

Introduction and product requirements

This section includes the following topics:

- [About this document](#)
- [Requirements](#)
- [Documentation errata](#)

InfoScale® 9.1 Release Notes - Windows - v9.1

About this document

This document provides the information that is specific to version 9.1 of the InfoScale products.

Review this entire document before using the following products:

- InfoScale Foundation
- InfoScale Storage
- InfoScale Availability
- InfoScale Enterprise

The information in this document supersedes the information that is provided in the product-specific documents.

You can download the latest version of this document from the Arctera Service and Operations Readiness Tools (SORT) website at:

[documents](#)

- The following documents provide further information that is common to all the InfoScale for Windows products:
 - *InfoScale Getting Started Guide*
 - *InfoScale Installation and Upgrade Guide*

- For information about the InfoScale product components and their capabilities, refer to the corresponding administrator's guides and agent guides.
- For information about configuring and administering your applications with the InfoScale products, refer to the application-specific Implementation Guides and Solutions Guides.

Requirements

For information about the operating system, hardware, and other general requirements for the InfoScale products, refer to the *InfoScale Installation and Upgrade Guide*.

For the latest information on the supported hardware and software, see the appropriate compatibility list at:

[home](#)

Click Documentation from the menu on the top, and on the Documentation tab select the appropriate values in the Product, Platform, Documentation Type, and Version filters.

Documentation errata

The information in this section overrides the related information in the product documentation.

The following online help systems have not been updated for this InfoScale release:

- InfoScale Solutions Configuration Center (SCC)
- Storage Foundation for Windows Online Help
- Volume Replicator Online Help
- VxSAS Online Help

For the most recent information, refer to the PDF and HTML versions of the corresponding guides.

Changes introduced in this release

This section includes the following topics:

- Updated release and support policy
- Auto-revival of service group dependency trees
- Enhanced support for gMSA users with clustering and replication configurations
- Support for cloud-based Key Management Service
- Tunable for automatically deleting stale KMS keys
- Tunable for scheduling SecureFS snapshots
- Enhanced support for replication using virtual host names
- Changes in InfoScale Volume Replicator Advisor (VRAdvisor)

Updated release and support policy

InfoScale 9.1 is a Short Term Support (STS) release. For details about the product life cycle and end-of-life policy, refer to the article at: <https://helpcenter.cloud.com/wolken-support/kbsearch-article-100076658>.

Auto-revival of service group dependency trees

The Cluster Server (VCS) component of InfoScale now provides an auto-revival feature for service group dependency trees. This feature can be enabled at the cluster level and at the service group level, both. When this feature is enabled, and a faulted resource comes online, the dependent resources and service groups - whose states were directly affected by the fault, and not due to other conditions or events - are automatically brought online.

The auto-revival of a service group dependency tree works only when the following prerequisites are met:

- The cluster protocol version is 11100 or higher.
- The cluster version is 9.1 or higher.
- The dependencies between the service groups are set to any of the following:

- online local hard
- online local firm
- online remote firm
- online global firm
- online site firm

For details, refer to the *InfoScale 9.1 Cluster Server Administrator's Guide*.

Enhanced support for gMSA users with clustering and replication configurations

A group Managed Service Account (gMSA) is a special type of domain account in Windows Server Active Directory designed to simplify credential management for services running across multiple servers.

With this InfoScale release, the existing support of gMSA users has been enhanced as follows to provide improved security for clustering and replication configurations:

- gMSA integration with the Cluster Configuration Wizard (VCW)
 - VCW now discovers and lists gMSA users on the User Selection page.
 - When selecting a gMSA user for HADHelper, no password entry is required.
 - Cluster Server (VCS) provides the ability to switch between a normal user and a gMSA user for HADHelper using VCW or the HadHelper utility.

For details, refer to the *InfoScale 9.1 Cluster Administrator's Guide*.


- Support for Support for the Volume Replicator (VVR) Security Service (VxSAS) - available from InfoScale 8.0.2 onwards
 - VxSAS can be configured to use gMSA via commands or via the VVR Security Configuration Wizard.
 - gMSA users can be added to the local administrators group and used as service logon accounts.

For details, refer to the *InfoScale 9.1 Volume Replicator Administrator's Guide*.

Support for cloud-based Key Management Service

The existing InfoScale support for volume encryption employs Key Management Interoperability Protocol (KMIP)-based Key Management Service (KMS) servers. It works with InfoScale deployments on-premises as well as on the supported cloud platforms.

In addition, InfoScale now supports volume encryption using cloud-based (non-KMIP-compliant, Software as a Service--SaaS) KMSs. Cloud KMS-based encryption is suitable for environments that support high availability and automated configurations. With this release, InfoScale supports AWS and Azure as cloud KMS providers for volume encryption.



With a cloud-based KMS, you can:

- Avoid managing complex passphrases manually.
- Assign a unique master key to each disk group.
- Recover the master key - within a defined waiting period - after a disk group is removed from the system.
- Use a single cloud KMS key to encrypt or retrieve the volume encryption key for all volumes within a disk group.
- Enable access control and auditing for KMS operations to enhance security and compliance.

For details, refer to the following documents:

- *InfoScale 9.1 Installation and Upgrade Guide*
- *InfoScale 9.1 Storage Foundation Administrator's Guide*

Tunable for automatically deleting stale KMS keys

InfoScale now provides a new registry entry - `DEL_KMS_KEYS` \- to automatically delete stale keys after a re-key operation or after disk group deletion.

By default, the value of this registry entry is set to 0 (zero). When the value of this registry entry is set to 1, any stale keys are automatically deleted from the KMS server, regardless of whether it is a KMIP-based or a cloud-based KMS configuration.

For details, refer to the *InfoScale 9.1 Storage Foundation Administrator's Guide*.

Tunable for scheduling SecureFS snapshots

InfoScale now provides a new registry entry - `SECUREFS_SNAP_INTERVAL` \- to set a minimum time interval limit in hours for scheduling a SecureFS snapshot. This registry-based tunable helps storage administrators configure the minimum time interval below which snapshots cannot be configured. It ensures that a central time policy is present for all users.

By default, the value of this registry entry is set to 12 (hours), which is the minimum permitted value. Snapshots cannot be scheduled if this value is less than 12.

Additionally, if this registry entry is set to a value higher than 12, then the same value or higher must be used to schedule a snapshot. For example, if `SECUREFS_SNAP_INTERVAL` is set to 15, the time interval must be set to 15 (hours) or higher when scheduling a snapshot using the VEA GUI or the `vxschadm` command.

For details, refer to the *InfoScale 9.1 Storage Foundation Administrator's Guide*.

Enhanced support for replication using virtual host names

The Volume Replicator (VVR) component has been enhanced to support replication across cloud availability zones and multi-subnet environments using virtual host names.

VVR now validates the source and destination IPs before use as follows:

- The use of the same IPs is enforced for both `vradmin` and replication.
- Hostname resolution includes iterating through local IPs to find a valid path.
- Disconnects trigger reevaluation of IPs to restore the functioning of `vradmin` and replication.
- In failover scenarios, the IPs are updated dynamically to ensure that the correct replication paths are followed.

Before you can configure replication using virtual hostnames, the following prerequisites must be met:

- Applications must be configured according to the VVR replication requirements.
- Dependencies between the application service groups must be defined appropriately.

The additional step for specifying the virtual hostnames needs to be performed after configuring replication using the Disaster Recovery (DR) wizard.

For details on configuring disaster recovery (DR) using the DR wizard, refer to the *InfoScale 9.1 Storage Foundation and High Availability Solutions - Solutions Guide*.

For details on using the VVR enhancement, refer to the *InfoScale 9.1 Volume Replicator Administrator's Guide*.

