in the Exchange 2007 service group to fail over to another cluster node. This happens if the snapshots are taken of volumes that belong to an Exchange 2007 storage group.

This issue is applicable to Exchange Server 2007 Service Pack 1 running on 64-bit Windows Server 2008 Service Pack 2 systems. (1727397)

Workaround: Install Update Rollup 9 for Microsoft Exchange Server 2007 Service Pack 1.

See: http://support.microsoft.com/kb/970162

MountV resource takes time to offline on Windows Server 2008 systems

The MountV resource may take considerable amount of time to go offline on Windows Server 2008 systems. (1189260, 1235123)

The MountV agent log displays the following message:

```
VCS WARNING V-16-10051-9023 MountV:<servicegroupname>-MountV:offline:
Failed to lock
volume [2:5]
```

Workaround: Symantec recommends that on Windows Server 2008 systems, set the value of the MountV attribute ForceUnmount to ALL. The MountV agent forcibly dismounts the volume irrespective of the type of access an application has to that volume.

File share service group may fail to come online on Windows Server 2008 systems

While configuring file shares if you select Bring the service group online option on the **File Share Configuration Wizard Completion** panel, the file share service group may fail to come online. (1204865)

The following message appears in the log:

```
VCS ERROR V-16-10051-10506
FileShare:fs-FileShare:online:Unknown error for folder
<sharename>
```

This issue occurs on Windows Server 2008 systems.

To work around this file share service group issue

1 Flush the file share service group.

- From the Java Console, click the cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the **View** panel.
- Click **Flush**, and then click the appropriate system from the menu.

2 Bring the file share service group online.

On a Windows Server 2008 or 2008 R2 node, installer may fail to restart the services after the installation is complete

The product installer may fail to restart the services, if the logged on user account does not have the required administrative privileges on the cluster node.

This issue occurs even if you choose to log on using the User Access Control (run as administrator) option. (2140506)

SharePoint 2010 resource faults with an initialization error

After configuring the SharePoint service group, the VCS SharePoint 2010 agent resource (SharePointServer) may fault.(2102270)

The SharePoint 2010 agent log may contain the following messages:

VCS ERROR V-16-10051-13583 SharePointServer:<SharePointcomponentname>:monitor:Provider Initialization failed [4, 0x800705AF]

VCS ERROR V-16-20083-105 SharePointServer:<SharePointcomponentname>:monitor:Provider Initialization failed [4, 0x800700A4].

Workaround: From the Windows Services MMC snap-in, restart the Windows Management Instrumentation (WMI) service and then probe the SharePoint 2010 resource.

This is a known Microsoft issue and a hotfix was not available at the time of this release.

MSMQ resource fails to come online if the MSMQ directory path contains double byte characters

The VCS MSMQ resource may fail to come online and may eventually fault if the MSMQ directory path contains Double Byte Character Set (DBCS) characters. (584162, 2121635)

The MSMQ agent log may contain the following message:

V-16-2-13066 Agent is calling clean for resource (MSMQresourcename) because the resource is not up even after online completed.

The Windows Event Viewer may display the following message:

The logger files cannot be initialized (Error: 0x800700003) The file <filename> in the <MSMQdirectory> folder is corrupted or absent. To start the Message Queuing service without losing consistency, you must correct or recover this file.

Workaround: This is a limitation from Microsoft. Do not use DBCS characters in the MSMQ directory paths.

When running Enterprise Vault Configuration Wizard, Enterprise Vault may fail to connect to SQL Server

In the Enterprise Vault Configuration Wizard, if SQL Server is previously configured for high availability, you must enter the name in the format *virtualservername**instancename*.

Rarely, when you click **Next** after specifying *virtualservername**instancename*, a message is displayed that Enterprise Vault failed to connect to SQL Server. If you receive this message, use the following workaround. (2432332)

Workaround: In the SQL Server name field, replace the virtual server name with the physical server name (entered as *physicalservername**instancename*) and click **Next**. When the wizard displays the next panel, click **Back**. Re-enter the name as virtualservername\instancename and click **Next**. Continue with the wizard.

Error while switching global service groups using Veritas Operations Manager (VOM) 3.0

The following issue may occur if you are using Veritas Operations Manager (VOM) 3.0 for administering VCS global service groups configured in secure clusters. (2084898)

If you try to switch global service groups between clusters, the operation fails with the following error:

```
VCS WARNING V-16-1-50824 Command (hagrp -switch <servicegroupname>
<targetsystemname> <targetclustername>) failed. At least Group
Operator privilege required on remote cluster <targetclustername>.
```

Workaround: VOM uses the Veritas Storage Foundation Messaging Service to run VCS commands. This service runs in the Local System account context. Configure this service to run in the Domain Administrator account context and then perform the switch operation.

Change the service account on each of the managed hosts in the clusters.

Perform the following steps on each of the cluster nodes (managed hosts):

1. Open the Windows Services MMC snap-in.
2. Right-click **Veritas Storage Foundation Messaging Service** and then click **Properties**.
3. Click the **Log On** tab and do the following:
   - Click **This account**, click **Browse**, and in the Select User dialog box specify a user account that has Domain Administrator privileges.
   - Click **OK**.
4. Type the account password in the Password and Confirm password fields and then click **OK**.
5. Proceed with the service group operations.

Global group fails to come online on the DR site with a message that it is in the middle of a group operation

When the node that runs a global group faults, VCS internally sets the MigrateQ attribute for the group and attempts to fail over the global group to another node within the local cluster. The MigrateQ attribute stores the node name on which the group was online. If the failover within the cluster does not succeed, then VCS clears the MigrateQ attribute for the groups. However, if the groups have dependencies which are more than one-level deep, then VCS does not clear the MigrateQ attribute for all the groups.(1795151)

This defect causes VCS to misinterpret that the group is in the middle of a failover operation within the local cluster and prevents the group to come online on the DR site. The following message is displayed:

```
VCS Warning V-16-1-51042 Cannot online group global_group.
Group is in the middle of a group operation in cluster
local_cluster.
```

Workaround: Perform the following steps on a node in the local cluster which is in the running state.

To bring the global group online on the DR site

1 Check whether the MigrateQ attribute is set for the global group that you want to bring online on the remote cluster. Type the following on the command prompt:

```
hagrp -display -all -attribute MigrateQ
```

This command displays the name of the faulted node on which the group was online.

2 Flush the global group that you want to bring online on the remote cluster. Type the following on the command prompt:

```
hagrp -flush global_group -sys faulted_node -clus local_cluster
```

where:

- *global_group* is the group that you want to bring online on the remote cluster.
- *faulted_node* is the node in the local cluster that hosted the global group and has faulted.
- *local_cluster* is the cluster at the local site.

The flush operation clears the node name from the MigrateQ attribute.

3 Bring the service group online on the remote cluster.

Type the following on the command prompt:

```
hagrp -online global_group -any -clus remote_cluster
```

VCS cluster configuration fails if Symantec Endpoint Protection 11.0 MR3 version is installed

The VCS Cluster Configuration Wizard (VCW) fails to configure the cluster on systems where Symantec Endpoint Protection (SEP) 11.0 MR3 version is installed.(1455690)

The following error is displayed:

```
Failed to start the cluster. Error=FFFFFFFF. Failed to start
services on all the nodes.
```

Perform the following workaround to resolve this error.

To resolve error message

1. Create a custom rule in the SEP firewall rules table. Specify the following details for the rule:
   - Rule type: Application
   - Application name: llt.sys
   - Action: allow
2. Move this rule to the top of the firewall rules table and then apply the firewall policy again.
3. Ensure that the SEP clients on the systems receive this policy and then proceed with the cluster configuration task.

Refer to the SEP documentation for detailed instructions on creating custom firewall rules.

SQL Server 2008 Analysis and Agent service resources may go in to an unknown state

The SQL Server 2008 Analysis service and SQL Server 2008 Agent service resources may go in to an UNKNOWN state when you bring the SQL Server 2008 service group online. (1466012)

The following error is logged on the GenericService agent log:

```
VCS ERROR V-16-10051-6012
GenericService:MSOlap-NEW:online:Failed to wait for the
service 'MSOLAP$NEW' to start. Error = 258
VCS ERROR V-16-10051-6012
GenericService:SQLServerAgent-NEW:online:Failed to wait for
the service 'SQLAgent$NEW' to start. Error = 258
```

Workaround: Probe the resources if they are in the UNKNOWN state.

Symantec Endpoint Protection security policy may block the VCS Cluster Configuration Wizard

While configuring a cluster, the VCS Cluster Configuration Wizard (VCW) may fail to ping systems that are selected to be a part of the cluster. As a result, you cannot configure the cluster. This may happen in case Symantec Endpoint Protection (SEP) client is installed on the selected systems. VCW uses Internet Control Message Protocol (ICMP) to ping systems and ICMP traffic is blocked in SEP, by default.(1315813)

Workaround: Create a custom rule in SEP to allow ICMP traffic in both directions.

Ensure that you create this rule on all the systems that are going to be part of the cluster. Refer to the SEP documentation for instructions.

Saving large configuration results in very large file size for main.cf

If your service groups have a large number resources or resource dependencies, and if the PrintTree attribute is set to 1, saving the configuration may cause the configuration file to become excessively large in size and may affect performance.(616818)

Workaround: Disable printing of resource trees in regenerated configuration files by setting the PrintTree attribute to 0.

AutoStart may violate limits and prerequisites load policy

The load failover policy of Service Group Workload Management may be violated during AutoStart when all of the following conditions are met:

- More than one autostart group uses the same Prerequisites.
- One group, G2, is already online on a node outside of VCS control. The other group, G1, is offline when VCS is started on the node.
- The offline group is probed before the online group is probed.

In this scenario, VCS may choose the node where group G2 is online as the AutoStart node for group G1 even though the Prerequisites load policy for group G1 is not satisfied on that node.

Workaround: Persistently freeze all groups that share the same Prerequisites before using `hastop -force` to stop the cluster or node where any such group is online. This workaround is not required if the cluster or node is stopped without the force option.

Trigger not invoked in REMOTE_BUILD state

In some situations, VCS does not invoke the in jeopardy trigger if the system is a REMOTE_BUILD state. VCS fires the trigger when the system goes to the RUNNING state.

Some alert messages do not display correctly

The following alert messages do not display correctly: (612268)

51030	Unable to find a suitable remote failover target for global group %s. Administrative action is required.

51031	Unable to automatically fail over global group %s remotely because local cluster does not have Authority for the group.
50913	Unable to automatically fail over global group %s remotely because clusters are disconnected and ClusterFailOverPolicy is set to %s. Administrative action is required.
50914	Global group %s is unable to failover within cluster %s and ClusterFailOverPolicy is set to %s. Administrative action is required.
50916	Unable to automatically failover global group %s remotely due to inability to communicate with remote clusters. Please check WAN connection and state of wide area connector.
50761	Unable to automatically fail over global group %s remotely because ClusterList values for the group differ between the clusters. Administrative action is required.
50836	Remote cluster %s has faulted. Administrative action is required.
51032	Parallel global group %s faulted on system %s and is unable to failover within cluster %s. However, group is still online/partial on one or more systems in the cluster
51033	Global group %s is unable to failover within cluster %s and AutoFailOver is %s. Administrative action is required.

NetBackup may fail to back up SQL Server database in VCS cluster environment

In a VCS cluster environment, backup of the SQL database with Symantec NetBackup may fail.

The batch (.bch) file generated by NetBackup for backing up a SQL Server database must contain the following keyword in a VCS cluster environment:

`BROWSECLIENT` *`VirtualServer`*

where *`VirtualServer`* is the SQL Server virtual server name used in the VCS SQL Server service group.

With NetBackup 7.1, the batch file generated for the SQL database has been observed to be missing this keyword, and as a result, the backup fails. (2415667)

Workaround: Manually add the missing `BROWSECLIENT` *`VirtualServer`* keyword to the batch file after it is created.

Issues related to the VCS engine

The following issues relate to the VCS engine.

Engine may hang in LEAVING state

When the command hares -online is issued for a parent resource when a child resource faults, and the `hares -online` command is followed by the command `hastop -local` on the same node, then the engine transitions to the LEAVING state and hangs.

Workaround: Issue the command `hastop -local -force`

Timing issues with AutoStart policy

Consider a case where the service group is offline and engine is not running on node 1. If you restart the engine on node 1 after HAD is killed on node 2 and before the engine is restarted on node 2, then VCS does not initiate the autostart policy of the group.

Issues related to Cluster Manager (Java Console)

The following issues relate the Cluster Manager (Java Console)

Delay in refreshing the VCS Java Console

You may observe a delay in refreshing the VCS Java Console, if the cluster is configured for Single Sign-on authentication.

This issue may occur because the Veritas Enterprise Administrator Service (VxSVC) service may sometime consume 100% of the CPU memory.(2570302)

Cluster connection error while converting local service group to a global service group

This issue occurs while converting a local service group into a global service group using the Global Group Configuration Wizard from the Cluster Manager (Java Console). While specifying the remote cluster information, if you choose the **Use connected clusters credentials** option for the cluster admin user, the wizard fails to validate the user credentials even if the logged on user is a cluster administrator. (1295394)

The following error is displayed:

```
VCS WARNING V-16-10-73 Following clusters had problems while
connection: Cluster <cluster name>: Connection Refused
```

Workaround: You must select the **Enter new credentials** option and manually specify the cluster administrator credentials.

Repaint feature does not work properly when look and feel preference is set to Java

When a user selects the **Java Look and Feel in the Preferences** dialog box and the look and feel has changed, repainting does not work in that the **Preferences** dialog box does not change as it should and the panel is not clearly visible.(1082952)

Workaround: After selecting the **Java Look and Feel in the Preferences** dialog box, close the Java GUI and then reopen it. You should then be able to select other tabs in the **Preference** dialog box.

Exception when selecting preferences

On Windows systems, selecting the Java (Metal) look and feel of the Java Console may cause a Java exception. (585532)

Workaround: After customizing the look and feel, close restart the Java Console.

Java Console errors in a localized environment

When connected to cluster systems using locales other than English, the Java Console does not allow importing resource types or loading templates from localized directories.

Workaround: Copy the types files or templates to directories with English names and then perform the operation.

Common system names in a global cluster setup

If both local and remote systems have a common system name in a global cluster setup, group operations cannot be performed on those systems using the Java console.

Workaround: Use command-line interface to perform group operations.

Agent logs may not be displayed

If VCS is installed at a different location (at a location other than the default location), the VCS agent logs may not be visible from the Java Console. (643753)

Workaround: Copy the `bmc` and `bmcmap` files to the location specified in Table 1-3:

Table 1-1 bmc and bmcmap file location

Copy from this directory	Copy to this directory
(For English) D:\Program Files\Veritas\messages\en Where, D: is the drive on which VCS is installed.	%VCS_HOME%\messages\en Where, %VCS_HOME% is the default installation directory for VCS, typically C:\Program Files\Veritas\Cluster Server.

Global service groups

The following are global service groups issues.

VCW configures a resource for GCO in a cluster without a valid GCO license

The VCS Configuration Wizard (VCW) enables you to configure a resource for global clustering, even if the cluster does not have a valid license for the Global Cluster Option (GCO). You can successfully bring a GCO resource online, take it offline, or switch it between nodes in a cluster. However, the following message is logged on the engine log if you attempt to connect to a remote cluster:

```
VCS WARNING V-16-3-18000 Global Cluster Option not licensed.
Will not attempt to connect to remote clusters
```

Workaround: Symantec recommends that you do not configure a global cluster resource in a cluster without a valid GCO license.

Group does not go online on AutoStart node

Upon cluster startup, if the last system on which the global group is probed is not part of the group's AutoStartList, then the group will not AutoStart in the cluster. This issue affects only global groups. Local groups do not experience this behavior.

Workaround: Ensure that the last system to join the cluster is a system in the group's AutoStartList.

Cross-cluster switch may cause concurrency violation

If the user tries to switch a global group across clusters while the group is in the process of switching within the local cluster (across systems), then the group will be online on both the local and remote clusters. This issue affects only global groups. Local groups do not experience this behavior.

Workaround: Ensure that the group is not switching locally before attempting to switch the group remotely.

Declare cluster dialog may not display highest priority cluster as failover target

When a global cluster fault occurs, the **Declare Cluster** dialog enables you to fail groups over to the local cluster. However, the local cluster may not be the cluster assigned highest priority in the cluster list.

Workaround: To bring a global group online on a remote cluster, from the Java Console, right-click the global group in the Cluster Explorer tree or **Service Group View**, and use the Remote Online operation to bring the group online on a remote cluster.

Fibre Channel adapters may require modified settings

The following issues apply to VCS with specific Fibre Channel host bus adapters.

Emulex Fibre Channel adapters

For servers configured with Emulex Fibre Channel host bus adapters, you must modify settings of the adapter. The default settings of the adapter do not ensure proper function of SCSI reserve and release.

Workaround: Be sure that the host bus adapter has the proper drivers installed.

Modify the Topology, ResetFF, and ResetTPRLO drive settings in the Emulex adapter BIOS settings, as instructed in the following workaround.

To workaround this issue

1. Locate and run the `Emulex` utility for changing Miniport driver settings.
2. Select **Configuration Settings**.
3. Select **Adapter Settings**.
4. Set the **Topology** parameters to 1, Permanent, and Global.
5. Set the **ResetFF** parameters to 1, Permanent, and Global.
6. Set the **ResetTPRLO** parameters to 1, Permanent, and Global.
7. Save the configuration.
8. Repeat step1 through step 7 for all Emulex adapters in each system.
9. Reboot the systems.

Note: When using EMC storage, you must make additional changes to Emulex host bus adapter settings. See TechNote 245039 on this topic at,

http://entsupport.symantec.com.

QLogic Fibre Channel adapters

When configured over QLogic Fibre Channel host bus adapters, the DiskReservation agent requires the Target Reset option of the adapter to be enabled. By default, this adapter option is disabled, causing the agent to hang during failover.

To workaround this issue

1. During system startup, press ALT+Q to access the QLogic adapter settings menu.
2. Select **Configuration Settings**.
3. Select **Advanced Adapter Settings**.

4 Set the **Enable Target Reset** option to Yes.

5 Save the configuration.

6 Repeat step 1 through step 5 for all QLogic adapters in each system.

7 Reboot the systems.

If VCS upgrade fails on one or more nodes, HAD fails to start and cluster becomes unusable

This issue may happen in cases where you are upgrading a multi-node VCS cluster. If the upgrade succeeds on at least one node but fails on one or more nodes in the cluster, the VCS High Availability Engine (HAD) may fail to start on the nodes on which the upgrade has failed.

The VCS installer does not let you remove VCS from those nodes with an error that those nodes are part of a cluster. The VCS Cluster Configuration Wizard (VCW) does not let you remove those nodes from the cluster with an error that the nodes have a different version of VCS installed.

As a result, you cannot perform any operations on the cluster. (1251272)

Workaround: To get the cluster running, you must manually remove the nodes on which VCS upgrade failed, from the cluster. Then, use the cleanup scripts to remove VCS from the nodes on which the upgrade failed, reinstall VCS, and add the nodes to the cluster.

Perform the following steps to remove the nodes on which the VCS upgrade failed, from the cluster:

To workaround this issue

1 Stop HAD and LLT on all the cluster nodes.

Type the following on the command prompt:

```
net stop had
```

```
net stop llt
```

2 On a node on which VCS was upgraded successfully, open the file `llthosts.txt` and delete the entries of all the cluster nodes on which the upgrade failed.

For example, consider a cluster with three nodes, N1, N2, and N3.