Changes in InfoScale Volume Replicator Advisor (VRAdvisor)

The InfoScale VRAdvisor is a planning tool that helps you collect and analyze data based on your application's I/O patterns and network parameters. VRAdvisor uses that data to recommend optimal replication configuration settings such as required network bandwidth, size of the Storage Replicator Log (SRL), and the replication mode (synchronous or asynchronous).

The following changes are introduced in this release:

- This release reintroduces the VRAdvisor tool with all the built-in data collection and data analysis features fully restored and functional as in previous InfoScale versions.
- The VRAdvisor is supported on the operating systems that are supported by InfoScale 9.1.
- VRAdvisor installer for UNIX platforms is included in the InfoScale 9.1 for Windows distribution package. You can find the tool in the `Tools > storage_foundation` directory.
- You can install the VRAdvisor on any Windows system.

For detailed instructions on VRAdvisor installation, data collection, and data analysis workflows, refer to the *InfoScale Volume Replicator Administrator's Guide*.

Limitations

This section includes the following topics:

- [Deployment limitations](#)
- [Cluster management limitations](#)
- [Storage management limitations](#)
- [Multi-pathing limitations](#)
- [Replication limitations](#)
- [Solution configuration limitations](#)
- [Internationalization and localization limitations](#)
- [Interoperability limitations](#)

Deployment limitations

This section lists the software limitations related to installing, upgrading, repairing, and uninstalling the InfoScale products.

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)

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
```

SnapDrive service fails to start after InfoScale Availability or InfoScale Enterprise is uninstalled

When InfoScale Availability or InfoScale Enterprise is uninstalled and the system is restarted, the SnapDrive service fails to start with a logon failure.

Workaround: Reset the password for the SnapDrive service account and then start the SnapDrive service.

Cluster management limitations

This section lists the software limitations related to the cluster management components; these components are available as part of the InfoScale Availability and InfoScale Foundation products.

EBS Multi-Attach support in AWS cloud and InfoScale service group configuration using wizards

The following limitations are applicable when you use these wizards on Windows:

- If the cluster nodes span across different subnets, the wizards configure the AWSIP agent resources using the private IP address (PrivateIP attribute) instead of using the overlay IP address (OverlayIP attribute).

In such a case, you must run the wizard until the end and complete the service group configuration. After the service group is configured, you must manually modify the service group and replace the PrivateIP entry with the OverlayIP attribute and the overlap IP address.

Shared disk support in Azure cloud and InfoScale service group configuration using wizards

InfoScale supports shared disks in an Azure cloud environment. You can use the InfoScale application service group configuration wizards on Windows platform to configure high availability for various applications such as SQL Server and file shares. The service groups automatically configure the required InfoScale Azure agents for managing the application and the network resources in the Azure cloud.

The following limitations are applicable when you use these wizards on Windows:

- If the cluster nodes span across different subnets, the wizards configure the AzureIP agent resources using the private IP address (PrivateIP attribute) instead of using the overlay IP address (OverlayIP attribute).

In such a case, you must run the wizard until the end and complete the service group configuration. After the service group is configured, you must manually modify the service group and replace the PrivateIP entry with the OverlayIP attribute and the overlap IP address.

- When you create new service groups in a shared storage environment, the wizards configure the Azure agents with managed identities for authenticating the Azure subscription. You cannot configure the AzureAuth agent to use service principal credentials for managing Azure authentication.

If required, you can manually modify the service group later and add the AzureAuth agent resources and dependencies. However, Arctera recommends that you use managed identities for Azure authentication.

Refer to the *Cluster Server Bundled Agents Reference Guide - Windows* for more information about Azure agents and user-assigned managed identities.

Undocumented commands and command options

VCS contains undocumented commands and command options intended for development use only. Undocumented commands are not supported for external use.

Unable to restore user database using SnapManager for SQL

While restoring the user database using SnapManager, user needs to select the checkbox which specifies that the snapshot was taken on other machine. The user is also required to select the system name. When this option is selected, it is possible to restore the database.

MountV agent does not work when Volume Shadow Copy service is enabled

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

WAN cards are not supported

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

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.

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 *Cluster Server Administrator's Guide* for more information.

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)

Service group dependency limitation - 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.

Unable to monitor resources if Switch Independent NIC teaming mode is used

VCS requires the MAC address of the team NIC to be static. In the default NIC teaming mode (Switch Independent), a static MAC address is not assigned to the team NIC. Therefore, if you plan to use NIC teaming, make sure that the Static Teaming mode is enabled. Otherwise, VCS will not be able to successfully monitor resources. The VCS agent will fail to identify the resources for which they are configured, and report the UNKNOWN state.(3011749)

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)

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.

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)

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.

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

VCS Authentication Service does not support node renaming

The VCS Authentication Service does not support renaming the cluster nodes.

An MSMQ resource fails to come online after the virtual server name is changed

After you configure an Microsoft Message Queuing (MSMQ) service group, you can change the virtual server name by running the MSMQ Configuration Wizard in the Modify mode. However, after you change the virtual server name, the MSMQ resource fails to come online. This issue occurs because some of the GUIDs related to MSMQ are tightly coupled with the virtual server name. The resource always looks for the coupling between such GUIDs and the original virtual server name, even though the name may have changed.

Recommendation: Do not change the virtual server name after an MSMQ service group is configured. If you inadvertently change the virtual server name, use the MSMQ Configuration Wizard to delete the service group and create it again.

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 limitation in one of the following ways.

- Use the appropriate wizard to delete the service group and create it again using resources with IPv6 addresses.
- Manually replace the IPv4 resources in the service group with corresponding IPv6 resources.

All servers in a cluster must run the same operating system

All servers in a cluster must run the same operating system. Systems that run different versions of Windows Server cannot form the nodes of a cluster.

Running Java Console on a non-cluster system is recommended

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

Cluster Manager console does not update GlobalCounter

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

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.

Storage management limitations

This section lists the software limitations related to the storage management components; these components are available as part of the InfoScale Foundation, InfoScale Storage, and InfoScale Enterprise products.

SFW does not support disks with unrecognized OEM partitions (3146848)

SFW does not support disks with an unrecognized original equipment manufacturer (OEM) partition or a partition that is not recognized by Microsoft Windows operating systems. Therefore, any operations on such a disk, such as creating a cluster disk group, are not supported. Refer to the article at:

<https://helpcenter.cloud.com/wolken-support/kbsearch-article-100029448>

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.

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 stop 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.

Database or log files must not be on same volume as SQL Server

When using the `vxsnapsql` utility, user-defined databases and logs must not be stored on the same volume as the SQL Server program files or system data files. (266133)

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)

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)

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 Arctera 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 Arctera VSS provider. If the Arctera 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

vxsnapsql restores all SQL Server databases mounted on the same volume

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

Workaround: To avoid this situation, configure each SQL Server database on its own separate dynamic volume.

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.

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.

Arctera 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.

Arctera 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.

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)

Storage migration of Hyper-V VM on cluster-shared volume resource is not supported from a Slave node

This issue occurs if you want to perform a storage migration from a Slave node of a Hyper-V virtual machine created on a cluster-shared volume resource. Currently, this is not supported and you cannot perform the storage migration. (3385754)

Workaround: As a workaround, you need to either perform a cluster-shared disk group (CSDG) level storage migration from the Master node or fail over the Hyper-V virtual machine to the Master and perform a VM-level storage migration.

In a CVM environment, disconnecting and reconnecting hard disks may display an error

This issue may occur if you have configured volume resources in a cluster-shared disk group, and brought them online on the Slave node. Now if you disconnect the disk, the volume resources fail over to the Master node. If you reconnect the disk to the Slave node, and fail back the resources to the Slave node, a dialog box may appear asking you to format the volume, even when the volume is accessible. This dialog box can be safely ignored. (3274685)

Limitations of SFW support for DMP

The limitations of SFW support for DMP are as follows.

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.

Multi-pathing limitations

This section lists the software limitations related to the multi-pathing components; these components are available as part of the InfoScale Foundation, InfoScale Storage, and InfoScale Enterprise products.

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.

Replication limitations

This section lists the software limitations related to the replication components; these components are available as part of the InfoScale Storage and InfoScale Enterprise products.

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)

Solution configuration limitations

This section lists the software limitations related to configuring, administering, and unconfiguring the various InfoScale solutions for the supported applications and hardware replication environments.

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.

However, the Fire Drill Wizard available from the Solutions Configuration Center enables you to set up and run a fire drill on a disaster recovery environment that uses Volume Replicator replication.

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.

Solutions 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 SFW 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.

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.

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: https://www.veritas.com/content/support/en_US/article.100023202

Wizard fails to discover SQL Server databases that contain certain characters

The SQL Server Agent Configuration Wizard fails to discover a SQL Server database if the database name contains any of the following characters:

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

Internationalization and localization limitations

This section lists the software limitations related to using the InfoScale products in locales other than U.S. English.

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.

Interoperability limitations

This section lists the software limitations related to the coexistence and usage of the InfoScale products with other software.

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)

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: Arctera 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.

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.

Known issues

This section includes the following topics:

- [Deployment issues](#)
- [Cluster management issues](#)
- [Storage management issues](#)
- [Multi-pathing issues](#)
- [Replication issues](#)
- [Solution configuration issues](#)
- [Internationalization and localization issues](#)
- [Interoperability issues](#)
- [Miscellaneous issues](#)

Deployment issues

This section lists the known issues related to installing, upgrading, repairing, or uninstalling the InfoScale products.

Reinstallation of an InfoScale product may fail due to pending cleanup tasks

When you attempt to reinstall an InfoScale product, the installation may fail due to pending clean up tasks. This issue occurs even though the clean up tasks for the previous installation are complete. (3806870)

During the installation, the validation tasks on the System Selection page fail and the wizard is unable to proceed with the installation.

Workaround:

To resolve the issue, perform the following steps:

1. Run the command prompt in run as administrator mode.
2. Enter the command `sc delete cleanupservice`

3. Reboot the system and then run the installation wizard again.

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)

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

During an InfoScale installation, if the Windows Installer Service is not accessible, the installer fails to proceed with the installation. (2497344)

This issue typically occurs if the Windows Installer Service, which is a native component of an operating system, is damaged or unregistered.

Workaround: Repair the operating system installation and launch the InfoScale installer again.

For details on the "Windows Installer Service could not be accessed" error and its workaround, refer to the Microsoft knowledge base.

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

The product installation wizard may fail to install an InfoScale product 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. (2365128)

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

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.

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 SFW 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.

In SFW with Microsoft failover cluster, a parl.exe error message appears when system is restarted after SFW installation if Telemetry was selected during installation

This issue occurs if you had installed SFW with Microsoft failover cluster option and chose to participate in the Arctera Product Improvement Program (Telemetry) during installation. In this case, a parl.exe application error message appears when you restart the system after installing SFW. The error message does not affect the product functionality and should be ignored. (3045916)

Workaround: Click OK in the error message dialog box and ignore the message. The message will not appear again.

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

While installing an InfoScale product, 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 .NET installation is corrupted. (2406978)

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

FlashSnap License error message appears in the Application Event log after installing license key

When an evaluation copy of SFW with the FlashSnap feature is installed on a system, and then later a permanent license key for SFW that does not have the FlashSnap feature is installed on it, an error message is logged periodically in the Application Event log: (1862627)

Storage Foundation FlashSnap License is not installed

Workaround: To avoid having the error message logged in the Application Event log, use the Add or Remove Programs of the Windows Control Panel to remove the demo license and add the permanent license key.

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 the product if licensing files are missing

This issue occurs if the licensing files are missing from the system, and you try to uninstall the product. (3330686)

The product installation wizard gives the following error on the system selection page:

License details for the following systems are not specified: <system_name>

Workaround: Use the Add or Remove Programs of the Windows Control Panel to add the product license and retry uninstalling the product.

The vxlicrep.exe may crash when the machine reboots after InfoScale Enterprise is installed

This issue may occur due to some internal error. (3800413)

Even if the vxlicrep.exe crash error appears, you can successfully run the vxlicrep command to generate the licensing report.

Workaround:

Ignore the error displayed.

Cluster management issues

This section lists the known issues related to the cluster management components; these components are available as part of the InfoScale Availability and InfoScale Foundation products.

Cluster Server (VCS) issues

This section lists the known issues related to the Cluster Server (VCS) component.

Cluster reconfiguration may fail post InfoScale installation (4117528)

This issue is observed after you install InfoScale Enterprise and configure a cluster. After configuring a disk group if you delete the cluster and then try to recreate the cluster again, the cluster configuration fails.

Workaround: If you want to delete the cluster, first restart the Arctera Enterprise Administrator (`vxsvc`) service and then proceed with any cluster reconfiguration operations.

Refer to the following for instructions:

<https://helpcenter.cloud.com/wolken-support/kbsearch-article-100047179>

The VCS Cluster Configuration Wizard may not be able to delete a cluster if it fails to stop HAD

When deleting a cluster, the Cluster Configuration Wizard needs to stop the HAD. If it fails to stop this service, the wizard is unable to complete the delete operation. (3298516)

Workaround: Forcefully stop the HAD using the following command:

```
hastop -all -force
```

Relaunch the Cluster Configuration Wizard to delete the cluster. The wizard prompts you to confirm whether to delete the cluster while the HAD is stopped. Click Yes to confirm the delete operation.

Deleting a node from a cluster using the VCS Cluster Configuration Wizard may not remove it from main.cf

When you attempt to delete a node from a cluster, the wizard unconfigures the cluster and stops the HAD on the specified system. The wizard also updates the cluster configuration to register the removal of this node. If the HAD stops before the cluster is unconfigured on the system, the wizard

fails to update the cluster configuration accordingly. Thus, the stale entry for that node remains in `main.cf` . (3249537)

Workaround: On the system from which you want to remove the cluster node, stop the HAD manually using the command:

```
hastop -local
```

Delete the node from the cluster using the command:

```
hasys - delete systemName
```

Update the cluster configuration using the command:

```
haconf -dump
```

NetAppSnapDrive resource may fail with access denied error

You may intermittently observe an issue where the NetAppSnapDrive resource faults with the access denied error. This issue typically occurs in a DR environment, due to the network connectivity issues.

Due to the connectivity issue, the NetAppSnapDrive resource fails to connect to the LUN and retrieve the Data ONTAP version running on the storage system.

You may observe that the issue is resolved during the subsequent probe. (2422479)

Mount resource fails to bring file share service group online

If you are using Windows Logical Disk Manager (LDM) for shared storage, you may encounter problems while bringing service groups online or switching service groups. Create file share service group and try to bring it online. Mount resource fails to bring the file share service group online or switch the service group. (1266158)

Workaround:

Set the attribute AutoFSClean of Mount resource agent to 1. The default value is 0. The agent cleans the file system by running Chkdsk/X on the volume being brought online.

Caution: Cleaning a file system may result in loss of data.

Mount agent may go in unknown state on virtual machines

On virtual machines having VCS for Windows installed while configuring the Mount agent, node may go in an unknown state. This happens because the attribute PartitionNo is Global and is seen differently on the shared disks. (1262346)

Workaround:

Perform the following steps

1. Take the service group offline.
2. Rescan the disks from all the nodes.
3. To retrieve information about the PartitionNo, type the following at the command prompt:

```
C:\>havo1 -getdrive -details
```

The information about the disk is retrieved and stored in a text file in the same path from where you executed the command.

1. Edit the attribute PartitionNo from Global to Per System. Assign the value for PartitionNo from the disk details retrieved in the above step.

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.

Enterprise Vault Configuration Wizard 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 `virtualServerName\instanceName` format. Sometimes, even though you specify the SQL Server instance name in the required format, you may encounter an error after you click Next. The error message states that Enterprise Vault failed to connect to the SQL Server instance. (2432332)

Workaround:

In the appropriate field, delete the virtual server name and add the physical server name in the `physicalServerName\instanceName` format.

Click Next, and after the next panel appears, click Back.

Re-enter the virtual server name in `virtualServerName\instanceName` format, and click Next.

Step through the wizard to complete the task.

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.

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 creating VMNSDg resources, the VMGetDrive command not supported to retrieve a list of dynamic disk groups

This issue occurs while creating VMNSDg (Volume Manager Non-Shared Diskgroup) resources in a VCS configuration. For creating the VMNSDg resources, the `vmgetdrive` command does not work in retrieving the list of dynamic disk groups. (2937411)

Workaround: To resolve this issue, use either the `vmgetdrive dynamicdg` command or the `vxdg list` command to retrieve the list of dynamic disk groups.

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.

Resource fails to come online after failover on secondary

During a migrate or failover on the secondary, the NetAppSnapDrive resources fail to come online.

Workaround: Mount the LUNs manually using the SnapDrive GUI or CLI and then probe the resources.

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 High Availability Engine (HAD).

Type the following at the command prompt:

```
net stop had
```

```
net start had
```

1. 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`.

HTC resource probe operation fails and reports an UNKNOWN state

This issue occurs after successfully installing the HTC agent and creating a HTC resource. If you try to probe the HTC resource, the probe operation fails and the state is reported as UNKNOWN. This is because the existing HTC.dll in the VCS bin needs to be replaced with the updated Default50Agent.dll

Workaround: To resolve this issue, you need to replace the HTC.dll with Default50Agent.dll

Perform the following steps using the CLI:

1. Stop the HTC agent using the following CLI command:

```
haagent -stop HTC -sys <systemname>
```

1. Delete the existing HTC.dll from %vcs_home%\bin\HTC
2. Copy Default50Agent.dll from %vcs_home%\bin to %vcs_home%\bin\HTC
3. Rename Default50Agent.dll to HTC.dll in the %vcs_home%\bin\HTC folder.
4. Start the HTC agent using the following CLI command:

```
haagent -start HTC -sys <systemname>
```

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.

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 and then bring it online.

Use the following commands:

```
hares -probe resource -sys system
```

```
hares -online resource -sys system
```

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.

The Cluster Server High Availability Engine (HAD) service fails to stop

This issue may occur if you try to stop the Cluster Server High Availability Engine (HAD) service, from the `services.msc` prompt or using the following command: (3273290)

```
net stop had
```

The task fails with the following error:

```
Failed to stop the HAD service. Node='systemname',  
  
Error=00000425."
```

This issue typically occurs in case of a secure VCS cluster containing a single system.

Workaround: Perform the following steps to resolve this error:

1. Forcefully stop the High Availability Engine (HAD) process on the system using the following command:

```
taskkill /f /im had.exe
```

- a. If HAD starts again, stop HAD using the following command:

```
hastop -local
```

2. Ensure that the HAD process is in the stopped state.

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.

The VCS Cluster Configuration Wizard (VCW) supports NIC teaming but the InfoScale High Availability Configuration Wizard does not

The VCS Cluster Configuration Wizard (VCW) supports Windows NIC teaming. However, the InfoScale High Availability Configuration Wizard does not support NIC teaming. (3048358)

If you wish to use NIC teaming you must use VCW to configure the VCS cluster and then configure the application using the application configuration wizards or the InfoScale High Availability Configuration Wizard.

Note: While using Windows NIC teaming you must select the mode as Static Teaming. Only the Static Teaming mode is currently supported.

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 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 High Availability Engine (HAD) on all the cluster nodes.

On each cluster node, type the following on the command prompt:

```
net stop had
```

1. 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`.

1. From `%VCS_home%\conf\config` directory, locate the configuration file **main.cf** and open it in a text editor.
2. In main.cf, search for the text "**SecureClus=1** " and delete that line altogether.
3. Save the file and close the text editor.

4. Start the VCS engine (HAD) locally on the node where you performed the earlier steps.

Type the following on the command prompt:

```
hastart
```

1. Set the cluster configuration to read/write mode.

Type the following on the command prompt:

```
haconf -makerw
```

1. Add a user to the cluster and assign it with cluster administrator privileges.

Type the following command:

```
hauser -add \<username\> -priv Administrator
```

1. Enter the password for the user, when prompted.
2. Save and make the cluster configuration read-only.

Type the following on the command prompt:

```
haconf -dump -makero
```

1. Start the 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. If required, run the VCS Cluster Configuration Wizard (VCW) and reconfigure the cluster to use single sign-on authentication.

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.

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.

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.

SQL Server service resource does not fault even if detail monitoring fails

This issue may occur when detail monitoring is configured and IMF is enabled for SQL Server agents. (2536297)

If detail monitoring fails (either the database becomes unavailable or there is a failure in the detail monitoring script), the SQL service resources fault 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 Server agents, the SQL service resources does not fault. Instead, the SQL Agent service resource (SQLServerAgent) faults.

This occurs because when detail monitoring fails, the SQL Server 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 Server service resource receives this VCS initiated offline and therefore does not fault in response to the original detail monitoring failure event.

Workaround: There is no known workaround for this issue at this time.

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)

NetBackup may fail to back up SQL Server database in VCS cluster environment

In a VCS cluster environment, backup of the SQL database with 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 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.

Cluster Manager (Java Console) issues

This section lists the known issues related to the Cluster Manager (Java Console).

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: bmc and bmcmap file location

COPY FROM THIS DIRECTORY	COPY TO THIS DIRECTORY
(For English)	%VCS_HOME%\messages\en
D:\Program Files\Veritas\messages\en	Where, %VCS_HOME% is the default
Where, D: is the drive on which VCS is installed.	installation directory for VCS,
	typically C:\Program
	Files\Veritas\Cluster
	Server.

Login attempts to the Cluster Manager may fail after a product upgrade

After upgrading InfoScale, you need to reconfigure any existing secure clusters so that they can use the 2048-bit key and SHA-256 signature certificates. However, even after you successfully reconfigure an existing cluster, login attempts to the Cluster Manager may fail. (3857832)

This issue occurs due to the cached credentials whenever the Cluster Configuration Wizard (VCW) is used to reconfigure a cluster. It does not occur if you configure a new cluster using the VCW after installing the patch.

Workaround: Log off from the system and log in again. You should then be able to log in to the Cluster Manager to perform further operations.

Global service group issues

This section lists the known issues related to global service groups created using the ProdFamilyNameS_InfoScale; products.

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: Arctera 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 over service groups to the local cluster. However, the local cluster may not be the cluster that is assigned the highest priority in the cluster list.

Workaround: Run the following command to bring a global service group online on a remote cluster:

```
hagrp -online service_group -sys system [-clus cluster]
```

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.

1. 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.

1. Bring the service group online on the remote cluster.

Type the following on the command prompt:

```
hagrp -online global_group -any -clus remote_cluster
```

VMware virtual environment-related issues

This section lists the known issues related to working with cluster management components of the InfoScale products in a VMware virtual environment.

Partial log entries for service group offline operations in case of application monitoring on single-node clusters (4188474)

If VMware happens to restart the node while a service group offline is in progress, you may find partial log entries for that offline operation. Such instances occur when a lot of service groups or resources are involved in the hierarchy, due to which, the offline operation may take too long to complete.

Workaround: None. Subsequently, if the node has restarted and the application service groups have been brought online successfully, you can ignore these old, partial log entries.

No corrective action is taken for a faulted service group even when its name is provided in the ServiceGroupName attribute of the AppMonHB resource (4187908)

This issue occurs when a mix of application monitoring configurations done manually and using the `ConfigAppMonHB` utility are present in a single-node cluster.

Workaround: None for the current mixed configuration.

You can proceed by updating the configuration in one of the following ways:

- Option 1: Delete the manually configured service groups and add them using `ConfigAppMonHB`.
- Option 2: Delete the service groups configured using `ConfigAppMonHB` and add them manually.

If the service group faults again, the corrective action should now be triggered successfully.

Recommendation: Configure the service groups on the node either manually or using `ConfigAppMonHB`, but not a combination of both. Corrective actions can be taken appropriately when configurations are done entirely manually or entirely using the utility. To avoid running into any such issues, Arctera recommends that you use `ConfigAppMonHB` for all application monitoring configurations on a single-node cluster.

Irrelevant INFO message about corrective action for a faulted service group (4178969)

When a service group faults and the `WaitBeforeCorrectiveAction` period elapses, an INFO message indicating that a corrective action has been initiated appears in the VCS logs. However, if the fault is resolved and the service group is brought online shortly after this period, the corrective action is not triggered, rendering the INFO message irrelevant.

Workaround: None. You can disregard this INFO message as no corrective action is necessary if the service group is already online.

VMwareDisks resource cannot go offline if VMware snapshot is taken when VMware disk is configured for monitoring

This issue occurs when the `VMwareDisks` agent is configured to monitor VMware disks of any type, and a VMware snapshot is taken. This issue occurs for both thin-provisioned disks and thick-provisioned disks. (2801599)

When the service group is taken offline or failed over, the `VMwareDisks` resource cannot go offline. The `VMwareDisks` agent cannot detach the disk, because the size of the snapshot disk is different from the size of the base disk. The following type of message appears 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'.'
```

Workaround: Delete all the VMware snapshots available for the VM.

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: Configure the "PingHostList" attribute for the VCS NIC agent. For more details, refer to the *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
```

- a. Run the following commands sequentially on each cluster system:

```
net stop VCScomm
```

```
net stop gab
```

```
net stop llc
```

1. Restart the VCS cluster using the following command:

```
hastart -all
```

VMWareDisks resource cannot probe (Unknown status) when virtual machine moves to a different ESX host

When Changed Block Tracking (CBT) is enabled in the VMWare virtual machine settings, the `vmkfstools -J getuuid` command returns an error, which causes the VMWareDisks resource to go into the Unknown state. (3915390)

Workaround: This issue is traced to VMware, and can be addressed by applying a patch. Perform the tasks described in the TechNote at:

<https://helpcenter.cloud.com/wolken-support/kbsearch-article-100032416>

Storage management issues

This section lists the known issues related to the storage management components; these components are available as part of the InfoScale Foundation, InfoScale Storage, and InfoScale Enterprise products.

Data corruption observed on snap plex when user reverts cow snapshot and creates the mirror snap. (4114185)

Data corruption observed on snap plex when user reverts cow snapshot and creates the mirror snap while performing below operations in sequence:

Revert COW snapshot > Shrink Volume > Add mirror > Break mirror > Evacuate disk.

Workaround: There is no workaround.

The GCO network selection page does not show network adapter cards. (4100888)

When cluster is configured using LLT over UDP configuration, the GCO network selection page does not show network adapters.

Workaround: Re-run the Cluster Server Configuration Wizard.

Storage Foundation issues

This section lists the known issues related to the Storage Foundation (SFW) component.

In Microsoft Azure environment, InfoScale Storage cannot be used in cluster configurations that require shared storage

To use InfoScale Storage in a cluster configuration, you must configure shared disks between the cluster nodes. Microsoft Azure does not support provisioning of shared disks. Due to this Microsoft limitation, InfoScale Storage cannot be used in cluster configurations that require shared storage. (IIP-24673)

In a Microsoft Azure environment SFW fails to auto discover SSD media type for a VHD

In a Microsoft Azure environment, the provisioned storage disks can be of SSD or HDD type. A hypervisor assigns these disks to a virtual machine as SCSI registered VHD disks. For the VHD disks, whether SSD or HDD, the media type information is identical. As a result, SFW fails to auto

discover whether or not is the underlying disk an SSD. Failing this, you cannot leverage SFW capabilities for an SSD disk. (3908299)

Workaround:

To leverage SFW capabilities for an SSD disk, using VEA, manually classify an SSD disk.

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

Incorrect message appears in the Event Viewer for a hot relocation operation

A hot relocation operation is performed successfully for a volume that has multiple subdisks for a single plex. However, the Event Viewer incorrectly displays a message stating that the relocation of some of the subdisks failed and only the last subdisk was relocated successfully.

You can manually verify whether all the subdisks are successfully relocated, and then ignore this message.

A volume state continues to appear as "Healthy, Resynching"

This issue may occur after initiating an Add Mirror operation on a volume. (3854819)

If a volume resource fails over when an Add Mirror operation is in progress, the status of that volume continues to appear as "Healthy, Resynching". This occurs because the Add Mirror operation is not reinitiated on the failover node.

Workaround:

Manually delete the Add Mirror task and initiate the operation again.

To manually delete the Add Mirror task, perform the following using VEA:

1. From the local host tree, select the Volume that shows "Healthy, Resynching" state.
2. In the right-side pane, select the Mirrors tab.
3. From the list of mirrors displayed, select the mirror that shows "Attaching" status.
4. In the context menu, click Remove Mirror.

After the selected mirror is removed, initiate the Add Mirror operation again.

After a mirror break-off operation, the volume label does not show on the Slave node

This issue occurs when you perform a mirror break-off operation. (3573381)

If you apply a volume label during the break-off operation, this new volume label does not show on the Slave node after the break-off operation.

Workaround:

To update the volume label, perform the following steps on the Slave node:

1. Bring the volume online.
2. Perform a Refresh operation.

CVM cluster does not support node names with more than 15 characters

This issue occurs while configuring a CVM cluster or adding a new node to it. If the node/host name exceeds the maximum limit of 15 characters, then you will not be able to configure it for the CVM cluster because it is not supported. (3351326)

Workaround: Ensure that the node name does not exceed the maximum limit of 15 characters.

For CSDG with mirrored volumes, sometimes disks incorrectly shows the yellow warning icon

This issue occurs if you have a cluster-shared disk group (CSDG) with mirrored volumes. After performing any SFW operation, one or more disks may incorrectly show the yellow warning icon. For example, this occurs when you deport and import a CSDG with mirrored volume. This does not suggest any problem or affect product functionality. (3385783)

Workaround: As a workaround to remove the yellow icon, right-click the disk and then select Reactive Disk and click Yes.

SSD is not removed successfully from the cache pool

While creating a cache area, the SSD is placed in the cache pool. If the operation fails, then SmartIO should remove the SSD from the cache pool to allow the user to retry creating cache area with the SSD. But if this issue occurs, then the SSD is not removed properly from the cache pool, and hence it cannot be added to the cache area.

Similarly, if a cache area is deleted, SmartIO should remove the SSD in the cache area from the cache pool. Now if the operation fails, the SSD is not removed successfully from the cache pool.

Workaround: Run the following command to remove the SSD from the cache pool:

```
vxdbg -gcachepool -f rmdisk <diskname>
```

On fast failover enabled configurations, VMDg resource takes a longer time to come online

This issue occurs on the fast failover enabled configurations, where the VMDg resource takes a longer time (around two to three minutes) to come online. (3274957)

Workaround: There is no workaround for this issue.

On some configurations, the VSS snapshot operation may fail to add volumes to the snapshot set

On certain configurations, this issue may occur when you perform a VSS snapshot operation. In this case, the snapshot operation fails to obtain the CLSID (class ID) of the VSS software plugin and displays the following error:

Volume Shadow Copy Service error: Error creating the Shadow Copy Provider COM class with CLSID \<Class ID\> Class not registered

Because of this, the snapshot operation fails to add volumes to the snapshot set. (3272033)

Workaround: As a workaround, re-register the VSS software plugin by running the following commands using the CLI, and then retry the operation:

1. `net stop "veritas vss provider"`
2. `vxvssprovider.exe /unregserver`
3. `vxvssprovider.exe /service`
4. `net start "veritas vss provider"`

Using Failover Cluster Manager you cannot migrate volumes belonging to Hyper-V virtual machines to a location with VMDg or Volume Manager Shared Volume resource type

While migrating volumes belonging to Hyper-V virtual machines using Windows Server's Failover Cluster Manager snap-in, if the destination storage location has Volume Manager Disk Group (VMDg) or Volume Manager Shared Volume resource type, then they are not displayed. Therefore, you cannot migrate Hyper-V virtual machines using Failover Cluster Manager. This is a known Microsoft issue. (3276940)

Workaround: As a workaround, you can migrate Hyper-V virtual machines using either Hyper-V Manager or System Center Virtual Machine Manager (SCVMM).

Performance counters for Dynamic Disk and Dynamic Volume may fail to appear in the list of available counters

This issue is observed in the Windows Performance Monitor tool (`perfmon`), typically after any of the following operations are performed:

- Rebooting the system on which the InfoScale product is installed
- Installing or upgrading the InfoScale product
- Restarting the Veritas Enterprise Administrator service (`VxSVC`).

After these operations are complete, the Windows performance counter library (`vmperf.dll`) may load before the `vmstat.dll` module of `VxSVC` is loaded. As a result, the `Disable Performance Counters` registry with the value 1 gets added at the following location:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\vxio\Performance
```

Since the registry value is set to 1, the performance counters for Dynamic Disk and Dynamic Volume fail to appear in the list of available counters. Even though these counters fail to appear in the list of available counters, the performance trend for SFW is not affected. (3791725)

Workaround: Reset the `Disable Performance Counters` registry key value to 0 or delete the registry key.

VSS snapshot of a Hyper-V virtual machine on SFW storage does not work from NetBackup

This issue occurs while taking VSS-aware snapshot or backup of a Hyper-V virtual machine using NetBackup and the virtual machine's VHD is on a dynamic volume belonging to an SFW cluster disk group.

During the VSS snapshot process, VSS service on the host machine (hypervisor) internally takes a snapshot of the VHDs belonging to the guest virtual machine (OS VHD as well as data VHDs) and exposes these snapshots to the host machine. As a result, we can see disk arrivals (snapshots of the guest VHDs) and the corresponding volume arrivals (partitions on the guest VHDs) on the host machine.

When SFW is installed, auto-mounting is disabled by default. But, for a successful backup of a Hyper-V virtual machine (guest), VSS on the host machine expects the arrived partitions to be automatically mounted. When VSS doesn't find these partitions to be mounted, it fails the backup operation of the guest virtual machine and NetBackup displays error with the code 156. (3342470)

Workaround: There is no workaround for this issue.

In some cases, when 80 or more mirrored volume resources fail over to another node in CVM, some volume resources fault causing all to fault

This issue may occur in a CVM cluster if you are failing over 80 or more mirrored volume resources in a cluster-shared disk group (CSDG) from one node to another. Sometimes, some resources for the mirrored volumes may fault during the failover, which causes even the successfully failed over other volume resources to fault. This happens because of an issue in the CVM message handling behavior. Microsoft failover cluster shows all the resources as faulted. Because of this issue, CVM hangs. (3429391)

Note: Chances of encountering this issue are less if you are performing failover to more than two nodes in the cluster.

Workaround: As a workaround, change the GAB message handling settings as follows:

1. On all the nodes in the cluster, run the following command from the command-line interface:

```
gabconfig -Q recvq:15
```

This command changes the value of the GAB receive queue length from 5 (default) to 15.

You can view the current value of the `recvq` parameter by running the `gabconfig -l` command.

1. Because running the command in Step 1 only makes changes for the running instance, you must also update the value of the command to 15 in the `gabtab.txt` file located under `C:\Program Files\Veritas\comms\gab`. This ensures that the changes are not lost if CVM or the machine restarts.

In CVM, if subdisk move operation for CSDG fails because of a cluster reconfiguration, it does not start again automatically

In a CVM cluster, this issue occurs while performing the subdisk move operation for a cluster-shared disk group (CSDG). If the subdisk move operation fails because of a cluster reconfiguration (such as new node joins or Master switch operation), then it does not start again automatically. (3429651)

Workaround: As a workaround, manually delete the plex that is in the "attaching" state, and restart the subdisk move operation.

Stale volume entries under MountedDevices and CurrentControlSet registries not removed after the volume is deleted

When you create a volume, its volume entry gets created under the following registry keys:

```
HKEY_LOCAL_MACHINE\SYSTEM\MountedDevices
```

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Enum\STORAGE\Volume
```

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\DeviceClasses\{53f5630d-b6bf-11d0-94f2-00a0c91efb8b}
```

This issue occurs when you delete a volume, during which its corresponding volume entries under the MountedDevices and CurrentControlSet registries are not deleted automatically. If these stale entries are not deleted, then the registry becomes full with them. Because of this, the volume arrival operation may take a long time or fail. (3329663)

Workaround: As a workaround, use the `vxscrub` utility to delete stale volume entries that are not needed.

Some arrays do not show correct Enclosure ID in VDID

Some of the arrays do not show correct Enclosure ID in Veritas Disk ID (VDID), which contains information that can identify the array that a disk is contained in.

These arrays are Fujitsu, Compellent, EMC VPLEX, HP MSA P2000, HP MSA 2040, IBM DS5020, IBM DS3524, Sun 6540, and Sun 7310. (3351984)

Workaround: There is no workaround for this issue.

In some cases, EV snapshot schedules fail

This issue may occur when you create snapshot set schedules for an Enterprise Vault (EV) component. In some cases, the snapshot schedules may fail. (3426945)

The following error message is displayed:

V-76-58657-13: The specified volume(s) are not supported for the requested operation by VSS provider

Workaround:

1. Run the `vxassist refresh` command on the SQL Server node (for the "hassnap" property to get updated).

2. Snapback (reattach) all the volumes on the EV server side that got broken in the failed snapshot attempt.

In CVM, Master node incorrectly displays two missing disk objects for a single disk disconnect

This issue occurs in a CVM cluster after one of the disks of a cluster-shared disk group (CSDG) gets disconnected. The Master node displays two missing disk objects even though only one disk is disconnected. (3284450)

Workaround: To resolve this issue, refresh all nodes of the CVM cluster.

DRL plex is detached across CVM nodes if a disk with DRL plex is disconnected locally from the node where volume is online

In a CVM cluster, this issue occurs if a disk having DRL plex is disconnected locally from the node where volume is online. Because of this, I/O writes to the DRL log fail and causes the DRL plex to detach on all the nodes in the cluster even if the disk is visible on the other node.

Workaround: To resolve this issue, reconnect the disk on the node where it is disconnected, and then reactivate the disk.

Error while converting a Microsoft Disk Management Disk Group created using iSCSI disks to an SFW dynamic disk group

This issue occurs while converting a Microsoft Disk Management Disk Group that was created using iSCSI disks to an SFW dynamic disk group. (3146530)

The following error appears because this operation is currently not supported:

This operation is not supported on disk groups created on iSCSI disks.

Workaround: There is no workaround for this issue.