The `llthosts.txt` file contains the following entries:

```
# This is program generated file, please do not edit.
0 N1
1 N2
2 N3
```

If the upgrade failed on N3, delete the last entry from the file.

So the modified `llthosts.txt` file should look like this:

```
# This is program generated file, please do not edit.
0 N1
1 N2
```

The `llthosts.txt` file is typically located at `C:\Program Files\VERITAS\comms\llt`.

Here C:\ is the drive on which VCS is installed.

3 On the node on which you performed step 2, open the `gabtab.txt` file and modify the entry to reflect the exact number of nodes in the cluster.

The `gabtab.txt` file contains the following entry:

```
#This is program generated file, please do not edit.
gabconfig -c -n <number of nodes in the cluster>
```

The *<number of nodes in the cluster>* should be the number of nodes on which VCS was upgraded successfully.

Considering the example in step 2 earlier, the `gabtab.txt` file contains the following entry:

```
#This is program generated file, please do not edit.
gabconfig -c -n 3
```

As the upgrade failed on one out of the total three nodes in the cluster, the entry should look like this:

```
#This is program generated file, please do not edit.
gabconfig -c -n 2
```

The `gabtab.txt` file is typically located at `C:\Program Files\VERITAS\comms\gab`.

Here C:\ is the drive on which VCS is installed.

4 From the Windows Services snap-in, change the startup type of the Veritas High Availability Engine (HAD) service to Manual.

5 Repeat step 2, step 3, and step 4 on all the nodes on which VCS was upgraded successfully.

6 On one of the nodes on which VCS was upgraded successfully, open the VCS configuration file main.cf in a text editor and remove the entries of all the cluster nodes on which the VCS upgrade failed.

The `main.cf` file is located at %VCS_Home%\conf\config.

The variable *%VCS_HOME%* is the default installation directory for VCS, typically C:\Program Files\VERITAS\Cluster Server.

7 Start HAD on the node on which you modified the VCS configuration file in step 6 earlier.

Type the following on the command prompt:

```
net start had
```

You can remove VCS from the affected nodes using the cleanup scripts that are provided with the software. These scripts are .bat files located in the \Tools\vpi

directory on the software DVD. Refer to the `readme.txt` file located in the directory for details on how to use the cleanup scripts. After removing VCS, install VCS using the product installer and then add the nodes to the cluster.

Contact Symantec Technical Support for more information.

Custom settings in the cluster configuration are lost after an upgrade if attribute values contain double quote characters

This issue may occur if attribute or argument values of the configured resources in the cluster contain double quote characters (“ “).

If double quotes are used and you upgrade the cluster, all the custom settings made in the cluster configuration are lost. The upgrade itself is successful and the cluster is able to start. But all the customized settings (custom agents, attributes values, arguments and settings) are lost.

Note that these double quotes are not those added by the VCS wizards or Cluster Manager (Java Console). Here’s an example of an agent attribute value:

```
StartProgram @CLUSSYSTEM1 = "\"C:\\ Windows \\ System32 \\
notepad.exe\""
```

The double quotes at the start and end of the entire path are valid. The double quotes included within the starting and ending double quotes cause this issue. (2837356)

Workaround: Symantec recommends that before you upgrade, you take a backup of the cluster configuration files, main.cf and types.cf. The files are located at:

`%vcs_home%\conf\config`.

Here *%vcs_home%* is the default VCS installation directory, typically `C:\Program Files\Veritas\Cluster Server`.

If there are custom settings made in the cluster configuration, then before upgrading the cluster you modify the resource attributes and argument values to remove the double quotes. If you have already upgraded the cluster, then you will have to modify the cluster again to include all the customizations required in the configuration. For the custom settings, you can refer to the cluster configuration files that you backed up before the upgrade.

Options on the Domain Selection panel in the VCS Cluster Configuration Wizard are disabled

While running the VCS Cluster Configuration Wizard (VCW), the options to retrieve a list of systems and users in the domain on the **Domain Selection** panel are available only for the first time you run the wizard. If you click **Next** and then

click **Back** to go back to the panel, all or some of these options appear disabled. (1213943)

Workaround: Exit and launch the wizard again.

Service group dependency limitations

The following are service group dependency limitations.

Secure clusters

The following issues relate to secure clusters.

Upgrading a secure cluster may require HAD restart

After upgrading a secure cluster, you may not be able to connect to the Cluster Manager Console (Java GUI) and may observe the following error in the VCS engine log: (849401, 1264386)

```
VCS ERROR V-16-1-50306 Failed to get credentials for VCS Engine(24582).
```

The following error is displayed if you run any VCS commands from the command line:

```
VCS ERROR V-16-1-53007 Error returned from engine:
HAD on this node not accepting clients.
```

To work around this upgrading a secure cluster issue

1 Restart the Veritas High Availability Engine (HAD).

Type the following at the command prompt:

```
net stop had
```

```
net start had
```

2 Verify that HAD is running.

Type the following at the command prompt:

```
hasys -state
```

The state should display as RUNNING.

New user does not have administrator rights in Java GUI

In a secure cluster, add a new domain user to the cluster from the command line with Cluster Administrator privileges. Try to log on into the Cluster Console (Java GUI) using the newly added user privileges. The new user is logged on as a `guest` instead of an *administrator*. (614323)

Workaround: When adding a new user to the cluster, add the user name without the domain extension. For example, if the domain is `vcstest.com` then the user name must be specified as `username@vcstest`.

VCS with Microsoft Exchange Server

The following issues apply to VCS with Microsoft Exchange Server.

Exchange service group does not fail over after installing ScanMail 8.0

This issue occurs when you try to install ScanMail 8.0 in an Exchange cluster. After installing ScanMail on one node in a cluster, when you switch the service group to another node to install ScanMail, the service group does not come online.

You can complete the ScanMail installation by making changes to the registry keys and bring the Information Store online. But the Exchange services continue to stop intermittently, causing the resources and the service group to fault and fail over.(1054793)

To make changes in the registry keys

1. Bring the Exchange service group online.
2. Click **Start** and then click **Run**.
3. In the dialog box, enter regedit and click **OK**.
4. In the Registry Editor, locate the following subkey in the registry:

   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\VirusScan
5. In the right pane, double-click **Enabled**.
6. Click **Decimal**, enter 0, and then click **OK**.
7. On the **File** menu, click **Exit** to quit Registry Editor.

Error while performing Exchange post-installation steps

After installing Exchange and rebooting the node, the Veritas High Availability Engine (HAD) may fail to start. As a result, while performing the Exchange post-installation tasks, the Exchange Setup Wizard may either fail to launch or may display the following error message:

```
Failed to get the cluster information. Make sure that VCS
Engine (HAD) is in running state. Start HAD and click Retry
to continue. Click Cancel to exit the wizard.
Error V-16-13-4207
```

This issue may occur in a secure cluster environment. (1211491)

To work around this issue

1 Restart the Veritas High Availability Engine (HAD).

Type the following at the command prompt:

```
hastop -local -force hastart
```

2 Verify that HAD is running.

Type the following at the command prompt:

```
hasys -state
```

The state should display as RUNNING.

3 Click **Retry** on the Exchange Setup Wizard panel and proceed with the Exchange post-installation steps.

Exchange Best Practices Analyzer health check fails with server name mismatch error

If you try to use the Exchange Server 2007 Best Practices Analyzer to run a Health Check on a clustered Mailbox server, the test fails with a server name mismatch error, as in the following example:

```
Server name mismatch Server: EVS1
There is a discrepancy with the Exchange server name EVS1 between
the registry and Active Directory. The computer may have been
renamed or third-party clustering software may be running. Host name
in registry: CNODE1.
```

You can ignore this error and proceed with the analysis. (1166985)

Exchange Setup Wizard does not allow a node to be rebuilt and fails during installation

The Exchange Setup Wizard does not allow a node to be rebuilt, and fails during installation. This is because the wizard stores all the information about the Exchange Virtual servers (EVS) that can fail over on a node, in the ExchConfig registry hive. The path in the registry hive is HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VCS\EXCHCONFIG.

Even if any of the failover nodes die, the corresponding entry still exists in the system list of the EVS. During installation, the Exchange Setup Wizard refers to this incorrect registry entry and fails. (256740)

Workaround: You will have to manually remove the registry entries of the nodes that are being rebuilt, from the system list of the Exchange virtual server on all nodes.

Warning: Incorrectly editing the registry may severely damage your system. Before making changes to the registry, make a backup copy.

To work around this issue

1. To open the Registry Editor, click Start > Run, type regedit, and then click OK.
2. In the registry tree (on the left), navigate to HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VCS\EXCHCONFIG
3. From the Exchange Virtual Server keys, delete the keys representing the nodes that are being rebuilt.
4. Repeat steps 1 to 3 for the Exchange virtual server on all the nodes in the cluster.
5. Exit the Registry Editor.

Resource for Exchange Information Store may take time to online

If the Microsoft Exchange database is in an inconsistent state and the enterprise agent for Exchange attempts to bring the resource for Microsoft Exchange Information Store (IS) service online, the IS service runs a recovery on the Exchange database. This recovery may take considerable time, depending on the number of transaction logs to be replayed.

As a default behavior, the enterprise agent for Exchange waits in the Online entry point and returns only when the IS resource starts or when the start operation fails. When IS service is delayed, the enterprise agent for Exchange logs the following message:

```
The Information Store service is not yet started.
It might be running recovery on the database.
```

In some cases, however, the IS service may not be running a recovery.

Workaround: If the IS service is stuck in the STARTING state, you can force the Online entry point to exit without waiting for IS service to start.

To work around this issue

1. Open the Registry Editor.
2. From the navigation pane, go to

   ```
   \\hkey_local_machine\software\veritas\vcs\
   exchconfig\parameters\msexchangeis
   ```

3. On the **Edit** menu, select **New**, and then click **DWORD Value**.

4 Name the value ForceExit.

5 Right-click the value and select **Modify**.

6 In the **Edit DWORD Value** dialog box, specify the value data as '1'.

Click **OK**

When the Online routine detects this value in the registry, it exits without waiting for the IS resource to start.