Snapshot schedules intermittently fail on the Slave node after failover

This issue occurs when, for a cluster-shared volume, you fail over a snapshot schedule from a Slave node to another node. A snapshot schedule involves snapback and snapshot operations that are executed at scheduled intervals. Sometimes, the snapshot schedule operation may fail intermittently after you fail over to another node or fail over back to the original node. (3372066)

Workaround: There is no workaround for this issue.

In VEA GUI, Tasks tab does not display task progress for the resynchronization of a cluster-shared volume if it is offline

This issue occurs if a resynchronization operation is started on a cluster-shared volume that is offline. In VEA GUI, the volume status is displayed as "Resynching", but the Tasks tab in the lower pane does not display the resynchronization task progress. This issue is GUI-related and does not affect the resynchronization task. Also, this issue does not occur if the volume is online. (3395586)

Workaround: There is no workaround for this issue.

Snapback operation from a Slave node always reports being successful on that node, even when it's in progress or resynchronization fails on Master

As per the CVM functionality, all operations are performed on the Master node, including those initiated on a Slave node that are then shipped to and performed on Master. When you run a task from Slave that takes a longer time to complete (for example, snapback), then the commands on Slave report that the operation was successful as soon as the task is submitted to Master. Therefore, the progress of such tasks should be monitored on Master as any failure during the task execution would not be reported back to the Slave.

This issue occurs when, for a cluster-shared volume, you perform the snapback operation from a Slave node. While the operation is still being performed on Master, on Slave it is reported as being successful. This happens even when the resynchronization operation fails. (3283523)

Workaround: There is no workaround for this issue.

For cluster-shared volumes, only one file share per Microsoft failover cluster is supported

In Microsoft Failover Clustering environment, this issue occurs when you try to create a file share on the Volume Manager Shared Volume resource for a cluster-shared volume. In this case, the file share always gets created under one file share server (Client Access Point) even if it is not part of that role.

Windows Explorer always returns the same Client Access Point name while sharing a folder for any Volume Manager Shared Volume resource. Arctera needs Microsoft's assistance to resolve this issue. (3290426)

Workaround: There is no workaround for this issue.

After cluster-shared volume resize operation, free space of the volume is not updated on nodes where volume is offline

This issue occurs after you perform either the Volume Shrink or Volume Expand operation on a cluster-shared volume. If the cluster-shared volume is offline on a node, then the available amount of free space (under "Free space" property) is not updated after the volume resize operation.

(3232264)

Workaround: As a workaround, you need to bring the cluster-shared volume online on the node to see the updated amount of free space that is available.

Issues due to Microsoft Failover Clustering not recognizing SFW VMDg resource as storage class

Microsoft Failover Clustering currently does not recognize SFW Volume Manager Disk Group (VMDg) agent resource as a storage class resource. As a result, several issues are observed when you configure file shares and Microsoft applications with SFW in a Microsoft Failover Clustering environment.

Microsoft has acknowledged this as an issue and Arctera is actively working with Microsoft to get this fixed in Microsoft Failover Clustering. For more information, refer to the following Microsoft KB article:

[2804526](#)

The following issues are observed because of this limitation in a Microsoft Failover Clustering environment:

- Unable to configure file shares on volumes that are managed using SFW

If you create disk groups and volumes using SFW and then try to configure a VMDg, RVG, or Volume Manager Shared Volume resource in a Microsoft failover cluster, the Failover Cluster Manager snap-in (Microsoft Failover Clustering GUI) does not display the SFW volumes. As a result, you cannot configure file shares on those volumes. For more information, refer to Microsoft KB2795993 and KB2796000.

Workaround: As a workaround, you can use Powershell cmdlets or Windows Explorer to configure file shares. Then, using Microsoft Failover Clustering, you can enable continuous availability.

- Unable to perform storage related operations on SFW VMDg resource from Microsoft Failover Cluster Manager

If you create disk groups and volumes using SFW and configure a VMDg resource in a Microsoft failover cluster, then the Failover Cluster Manager snap-in does not list any storage related operations for the VMDg resource. This is seen when the VMDg resource is part of Available Storage group. The storage operations are enabled if the resource is part of a role. (2999555, 3000675)

For more information, refer to Microsoft KB2795997.

Note that these issues are restricted to file shares and Microsoft applications. You can however use SFW and configure custom applications in a Microsoft Failover Clustering environment. This issue does not affect configuration and failover operations for custom applications in Microsoft Failover Clustering.

If fast failover is enabled for a VMDg resource, then SFW volumes are not displayed in the New Share Wizard

This issue occurs while adding a file share using the New Share Wizard in a Microsoft Failover Clustering environment. If fast failover is enabled for the Volume Manager Disk Group (VMDg) resource, then the SFW volumes are not displayed in the wizard while adding the file share. (3049048)

Workaround: To resolve this issue, set the FastFailover attribute to False, create the file share, and then set the attribute to True again after creating the file share.

Volume information not displayed for VMDg and RVG resources in Failover Cluster Manager

This issue occurs while viewing volume information for a Volume Manager Disk Group (VMDg) or Replicated Volume Group (RVG) resource in a Microsoft Failover Clustering environment. The volume information is not displayed in the Failover Cluster Manager snap-in. This is a known Microsoft issue and Arctera has submitted a Design Change Request (DCR) to Microsoft for resolving the issue. (3004078, 3011313)

Workaround: There is no workaround for this issue.

Failover of VMDg resource from one node to another does not mount the volume when disk group volume is converted from LDM to dynamic

This issue occurs while performing failover of a Volume Manager Disk Group (VMDg) resource from one node to another. If a basic disk with partition created and mounted is added to the SFW dynamic disk group, then the partition gets converted from the Logical Disk Manager (LDM) disk group volume to the dynamic volume. In this case, if the disk group is configured under cluster and

FastFailover is set to "true" for the VMDg resource, then it fails to mount the volume after failback to the first node. (3027037)

Workaround: To resolve this issue, deport and import the dynamic disk group after converting the partition to dynamic volume, but before enabling fast failover for the VMDg resource.

vxverify command may not work if SmartMove was enabled while creating the mirrored volume

This issue occurs while using the `vxverify` command only if you had enabled SmartMove while creating the mirrored volume. When you use the `vxverify` command to compare mirrored volumes, it may return volume comparison errors in the output. However, these errors do not impact the functionality of the product and can be ignored. (3021565)

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 and High Availability Solutions (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 Clustering. (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.

Some operations related to shrinking or expanding a dynamic volume do not work

The following issues are observed while performing volume shrink or volume expand related tasks on a dynamic volume:

- While shrinking (decreasing) or expanding (increasing) a dynamic volume's size, using the Shrink Volume and Expand Volume dialog boxes, respectively, you can click the Max Shrink and Max Size buttons to know the maximum amount by which a volume can be shrunk or expanded. These buttons do not work as expected for both the operations because of a Microsoft Virtual Disk Service (VDS) error.
- While performing the `vxassist shrinkby` or `vxassist querymax` operation for a newly-created volume, the operation fails with the "Invalid Arguments" error.

- While attempting the volume shrink operation, the "Invalid Arguments" error occurs and the Event Viewer displays a VDS provider failure error.

(2998422, 2411143, 2405311)

Workaround: To resolve any of these issues, restart VDS by using the following commands, and then try again:

1. `Net stop vds`
2. `Taskkill /f /im vxvds.exe`
3. `Taskkill /f /im vxvdsdyn.exe`
4. `Net start vds`
5. `Vxassist refresh`

VEA GUI displays error while creating partitions

When you create partitions using the VEA, sometimes the GUI displays an error indicating that the operation has failed.

VEA displays the following messages:

- Failed to format volume. Investigate further based on operation return status.
- Failed to format volume \Device\Harddisk#\Partition#

The partition is created successfully, however VEA sometimes is unable to format the partition and displays these errors. This issue occurs intermittently. You can reformat the partition using the command line. (3000941)

Note that this issue is limited only to the VEA GUI; this issue does not occur when you perform these operations using the command line.