Note: To restore the default behavior of the agent, set the ForceExit value to zero.

VCS with Oracle

The following issues apply to VCS with Oracle

Oracle Enterprise Manager cannot be used for database control

In this release, you cannot use Oracle Enterprise Manager for database control. (364982)

For more information, see TechNote 277440 at:

http://seer.entsupport.symantec.com/docs/277440.htm

VCS Hardware Replication Agent for EMC MirrorView

The following issues apply to VCS Hardware Replication Agent for EMC MirrorView.

MirrorView resource cannot be brought online because of invalid security file

If a configured MirrorView resource cannot be brought online successfully, the problem may be an invalid security file. Review the steps for executing the addArrayuser action in the *Veritas Cluster Server Hardware Replication Agent for EMC MirrorView Configuration Guide* and verify that the steps were followed correctly. If you did not specify a password as an Action Argument when executing the addArrayUser action, an invalid security file for the SYSTEM user is created on the local and remote arrays. Executing the addArrayuser action again with a valid password does not overwrite the invalid security file.

To resolve this issue, you must modify the `addArrayUser.pl` action script and re-execute it to remove the invalid security file. The `addArrayUser.pl` script is located in the directory, `%ProgramFiles%\Veritas\cluster server\bin\MirrorView\actions`.

Make a copy of the original addArrayUser.pl script before you make any changes to the script.

The following procedure removes the security file created for the SYSTEM user. (769418)

To remove the security file created for the SYSTEM user

1 In the `addArrayUser.pl` script, replace the line:

```
my $cmd = "\"" . $java_home . "\\java\" -jar \"" . $NaviCliHome
. "\\navicli.jar\" -h " . $LocalArraySPNames[$i] . "
-AddUserSecurity -Password $arrayPasswd -Scope 0";
```

with the line:

```
my $cmd = "\"" . $java_home . "\\java\" -jar \"" . $NaviCliHome
. "\\navicli.jar\" -h " . $LocalArraySPNames[$i] . "
-RemoveUserSecurity";
```

2 In the addArrayUser.pl script, replace the line:

```
my $cmd = "\"" . $java_home . "\\java\" -jar \"" . $NaviCliHome
. "\\navicli.jar\" -h " . $RemoteArraySPNames[$i] . "
-AddUserSecurity -Password $arrayPasswd -Scope 0";
```

with the line:

```
my $cmd = "\"" . $java_home . "\\java\" -jar \"" . $NaviCliHome
. "\\navicli.jar\" -h " . $RemoteArraySPNames[$i] . "
-RemoveUserSecurity";
```

3 After you have modified the `addArrayUser.pl` script, save the changes.

4 Execute the addArrayUser action to remove the invalid security file. Consult the *Veritas Cluster Server Hardware Replication Agent for EMC MirrorView Configuration Guide* for more details on executing the addArrayUser action. You do not need to specify an Action Argument.

5 The action should complete successfully. If an error is returned, verify that the changes to the `addArrayUser.pl` script were made correctly and verify that the script is in the correct location.

6 After the invalid security file has been removed, revert the modified `addArrayUser.pl` script back to the original script, and follow the procedure for executing the addArrayUser action again.

Disaster Recovery Configuration Wizard

The following are Disaster Recovery Configuration Wizard issues.

The DR Wizard does not provide a separate “GCO only” option for VVR-based replication

The Disaster Recovery Configuration Wizard provides a “GCO only” option for hardware array-based replication only, not for VVR-based replication. If this option is selected, before proceeding to GCO configuration, the wizard creates a storage and service group configuration intended for use in hardware array-based replication and incorrect for a VVR configuration. For VVR replication you should instead choose the option to configure both VVR Replication and GCO. (1184660)

If you do not want the wizard to configure the VVR replication but only GCO, you do the following:

To configure GCO only

1. Select the option **Configure Veritas Volume Replicator (VVR)** and the **Global Cluster Option (GCO)**
2. Exit the wizard after configuring the service group.
3. Configure VVR replication without using the wizard.
4. Restart the wizard and select the same VVR and GCO replication option.

   The wizard recognizes that the VVR replication settings are complete, and enables you to proceed to GCO configuration.

The Disaster Recovery Wizard fails if the primary and secondary sites are in different domains or if you run the wizard from another domain

The Disaster Recovery Wizard requires that the primary and secondary sites be in the same domain. In addition, you must launch the wizard from within the same domain as the primary and secondary sites.

Otherwise, when you select the secondary site system, the wizard returns the error that it was unable to perform the operation and that it failed to discover Veritas Cluster Server. (853259)

The Disaster Recovery Wizard may fail to bring the RVGPrimary resources online

During the final stage of disaster recovery configuration with the Disaster Recovery Wizard, the last action is to bring the RVGPrimary resources online. In some cases, the wizard displays an error on its final panel and notifies you to bring the resources online manually. (892503)

Workaround: Use the Cluster Manager (Java console) to manually bring online the RVGPrimary resources of the selected application service group and any dependent group.

The Disaster Recovery Wizard requires that an existing storage layout for an application on a secondary site matches the primary site layout

The Disaster Recovery Configuration Wizard is designed to use for a new installation on the secondary site. Because it clones the storage, you do not need to configure the storage at the secondary site.

If you configure disk groups and volumes at the secondary site and install the application before you run the Disaster Recovery Wizard, the following limitations apply:

The wizard recognizes the storage at the secondary site only if it exactly matches the layout on the primary site. If there is a mismatch in volume sizes, the wizard can correct this. Otherwise, if the layout does not match, the wizard will not recognize that a storage layout already exists.(781923)

If it doesn't find a matching storage layout, the wizard will clone the storage from the primary site, if there is enough disk space. The result is two sets of disk groups and volumes:

- The set of disk groups and volumes that you created earlier
- The different set of disk groups and volumes that the wizard created by cloning the primary storage configuration

Workaround: If you have already created the storage layout at the secondary site and installed the application, use the Disaster Recovery Wizard only if the layout exactly matches the layout on the Primary site.

Otherwise, if the wizard creates a different set of disk groups and volumes than what you have created earlier, you must set up the application to use the disk groups and volumes created by the Disaster Recovery Wizard before you can continue with the wizard.

The Disaster Recovery Wizard may fail to create the Secondary Replicator Log (SRL) volume

If the VMDg resource is not online on the selected Secondary system, the Disaster Recovery Wizard fails to create the SRL volume. This can occur if the disk group for the selected service group has not been imported on the selected secondary system so that the VMDg resource is not online. (896581)

Workaround: Exit the wizard. Bring the VMDg resource for the selected service group online at the secondary node where you are configuring replication. Then run the Disaster Recovery Wizard again.

The Disaster Recovery Wizard may display a failed to discover NIC error on the Secondary system selection page

The Disaster Recovery Wizard may display a failed to discover NIC error on the secondary system selection page. This can occur if it encounters a problem with the Windows Management Instrumentation (WMI) service on one of the cluster nodes. (893918)

Workaround: Exit the wizard and check if the Windows Management Instrumentation (WMI) service is running on the node identified in the error message. If not, start the service and restart the wizard.

If the error repeats, you can troubleshoot further by checking if there is a problem with the WMI repository on the node. To check for problems, use the WMI test program `wbemtest.exe` to enumerate instances of Win32_NetworkAdapterConfiguration and Win32_NetworkAdapter. If they do not enumerate successfully, fix the problem with the WMI repository before restarting the wizard.

Service group cloning fails if you save and close the configuration in the Java Console while cloning is in progress

While the Disaster Recovery Wizard is cloning the service group, if you save and close the configuration in the Java Console while cloning is still in progress, the cloning fails with an error. (1216201)

Workaround: Delete the service group on the secondary site. Run the wizard again to clone the service group.

If RVGs are created manually with mismatched names, the DR Wizard does not recognize the RVG on the secondary site and attempts to create the secondary RVG

The Disaster Recovery Wizard configures VVR replication for you. However, if you choose to configure the replication outside of the DR Wizard, ensure that you use the same names for the RDS and RVG on both sites. Otherwise, if the secondary site has a different RVG name than the primary, when you run the wizard, the wizard finds the primary site RVG information but does not recognize the misnamed secondary site RVG. On the replication action page, creation of the secondary RVG fails. (1214003)

Workaround: Rename the misnamed RVG on the secondary site to match the primary site. You can run the wizard again and continue with GCO configuration. Refer to the *Veritas Volume Replicator Administrator's Guide* for more information on implementing VVR manually.

Cloned service group faults and fails over to another node during DR Wizard execution resulting in errors

After service group cloning is complete, a resource fault may occur in the service group on the secondary site, causing the cloned service group to fault and fail over to the other cluster node. As a result, when the wizard proceeds to the replication Implementation stage, implementation actions may fail because the resource is online on the other node. (1177650)

Workaround: If you discover that the cloned service group has failed over to another node resulting in any failure of the actions shown on the wizard Implementation page, delete the cloned service group completely and run the DR Wizard again.

DR wizard may display database constraint exception error after storage validation in EMC SRDF environment

The DR wizard storage validation on the secondary site may result in a constraint exception error (duplicate database objects) shown on the Storage Validation page of the wizard. This error can occur because the array information and the Volume Manager information cached in the VEA are not in synch. This error is most likely to happen in an EMC SRDF environment. Rescanning the storage on the secondary node to update the Volume Manager information can often resolve this error. (1127959)

Workaround: Check the storage configuration on the secondary site for any errors. Using the VEA, rescan the storage on the secondary node on which the error occurred.

DR wizard creation of secondary RVGs may fail due to mounted volumes being locked by the OS

This volume lock issue can result in the DR wizard failing to create secondary RVGs. This issue is more likely to occur if there are many disk groups and volumes in the configuration. In such a case the wizard may successfully complete configuring some but not all RVGs. If the wizard is then run again to complete the RVG configuration, the wizard is unable to complete setting up the RLINKs for the RVGs that were configured earlier. (1299615)

Workaround: Offline all mountV resources at the secondary site before using the wizard to configure replication and GCO. If a failure occurs while configuring secondary RVGs, delete any existing secondary site RVGs before you re-run the wizard.

DR wizard with VVR replication requires configuring the preferred network setting in VEA

The DR wizard passes the host name rather than an IP address for the secondary host. By default VVR will attempt to resolve the host name using the IPv4 protocol. In this case, if the primary site is IPv6, the secondary replicated volume group (RVG) configuration will fail. (2515518)

Workaround: To ensure that VVR uses the correct protocol to resolve host names, use Veritas Enterprise Administrator (VEA) (Control Panel > VVR Configuration > IP Settings tab) to specify the IP preference before you run the wizard. The default setting is IPv4.

DR wizard displays error message on failure to attach DCM logs for VVR replication

In a configuration with a large number of disk groups and a large number of volumes using VVR replication, the DR wizard may display an error message on the implementation page. (2576420)

The message displayed is as follows:

```
The wizard failed to attach DCM log on primary RVGs
```

However, the wizard continues with implementation steps and the DCM log operation eventually completes. No further action is required.

Fire Drill Wizard

The following are Fire Drill Wizard issues.

Fire drill may fail if run again after a restore without exiting the wizard first

After using the Fire Drill wizard to run a fire drill and restore the fire drill configuration, you then try to run the fire drill again without exiting the wizard. The MountV fire drill resources may fail to come online and the fire drill may fail. (2563919)

Workaround: To run a fire drill again after restoring the configuration, restart the wizard first, or run the fire drill in a new instance of the wizard.

Fire Drill Wizard may fail to recognize that a volume fits on a disk if the same disk is being used for another volume

When using the Fire Drill Wizard to prepare the fire drill configuration, you can assign disks for the snapshot volumes. If you assign more than one volume to the same disk, the Fire Drill Wizard requires that the disk size be large enough to accommodate the size of both volumes combined, even if one of the volumes is being assigned to another disk as well. For example, if you have a 10-GB volume

assigned to disk A and disk B, and a 5-GB volume assigned to disk B, the Fire Drill Wizard only allows this assignment if disk B has at least 15 GB free. (893398)

Workaround: Assign volumes to separate disks or ensure that if more than one volume is assigned to a disk then it is large enough to accommodate all the volumes assigned.

Fire Drill Wizard may time out before completing fire drill service group configuration

In some larger application service group configurations with many resources, the Fire Drill Wizard may time out before it is able to complete the fire drill service group configuration. (1296532)

Workaround: The default value for the wizard time-out is 600000 milliseconds, the equivalent of 10 minutes. If the wizard times out, you can reset the default time value in the Windows registry to a longer time, for example to 20 minutes.

Modify the following registry setting:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\winsolutions\TimeLimit
```

RegRep resource may fault while bringing the fire drill service group online during "Run Fire Drill" operation

Occasionally, when you run a fire drill using the Fire Drill Wizard, the RegRep resource faults and the fire drill service group fails to come online. This occurs due to a VCS error (V-16-10051-5508).

To work around this issue

1. Stop the Fire Drill Wizard.
2. In the VCS Java Console, bring the fire drill service group offline.
3. In the fire drill service group, bring online the MountV resource on which the RegRep resource depends.
4. Copy the contents of the primary RegRep volume to the secondary RegRep volume.
5. Bring online the entire fire drill service group. If no other problem exists, the service group comes online.
6. Run the Fire Drill Wizard again, selecting the **Restore to Prepared State** option. You can then select the **Run Fire Drill** option to run the fire drill again.
7. Proceed as during a normal run of the Fire Drill Wizard.

Fire Drill Wizard in an HTC environment is untested in a configuration that uses the same horcm file for both regular and snapshot replication

In a Hitachi TrueCopy hardware replication environment, the Fire Drill Wizard has only been tested using two separate `horcm` files on the secondary site for the snapshot replication. (1703762)

In other words, it has been tested in a configuration with four `horcm` files as follows:

- Two matching `horcm` files on the primary and secondary site used for replication
  For example, a `horcm10.conf` on the primary site and an identical `horcm10.conf` on the secondary site
- Two additional `horcm` files (`horcm11.conf` and `horcm12.conf`) on the secondary site, used for the fire drill snapshot replication. The `horcm11.conf` file is the same as `horcm10.conf` except that it uses the secondary site IP address.

This configuration has been tested on the following array:

- Hitachi Thunder 9570 (Micro Code version – 065F/D)

The following snapshot configuration has not been tested and therefore results are unknown:

- Two matching `horcm` files (for example, `horcm10.conf`) on the primary and secondary site used for replication
- One additional `horcm` file (`horcm11.conf`) used along with the `horcm10.conf` file on the secondary site for the fire drill snapshot replication

FireDrill attribute is not consistently enabled or disabled

When running the Fire Drill Wizard or when performing a fire drill using the VOM console or Java GUI, the FireDrill attribute is not consistently enabled or disabled. The Run operation of the Fire Drill Wizard enables the FireDrill type-level attribute when bringing a resource online, and disables it when taking a resource offline. However, the VOM console or Java GUI cannot change the value of this attribute. (2735936)

Therefore, if you use the VOM console or Java GUI to run a fire drill, make sure that you do the following for all the resource types other than IP, Lanman, and VMDg:

- Before you run a fire drill, set the FireDrill attribute to `TRUE`.
- After you restore a fire drill, set the FireDrill attribute to `FALSE`.

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

MountV resource state incorrectly set to UNKNOWN

When a fire drill service group comes online, the MountV resource of the corresponding application service group goes into the UNKNOWN state. After the fire drill service group goes offline, the UNKNOWN state of the MountV resource is cleared. (2697952)

Remember to restore the fire drill configuration immediately after you perform a fire drill.

Other issues

The following are other issues.

Resources in a parent service group may fail to come online if the AutoStart attribute for the resources is set to zero

This issue occurs with service groups in a parent-child relationship linked with online local firm dependency and when the AutoStart attribute for all the resources of the parent service group is set to 0 (false). The AutoStart attribute of the parent service group is set to 1 (true).

If you take the parent service group resources offline and then switch or fail over the child service group to another node in the cluster, the child service group comes online on the node but the parent service group resources do not come online on that node. (1363503)

The following error is displayed in the parent service group's resource logs:

```
VCS WARNING V-16-1-10285 Cannot online: resource's group is frozen
waiting for dependency to be satisfied
```

Workaround: In such a scenario, while taking the parent service group resources offline, use the following command for the last resource:

```
hagrp -offline service_group -sys system_name -clus cluster_name
```

Here, *service_group* is the name of the parent service group.

You can also take the resource offline using the Cluster Manager (Java Console). This action ensures that the parent service group resources come online on the node on which the child service group is switched or failed over.

VCS wizards may fail to probe resources

While creating resources and service groups using VCS wizards, if you choose to bring the resources or service groups online, the wizards may fail to probe the resources. (1318552)

The following error is displayed:

```
Failed to online <resourcename> on system <nodename>
Resource has not been probed on system <nodename>
```

Workaround: In such cases, complete the wizards and then probe the resource manually from the Cluster Manager (Java console) and then bring it online.

Backup Exec 12 installation fails in a VCS environment

If you try to install Backup Exec 12 on systems where VCS is already configured, the installation may fail. This failure happens on 64-bit systems. (1283094)

Workaround: Stop the Veritas High Availability Engine (HAD) on all the cluster nodes and then proceed with the Backup Exec installation.

Changes to referenced attributes do not propagate

This behavior applies to resources referencing attributes of other resources; that is, the ArgList of one resource (A) passes an attribute of another resource (B). If resource B is deleted from the group, or if the SystemList of the group containing resource B does not contain a system defined in the SystemList of the group containing resource A, the VCS engine does not propagate these changes to the agent monitoring resource A. This failure to propagate the changes may cause resource A to fault because it does not receive the appropriate attribute values from resource B.

In such situations, you must reset the value of resource B in the attribute definition of resource A or restart the agent managing resource A.

For example, the ArgList of the MountV resource contains the DiskGroupName attribute of the VMDg resource. If you change the VMDg resource name or the SystemList, the VCS engine does not communicate the change to the MountV agent, causing it to fault. In such a situation, you can reconfigure the MountV agent using one of the following methods:

- Refresh the VMDgResName attribute for the MountV resource. Set the attribute to an empty string "" first, then reset it to the new VMDg resource name.
- Stop and restart the MountV agent on the system.

ArgListValue attribute may not display updated values

When you modify a resource type that has localizable attributes, the agent log warns that ArgListValues cannot be localized. You can safely ignore the warning message about ArgListValues.

After you modify values for a resource that has localizable attributes, the command hares -display does not display the updated ArgListValues.

Known behavior with disk configuration in campus clusters

The campus cluster configuration has the same number of disks on both sites and each site contains one plex of every volume. Note that an environment with an uneven number of disks in each site does not qualify as a campus cluster.

If a site failure occurs in a two-site campus cluster, half the disks are lost. The following cases may occur:

- The site in which the service group is not online fails.
- The site in which the service group is online fails.

The behavior and possible workarounds for these conditions vary.

AutoStart may violate limits and prerequisites Load Policy

The load failover policy of Service Group Workload Management may be violated during AutoStart when all of the following conditions are met:

- More than one autostart group uses the same Prerequisites.
- One group, G2, is already online on a node outside of VCS control, and the other group, G1, is offline when VCS is started on the node.
- The offline group is probed before the online group is probed.

In this scenario, VCS may choose the node where group G2 is online as the AutoStart node for group G1 even though the Prerequisites load policy for group G1 is not satisfied on that node.

Workaround: Persistently freeze all groups that share the same Prerequisites before using `hastop -local -force` command to stop the cluster or node where any such group is online. This workaround is not required if the cluster or node is stopped without the force option.

VCS Simulator installation may require a reboot

While installing the VCS Simulator, the installer may display a message requesting you to reboot the computer to complete the installation. Typically, a reboot is required only in cases where you are reinstalling the VCS Simulator. (851154)

Unable to output correct results for Japanese commands

When the Veritas Command Server starts up on a Windows setup, it runs as a Windows service on a local system. A Windows service generally runs in the same locale as the base Operating System's locale, and not the systems locale. For example, if a system is running an English version of Windows with a Japanese locale, then the CmdServer service runs in an English locale and not Japanese. Thus, when user commands are issued in Japanese the command server is confused

when performing the Uniform Transformation Format (UTF) conversions and is unable to output the correct results. (255100)

Configuration wizards do not allow modifying IPv4 address to IPv6

The VCS Cluster Configuration Wizard (VCW) and the service group configuration wizards do not allow you to modify IPv4 addresses of resources in an existing service group to IPv6. (2405751)

Workaround: You can work around this issue in one of the following ways.

- Use the appropriate wizard to delete the service group and create it again using resources with IPv6 addresses.
- Use either the Java GUI or CLI to replace the IPv4 resources in the service group with corresponding IPv6 resources.

Veritas Volume Replicator

This section provides information for Veritas Volume Replicator known issues.

VVR replication may fail if Symantec Endpoint Protection (SEP) version 12.1 is installed

SEP may block VVR replication if the replication packet size is set to greater than 1300 bytes. (2598692)

Workaround:

- Ensure that the required ports and services are not blocked by the firewall. Refer to the Installation and Upgrade Guide for list of ports and services used by SFW HA.
- Configure the SEP firewall to allow IP traffic on the systems. On the SEP client's Network Threat Protection Settings dialog box, check **Allow IP traffic** check box.

RVGPrimary resource fails to come online if VCS engine debug logging is enabled

This issue occurs when trying to bring the RVGPrimary agent resource online when VCS engine debug logging is enabled. This happens because RVGPrimary cannot parse command line output when the debug logging is enabled. However, this issue does not affect the monitoring of the RVGPrimary resource. (2886572)

Workaround: As a workaround, disable debug logs for the VCS engine by deleting the VCS_DEBUG_LOG_TAGS environment variable and its values. Once the

RVGPrimary resource is online, enable the debug logging again by creating the VCS_DEBUG_LOG_TAGS variable with the values that you had set before.

"Invalid Arguments" error while performing the online volume shrink

This issue occurs while performing the online volume shrink operation. While attempting the volume shrink operation, the "Invalid Arguments" error occurs and the Event Viewer displays a Microsoft Virtual Disk Service (VDS) provider failure error. (2405311)

Workaround: To resolve this issue, restart VDS, and then run the `vxassist refresh` command.

`vxassist shrinkby` or `vxassist querymax` operation fails with "Invalid Arguments"

This issue occurs while performing the `vxassist shrinkby` or `vxassist querymax` operation for a newly-created volume. The issue occurs because Veritas VDS Dynamic Provider is not updated properly. The `vxassist shrinkby` or `vxassist querymax` command fails with the "Invalid Arguments" error. (2411143)

Workaround: To resolve this issue, you need to restart the VDS components by performing the following steps using the CLI:

1. `net stop vds`
2. `taskkill /f /im vxvds.exe`
3. `taskkill /f /im vxvdsdyn.exe`
4. `net start vds`
5. `vxassist refresh`

In synchronous mode of replication, file system may incorrectly report volumes as raw and show “scan and fix” dialog box for fast failover configurations

In a fast failover configuration, this issue occurs when replication is configured in hard synchronous mode in Windows Server Failover Clustering and the service group is made offline and online or moved to another node. Since the RLINK is in hard synchronous mode, it may not be connected when the volume arrives after the service group is made offline and online or moved to another node, and the I/Os may fail.

In such cases, the file system may incorrectly report the volume as raw and the “scan and fix” dialog box may appear to help fix the volume’s file system. The Event Viewer may also display NTFS errors. However, please note that the volumes are not corrupted by this issue.

However, this issue would not occur if the FastFailover attribute of the Disk Group resource is set to False. (2561714)

Workaround: To resolve this issue, choose either the synchronous override or asynchronous mode of replication.

VxSAS configuration wizard fails to discover hosts in IPv6 DNS

This issue occurs while configuring the VxSAS service using the VVR Security Service (VxSAS) Configuration Wizard. If the DNS is configured for Internet Protocol Version 6 (IPv6), then the wizard fails to automatically discover hosts in the domain. (2412124)

Workaround: To resolve this issue, manually provide the IP address or name of the host in the wizard.

File system may incorrectly report volumes as raw due to I/O failure

This issue occurs if the Storage Replicator Log (SRL) overflow protection attribute—`srlprot`—is set to "fail", the RLINK is disconnected, and heavy I/O operations are performed that fill up the SRL. Once the SRL becomes full, any further I/O operations to the data volumes that are part of the RVG fails. In such cases, the file system may incorrectly report the volumes as raw. (2587171)

Workaround: To resolve this issue, set the `srlprot` attribute of the corresponding RLINK to “autodcm” or ensure that the RLINK is connected, which will cause the I/Os to flow to Secondary and reduce the SRL usage.

NTFS errors are displayed in Event Viewer if fast-failover DR setup is configured with VVR

This issue occurs if Disaster Recovery (DR) setup is configured with VVR and you fail over an Application Service Group to remote cluster. The Event Viewer displays NTFS errors long after the MountV and application service resources have successfully offlined the volumes. However, please note that this does not have any impact on the data or failover. (2570604)

Workaround: There is no workaround for this issue.

Volume shrink fails because RLINK cannot resume due to heavy I/Os

This issue occurs while shrinking a data volume. While the volume shrink is in progress, if you perform heavy I/O operations, then the paused RLINK times out and fails to resume, and therefore, the volume shrink operation fails. (2491642)

Workaround: To prevent the timeout, increase the AE_TIMEOUT value as follows:

1. Open the Registry Editor by typing `regedit` in the Run menu.
2. Navigate to the following location:

   `HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager\constants`
3. Modify the registry DWORD value for the AE_TIMEOUT entry, from the default value of 30 seconds to 60 seconds or higher.
4. In order for the registry key change to take effect, type the following at the command prompt:

   ```
   vxassist refresh
   ```

VxSAS configuration wizard doesn't work in NAT environments

This issue occurs while configuring the VxSAS service using the VVR Security Service (VxSAS) Configuration Wizard in a Network Address Translation (NAT) environment. VxSAS uses DCOM APIs that internally use port 135. If this port is not forwarded, then the VxSAS configuration wizard fails. (2356769)

Workaround: To resolve this issue, manually configure the VxSAS service in a NAT environment.

Online volume shrink operation fails for data volumes with multiple Secondaries if I/Os are active

This issue occurs while performing the online volume shrink operation on a data volume that has multiple Secondaries. The volume shrink operation fails in this case if the I/Os are active. (2489745)

Workaround: Ensure that the RLINKs are up-to-date and that there is no application I/O active when the online volume shrink operation is performed.

RLINKs cannot connect after changing the heartbeat port number

This issue occurs when you change the replication heartbeat port number after a Secondary RVG (Replicated Volume Group) is added. Because of this, the existing RLINKs cannot connect. (2355013)

Workaround: To resolve this issue, delete and then add the Secondary RVGs again.

VVR replication fails to start on systems where Symantec Endpoint Protection (SEP) version 12.1 is installed

This issue may occur if VVR replication is set up on systems in an IPv6 environment where Symantec EndPoint Protection (SEP) version 12.1 is installed.

The Replication Status in the VEA GUI displays as “Activating” and the replication may fail to start. (2437087)

Workaround: Ensure that the VVR ports are not blocked by the firewall. Refer to the SFW HA Installation and Upgrade Guide for list of ports and services used by SFW HA.

Configure the SEP Firewall policy to allow IPv6 traffic on the cluster nodes.

Edit the Firewall Rules table and edit the following settings:

- **Block IPv6**: Change the Action field value to “Allow”.
- **Block IPv6 over IPv4 (Teredo)**: Change the Action field value to “Allow”.
- **Block IPv6 over IPv4 (ISATAP)**: Change the Action field value to “Allow”.

Refer to the SEP documentation for detailed instructions on how to edit the firewall policy.

On a DR setup, if Replicated Data Set (RDS) components are browsed for on the secondary site, then the VEA console does not respond

If RDS and RVG items are browsed for on the secondary site on a DR setup, then the Veritas Enterprise Administrator (VEA) console hangs up and does not respond. This is noticed only on the secondary site and not on the primary site.

Workaround: Create reverse lookup entries in the DNS for VVR IP and physical host.

Secondary host is getting removed and added when scheduled sync snapshots are taken

Schedule a synchronized snapshot on the secondary host. It is noticed that sometimes when a synchronized snapshot happens on the secondary, the secondary gets removed and added. This is because the snapshot operation is taking too long to complete. To avoid this, increase the AE_TIMEOUT value in the registry to one minute. Its default value is set to 30 secs. (2010491)

Replication may stop if the disks are write cache enabled

In some hardware configurations, if the standard Windows write back caching is enabled on the Secondary, replication may stop for prolonged time periods. In such cases, update timeout messages appear in the primary system event log. Because the Secondary is slow to complete the disk writes, a timeout occurs on the Primary for acknowledgment for these writes. (343556)

Workaround: Before setting up replication, disable write caching for the disks that are intended to be a part of the RDS. You can configure write caching through Windows Device Manager by right-clicking the disk device under the Device drives node and selecting **Properties** > **Policies**.

Discrepancy in the Replication Time Lag Displayed in VEA and CLI

When the Secondary is paused, you may note a discrepancy in replication time lag reported by the `vxrlink status` command, the Monitor view, and the`vxrlink updates` command. The `vxrlink status` command and the Monitor view display the latest information, while the information displayed by the `vxrlink updates` command is not the latest. (299684)

The vxrlink updates command displays inaccurate values

When the Secondary is paused and is behind the Primary, the `vxrlink updates` command may show inaccurate values. While the Replicator Log is receiving writes, the status displayed remains the same as before the pause. However, if the Replicator Log overflows and the Data Change Map (DCM) are activated, then the `vxrlink updates` command output displays the correct value by which the Secondary is behind. In DCM mode, the Primary reconnects the Secondary RLINK and sends updated information, including the time associated with the last update sequence number on the Primary. (288514)

Some VVR operations may fail to complete in a cluster environment

If an RVG is a part of a VCS cluster and the cluster resource for this RVG exists, then VVR fails the Delete RDS, Delete Secondary RVG, Delete Primary RVG, Disable Data Access, Migrate, or Make Secondary operations with the following error:

```
Cannot complete operation. Remote node closed connection.
```

This is a timing issue. The VVR VRAS module times out before completing the check to determine if the RVGs participating in the operation already have a resource created. (309295, 2603103)

Workaround: To prevent the timeout, make the following change on all cluster nodes of the Primary and Secondary cluster:

To change the timeout value

1 Open the Registry Editor, and then navigate to the following location:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager\constants
```

2 Modify the registry DWORD value for the AE_TIMEOUT entry, from the default value of 30 seconds to 60 seconds or higher.

3 In order for the registry key change to take effect, run the following:

```
vxassist refresh
```

IBC IOCTL Failed Error Message

At times, the `vxibc register` or the `vxibc unregister` command may display the following error message: (496548)

```
Error V-107-58644-932: IBC IOCTL failed
```

Workaround: Verify that you have specified the correct RVG or disk group name with the command.

Pause and Resume commands take a long time to complete

At times, the pause and resume operation can take a long time to complete due to which it appears to be hung. (495192)

Workaround: Wait for some time till the operation completes, or manually disconnect and reconnect the network that is used for communication to enable the operation to complete.

Replication keeps switching between the pause and resume state

In a setup that is configured for Bunker replication, if a failure occurs at the primary site, then the Bunker is used to replay the pending updates to the secondary. Later, when the primary node becomes available again, the Bunker can be deactivated and replication can be started from this original primary to the secondary. However, performing any other intermittent operations such as detaching or attaching the RLINK, before starting replication from the original primary can cause the replication to switch between the pause and resume state. (638842, 633834)

Workaround: Recreate the Secondary RVG.

VEA GUI has problems in adding secondary if all NICs on primary are DHCP enabled

When VEA is connected to the Primary host using "localhost" as the hostname and all the NICs on the primary server have DHCP enabled on them, then the Add Secondary Wizard fails to identify that it is connected to the Primary host and does not proceed further.(860607)

Workaround: To avoid this issue, connect to the Primary host using either the hostname or the IP address of the server.

Pause secondary operation fails when SQLIO is used for I/Os

Pausing replication with checkpoints from the secondary host may fail for heavy I/Os and low-bandwidth network. If the secondary's request for RLINK checkpoint for Pause to primary times out before the primary's acknowledgment to the request, the pause operation would fail. (1278144)

Workaround: To avoid this issue, perform one of the following procedures:

- Pause the secondary RVG by selecting the Pause Secondary option from the secondary RVG right-click menu. If it fails, slow down the I/O to the Primary host and retry. Secondary initiated pause lets you specify a checkpoint and maintains the connection between Primary and Secondary.
- Select the Pause Secondaries from Primary option from the Primary RVG right-click menu. If it succeeds, it can be used instead of using the pause replication from the Secondary host. In a Primary initiated pause, the Secondary host gets disconnected and checkpoints cannot be specified.

Performance counter cannot be started for VVR remote hosts in perfmon GUI

Performance monitoring cannot be started if the file is saved under **Performance Logs and Alerts** > **Counter Logs**. (1284771)

Performance counters can be started as follows:

To start performance monitoring

1. To start the file from the details pane, right-click and select the **Properties** dialog box. Then, select the **General** > **Run As** option.
2. In the **Run As** text box enter a username that has administrative privileges on your local computer. Select the **Set Password** tab to enter the password. If your computer is connected to a domain, then use the Domain Admin Group privileges.

VVR Graphs get distorted when bandwidth value limit is set very high

When bandwidth value is set to a very high value, VVR graphs get distorted. (1801004)

BSOD seen on a Hyper-V setup

When a virtual machine resource group is failed over, BSOD is noticed on Hyper-V. (1840069)

Workaround: Run the cluster tunable command as shown. Symantec recommends that you set the value of x to 1:

```
cluster/cluster:clustername/prop HangRecoveryAction=x
```

Here x can take the following values:

- 0=disables the heartbeat and monitoring mechanism.
- 1= logs an event in the system log of the Event Viewer.
- 2=terminates the cluster services.
- 3=causes a Stop error (Bugcheck) on the cluster node.

Unable to start statistics collection for VVR Memory and VVR remote hosts object in Perfmon

Using Perfmon's alerts and counter logs, try to create a new log by selecting VVR memory or VVR remote hosts as objects. The log gets created; however, when we try to start statistics collection by selecting the log, it does not start. (1670543)

Workaround: Use Perfmon's System Monitor page to directly add the VVR counters.

Bunker primary fails to respond when trying to perform stop replication operation on secondary

If there are pending writes along with IBC messages on a bunker host that has multiple secondaries, then while replaying the pending writes from the bunker to the secondary site, the bunker host can experience a hang-like situation. (1544680)

Workaround: Stop replication from the bunker host and either do a takeover on the secondary or synchronize with the existing primary by restarting replication.

Fixes and enhancements for 6.0.1

Fixed issues and software enhancements are referenced by Symantec incident number and described briefly below for the 6.0.1 release.

Installation, upgrades, and licensing

This section lists the fixes and enhancements related to installation, upgrade, and licensing for 6.0.1.

Table 1-2 Installation, upgrades, and licensing fixed incidents

Fixed Incidents	Description
2578521	Clearing the default selection does not remove the DSMs
2605960	Issues due to remote installation of client only components
1862627	FlashSnap License error message appears in the Application Event log after installing license key
774442, 2612475	VMDg resource may not fail over in a cluster environment

Veritas Storage Foundation

This section lists the fixes and enhancements in Storage Foundation for Windows for 6.0.1.

Table 1-3 Storage Foundation for Windows (SFW) fixed incidents

Fixed Incidents	Description
2293995	Though disk partition gets created successfully, a failure message is displayed in the Event Viewer
2585792	VEA console may display timeout error
2482133	Volume labels may disappear for fast failover-enabled disk groups
2534592	Drive letters are getting reassigned when Volume Manager resource comes online even when drive letters are removed from the Veritas Enterprise Administrator (VEA) console
2115125, 2126918	Snapshot operation might fail with "Invalid Arguments" Error

Table 1-3 Storage Foundation for Windows (SFW) fixed incidents *(continued)*

Fixed Incidents	Description
1878059	On a Windows Server 2008 R2 x64 setup, after destroying a Storage Foundation for Windows (SFW) disk group, cannot create a Logical Disk Manager (LDM) disk group on the same disks
2028805	Sometimes VSS snapshot schedules are not visible under the Exchange node in the VEA console tree-view but are visible on the right pane of the console
2028846	If a VCS resource database is configured for an Exchange 2010 mailbox database, the same database node appears in the VEA console tree-view more than once
999290, 999316	Scheduled snapshots do not occur at correct times after installing patch for Daylight Savings Time
2799388	DMP doesn't use all paths of an active-active array in a round robin load balance policy
2774415	The "bogus path" system event log issue when a path failover occurs on an EMC Symmetrix data disk where the data disk shares only a subset of paths with an EMC Symmetrix gate keeper disk

Veritas Volume Replicator

This section lists the fixes and enhancements in Veritas Volume Replicator for 6.0.1.

Table 1-4 Veritas Volume Replicator (VVR) fixed incidents

Fixed Incidents	Description
2534585	RVG resource incorrectly displays all volumes in Failover Cluster Manager

Veritas Cluster Server

This section lists the fixes and enhancements in Veritas Cluster Server for 6.0.1.

Table 1-5 Veritas Cluster Server (VCS) fixed incidents

Fixed Incidents	Description
1466183	Memory leak occurs in the VCS agent for SQL 2008
2867806	The application configured in a mixed-mode environment fails to start

Fixes and enhancements for 6.0

Fixed issues and software enhancements are referenced by Symantec incident number and described briefly below for the 6.0 release.

Veritas Storage Foundation

Table 1-6 describes the Veritas Storage Foundation for Windows (SFW) issues that were fixed in this release.

Table 1-6 Storage Foundation for Windows (SFW) fixed incidents

Fixed Incidents	Description
814881	Shrinking an NTFS volume that is greater than 2 TB is not supported
1863910	Copy On Write (COW) snapshots are automatically deleted after shrink volume operation
1093454	Certain operations on a dynamic volume cause a warning SFW volume operations fail on Windows Server 2008
1592758	Shadow storage settings for a Copy On Write (COW) snapshot persist after shrinking target volume
2138005	In a disk group with Thin Provisioning (TP) disks, deporting forcibly while reclaiming at the disk group level causes the reclaim operation to never return back.
2203640	While creating a disaster recovery configuration, the Disaster Recovery Configuration Wizard fails to discover the cluster node at the primary site where the service group is online.
2207263	Disk group deport operation hangs due to deadlock situation in the storage agent. The vds provider makes PRcall to other providers after acquiring the Veritas Enterprise Administrator (VEA) database lock.

Table 1-6 Storage Foundation for Windows (SFW) fixed incidents *(continued)*

Fixed Incidents	Description
2245816	Windows Operating System by default fails any sector level reads/writes greater than 32MB on Windows Server 2003. Hence, IOs are split into multiple IO for the WriteSector function in fsys provider.
2267265	Issue 1: Attempting to bring the VMDg resource online during MoveGroup operation on a Windows Failover Cluster (WFC) results in an RHS deadlock timeout error: RHS] RhsCall::DeadlockMonitor: Call ONLINERESOURCE timed out for resource. Issue 2: Distributed File System Replication (DFSR) in a Microsoft cluster is not working properly with the Volume Manager Disk Group (VMDg) resource.
2087139	Hotfix installer fails to perform prerequisite operations intermittently. Failure occurs when the installer tries to stop the vxvm service and reports error. When QUERY VXVM is used to query the state of service, status is shown as stopped. While stopping the VXVM service, iSCSI and Scheduler providers perform certain operations which result in failure.
2290214	In a clustering environment, a VxBridge client process may either crash or there could be a memory corruption. This occurs because VxBridge.exe tries to read beyond the memory allocated to a [in,out,string] parameter by its client.
2321015	The bug check 0x3B may happen when removing a disk of a cluster dynamic disk group.
2330902	VxVDS refresh operation interferes with disk group import/deport operation resulting in timeout and delay.
2318276	In a NAT environment, VVR sends heartbeats to the remote node only if the local IP address mentioned on the RLINK is online on that node.
2364591, 2371250	Issue 1: Storage Agent crashes while releasing memory for the buffer passed to collect information. Issue 2: Mirror creation failed with auto selection of disk and track alignment enabled, even though enough space was there.
2368399	Volume shrink operation may cause file system corruption and data loss.

Table 1-6 Storage Foundation for Windows (SFW) fixed incidents *(continued)*

Fixed Incidents	Description
2397382	VVR primary hangs and a BSOD is seen on the secondary when stop and pause replication are performed on the configured RVGs.
2406683, 2329130, 2258124, 2244093, 2114928, 2225878	Issue 1: Mirror creation failed with auto selection of disk and track alignment enabled, even though enough space was there. Issue 2: When using GUI or CLI to create a new volume with a mirror and DRL logging included the DRL log is track aligned and the data volume is no longer track aligned. Issue 3: On a VVR-GCO configuration, when the primary site goes down, the application service group fails over to the secondary site. It is observed that MountV resource probes online on a failed node after a successful auto takeover operation. Issue 4: On a SFW DMP DSM environment with SFW SCSI-3 settings, reconnecting a detached disk of an imported cluster disk group can cause the SCSI-3 release reservation logic to get into a loop or operation may take a long time to get completed. Issue 5: Data corruption while performing subdisk Move operation. Issue 6: During Windows Failover Cluster Move Group operation, cluster disk group import fails with Error 3 (DG_FAIL_NO_MAJORITY 0x0003).
2218963	If the sub disks created on a disk are moved and are not aligned to 8 bytes, then there is a possibility of missing some disk blocks while syncing with the new location. This may result in data corruption.
2426197	Disks fail to appear after the storage paths are reconnected, until either the vxvm service is restarted or the system is rebooted.
2440099	In case the original SFW product license fails, SFW now tries to find license with a different API.
2477520, 2210586	Installing cumulative patch 1 (CP1) over SFW 5.1 SP2, does not allow to configure dynamic cluster quorum in Microsoft Cluster.
2376010	In a split brain scenario, the active node (defender node) and the passive nodes (challenger nodes) try to gain control over the majority of the disks. If the disk group contains a large number of disks, then it takes a considerable amount of time for the defender node to the reserve the disks. The delay may result in the defender node losing the reservation to the challenger nodes.

Table 1-6 Storage Foundation for Windows (SFW) fixed incidents *(continued)*

Fixed Incidents	Description
2415517, 2270478	VVR Primary goes into a hang state while initiating a connection request to the Secondary. This is observed when the Secondary machine returns an error message which the Primary is unable to understand.
2489881	Thin Provisioning Reclaim support for HP P9500 array
2512482	SFW disables the automount feature on a system which leaves the default system volume offline after each reboot. Cluser validation checks for access to all volumes and fails for the offline system volume with error 87.
2372049	Unable to create enclosures for SUN Storage Tek (STK) 6580/6780 array.
2400260	If a volume or any of its snapshot is reclaimed, then performing the snapback operation on such a volume causes data corruption.
2536009	Volumes can be mounted externally even after the MountV offline is complete.
2554039	After importing a disk group, a rescan fired from the ddlprov.dll when the VDID for a device changes, appears in the VEA task bar.
2536342	Dynamic disk groups created using GPT disks have the same signature and ID. This causes an issue where you cannot use MSCVMM to deploy additional virtual machines on these disk groups.
2564914	VxVDS.exe process does not release handles post CP2 Updates.
2372164	Multiple buffer overflows in SFW vxsvc.exe causes vulnerability that allows remote attackers to execute arbitrary code on vulnerable installations of Symantec Veritas Storage Foundation. Authentication is not required to exploit this vulnerability.

Veritas Cluster Server

This section lists the fixes and enhancements in Veritas Cluster Server for 6.0.

Table 1-7 describes the Veritas Cluster Server (VCS) issues that were fixed in 6.0.

Table 1-7 Veritas Cluster Server (VCS) fixed incidents

Fixed Incidents	Description
1876562	Unable to access SQL Server 2008 databases upon failover
1248877, 2479890	Print Share wizard may fail to discover the PrintSpool resource if NetBIOS over TCP/IP is disabled
2220352	File shares on Windows Server 2003 are accessible using the virtual name (Lanman) or the IP address (administrative IP or virtual IP). On Windows Server 2008 systems, due to file share scoping, the file shares configured with VCS are accessible only using the virtual server name (Lanman). These file shares are not accessible using the IP address.
2237219	Before making an IP address usable, Windows server performs a Duplicate Address Detection (DAD) test to check whether the IP address is unique on the network. During this check, the state of the IP address (ipconfig) on the system remains as "Tentative". The VCS IP agent does not wait for the DAD check to complete and reports the status of the IP resource as online, even though the configured IP address is not yet usable.
2231216	In a Global Cluster Configuration (GCO) cluster environment, PrintShare resources come online on the nodes at the primary and secondary site at the same time. Taking the resources offline at the secondary site results in a service group fault at the primary site. Bringing the faulted resources online on the primary site also brings the corresponding resources online at the secondary site leading to a concurrency violation.
2401163, 2251689, 2239785, 2482820, 2392336	Issue 1: The VCS High Availability Engine (HAD) fails to acknowledge resource state changes. As a result, service group resources do not go offline or fail over and remain in the waiting state forever. Issue 2: In a online global soft group dependency, if more than one parent service group is ONLINE in a cluster, then you cannot switch a child service group from one node to another. Issue 3: The VCS High Availability Engine (HAD) rejects client connection requests after renewing its authentication credentials as per the CredRenewFrequency cluster attribute. Issue 4: The HA commands spike CPU usage to almost 100% for about 10 seconds or more. Also there is a delay in getting the results from the HA commands.

Table 1-7 Veritas Cluster Server (VCS) fixed incidents *(continued)*

Fixed Incidents	Description
2497664	Exchange 2010 database resources fail to come online in first attempt because Exchange Active Manager takes time to update.
2495109	While configuring a cluster VCW fails to discover the systems due to an access denied error.
2522314	LLT spinlock causes a deadlock on multi-CPU servers because a spinlock is held while sending out packets. This deadlock causes the system to hang.
2513842	The VCS FileShare agent is unable to probe FileShare resources on passive nodes if the fileshare is configured for a root of a volume that is mounted as a folder mount on a folder on a local system drive.
2207404	VCS SRDF agent resource fails to come online if 8.3 short file naming convention is disabled on the cluster node and the EMC SRDF SymHome path name contains spaces.
2271435	VCS EMC MirrorView resource fails to come online in the cluster.
2338437	VCS VMDg agent reboots a cluster node if it is unable to start the Veritas Storage Foundation Service (vxvm). This occurs even if the agent is not configured to reboot the node.
2391453b	The VCS agent for Hitachi TrueCopy fails to detect the correct status of the HTC configuration and thus fails to invoke the "pairresync" command. On a faulted remote node, the horctakeover returns an error with fence level “never”. This causes the VCS agent to freeze the service group in the cluster, and a critical message is logged in the engine log file. The message indicates that a manual recovery operation is required on the array before the service group is failed over to another node in the cluster.
2422907	Issue 1: The VCS MountV resources take a long time to go offline if there are applications accessing the clustered mount points. Issue 2: The VCS MountV resource faults or if the OnlineRetryLimit is set to a non-zero value, the MountV resource takes a long time to come online.
2203640	While creating a disaster recovery configuration, the Disaster Recovery Configuration Wizard fails to discover the cluster node at the primary site where the service group is online.

Table 1-7 Veritas Cluster Server (VCS) fixed incidents *(continued)*

Fixed Incidents	Description
2378712	While configuring replication and global clustering in an EMC SRDF environment, if the DR wizard is run by two different domain users one after the other, then the DR wizard blocks the second user with the SRDF discovery error.
2207263	Disk group deport operation hangs due to deadlock situation in the storage agent. The vds provider makes PRcall to other providers after acquiring the Veritas Enterprise Administrator (VEA) database lock.
2245816	Volume turning RAW due to Write failure in fsys.dll module
2087139	Hotfix installer fails to perform prerequisite operations intermittently. Failure occurs when the installer tries to stop the vxvm service and reports error. When QUERY VXVM is used to query the state of service, status is shown as stopped.
2290214	Memory corruption or a crash in a VxBridge client process in the clustering environment.
2321015	The bug check 0x3B may happen when removing a disk of a cluster dynamic disk group.
2536009	If a global service group that contains multiple volumes is switched to another node either at the local site or to the remote site, some of the MountV resources in the service group fail to come online and may eventually fault.
2530236	rhs.exe crashes on mscsrvgresource.dll with STATUS_ACCESS_VIOLATION(c0000005) while OS booting.