For dynamic disk group configured as VMNSDg resource, 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, if the dynamic disk group is configured as a VMNSDg resource, then 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.

In Microsoft failover cluster, 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 Clustering 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.

In some cases, updated VSS components are not displayed in VEA console

This issue occurs while adding or removing the VSS components or when connecting to the VEA console for the first time. During this, the updated VSS components are not displayed in the VEA console.

Workaround: To resolve this issue, you must manually refresh the VEA using either the `vxsnap refresh` command or the Refresh option in the VEA console.

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\(W.S.10\).aspx](http://technet.microsoft.com/en-us/library/dd183729(W.S.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:

```
$HypervVGuest = $args[0]
```

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

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

```
{  
  
$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.

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 Arctera 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).

For more information, refer to *Storage Foundation and Disaster Recovery Solutions for Microsoft Hyper-V*.

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

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. (2509422)

In such cases, the following is observed:

- The SFW messaging for Live Migration support will not work between the remaining nodes
- The VMDg 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.

For more information, refer to *Storage Foundation and Disaster Recovery Solutions for Microsoft Hyper-V*.

A remote partition is assumed to be on the local node due to Enterprise Vault 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 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.

Messaging service does not retain its credentials after upgrading SFW and SFW HA

After upgrading SFW and SFW HA from version 5.x to its latest version, the Veritas InfoScale Messaging Service does not retain its credentials, if they are configured on a Domain in the Active Directory. (2529295)

Workaround: Manually reset the service credentials after upgrade.

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 Veritas 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.

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 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)

Microsoft failover cluster 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 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 Microsoft failover 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`

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)

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 *Volume Replicator Administrator's Guide*.

1. 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.
2. Finish running the DR wizard to complete the service group cloning, replication configuration, and global cluster option (GCO) configuration.
3. 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.

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

In case of a clustered (VCS or DAG/Microsoft failover cluster) 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.

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.

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. In regedit, navigate to
`HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\lanmanserver`
2. Right-click the lanmanserver node and select **New>Multi-String Value** to enter a new REG_MULTI_SZ entry.
3. 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.)
4. 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).

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.

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 be 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)

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).

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)

Mirrored volume in Microsoft Disk Management Disk Group does not resynchronize

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)

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 *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 *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.

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.

VEA console issues

This section lists the issues related to the Veritas Enterprise Administrator (VEA) console.

Login to VEA on an IPv6-enabled system with the Logged On User on this computer option may cause incorrect privileges to be assigned

If you log in to VEA on an IPv6-enabled system by using the Logged On User on this computer option, you may sometimes not be able perform any operations. (3940157)

When you log in to VEA by using the Logged On User on this computer option, you typically do not need to provide any credentials. However, when you log in to VEA by using this option on an IPv6-enabled system, you may sometimes not be able to perform any further operations. This issue occurs because of the lack of appropriate privileges, which might have to do with AD issues in an IPv6 domain.

Workaround:

If you encounter this issue, log in to VEA again by using the Connect using a different user account option. The appropriate privileges are assigned when you provide the user name and password for the login.

VEA may fail to start when launched through the SCC, PowerShell, or Windows Start menu or Apps menu

This issue occurs if a space character followed by a hyphen is included in the product installation directory path.

By default, VEA is installed with the following products:

- InfoScale Foundation
- InfoScale Storage
- InfoScale Enterprise

While installing any of these products, you can specify a custom installation directory. If that directory path includes a space character followed by a hyphen, VEA fails to start when you launch it through the SCC, PowerShell, or the Windows Start menu or Apps menu. (3796076)

Workaround: Navigate to the appropriate location in the product installation directory, and launch VEA from the command line (`vea.exe`).

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.

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 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.

Diskgroup creation fails on Ultradisk in Azure cloud for logical sector size 4096-byte sector (4101641)

In Azure cloud if you create the dynamic cluster diskgroup on shared Ultradisk with logical sector size 4096 bytes, then VEA gives the following error: "Disk sector size is not supported by Volume Manager".

Note: In general, 4K sector is not supported on SFW (irrespective of cloud)

Snapshot and restore issues

The section lists the known issues related to the snapshot and restore operations.

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.

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 on SQL Server component with missing volume fails

A restore operation on a SQL Server component that has a missing volume fails. This issue occurs 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: Before you start the restore operation, change the drive letter of the snapshot volume to the drive letter of the original volume.

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](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.

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.

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)

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. Refer to the Microsoft Knowledge Base for details.

vxsnapsql restore may fail to restore SQL Server database

When the VCS, VVR, and GCO options are configured on a system, if you use the `vxsnapsql restore` command to restore a SQL Server database, the operation may fail. (895239)

The following error message is logged when the operation fails:

```
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: First, take all the SQL Server and the MountV resources offline for the volume that contains the SQL Server database and the logs. Then, bring them back online. Thereafter, the `vxsnapsql restore CLI` command works correctly.

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.

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 section lists the known issues related to the snapshot scheduling operations.

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.

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 Recovery Wizard 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>scs). 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 a Microsoft failover cluster node while another cluster node is shutting down

If you are creating or deleting a snapshot schedule on a Microsoft failover 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 (`VxSVC` 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.

Multi-pathing issues

This section lists the known issues related to the multi-pathing components; these components are available as part of the InfoScale Foundation, InfoScale Storage, and InfoScale Enterprise products.

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)

vxddmpadm's deviceinfo and pathinfo with disk specified in p#c#t#l# parameter displays information only by one path

The `deviceinfo` and `pathinfo` commands of `vxddmpadm` work with only one `p#c#t#l#` parameter shown with the disk in the `vxddmpadm disk list` command. Even if the disk has more than one path, the `deviceinfo` and `pathinfo` commands with `p#c#t#l#` values of other paths fail with the Invalid Argument error. (2162670)

Replication issues

This section lists the known issues related to the replication components; these components are available as part of the InfoScale Storage and InfoScale Enterprise products.

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.

VVR replication fails to start on systems where Symantec Endpoint Protection (SEP) version 12.1 or 12.1 RU2 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 or 12.1 RU2 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.

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 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.

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 - `sr1prot` - 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 `sr1prot` 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`

1. Modify the registry DWORD value for the AE_TIMEOUT entry, from the default value of 30 seconds to 60 seconds or higher.
2. In order for the registry key change to take effect, type the following at the command prompt:

```
vxassist refresh
```

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.

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
```

1. Modify the registry DWORD value for the AE_TIMEOUT entry, from the default value of 30 seconds to 60 seconds or higher.
2. 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

A pause replication operation with checkpoints from the Secondary host may fail for heavy I/Os and low-bandwidth network. If the request of the Secondary for rlink checkpoint for Pause to Primary times out before the Primary acknowledges the request, the pause operation fails. (1276014)

Workaround:

To avoid this issue, perform one of the following tasks:

- Pause the secondary RVG by selecting the Pause Secondary option from the secondary RVG context menu. If it fails, slow down the I/O to the Primary host and retry. The 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 context 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. Arctera 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.

CLI shows the "Volume in use" error when you dismount the ReFS data volumes on the Secondary RVG.

When you dismount the ReFS data volumes on the Secondary RVG, you may get the following error message "Volume in use" even when no drive letter has been assigned. This happens because, when an RLINK is attached on the Secondary RVG, the volume is marked as read-only, and VVR is unable to lock the volume. (3789197)

Note: This is an operating system issue and not a VVR issue.

Solution configuration issues

This section lists the known issues related to configuring, administering, and unconfiguring the various InfoScale solutions for the supported applications and hardware replication environments.

Oracle Enterprise Manager cannot be used for database control

In this release, you cannot use Oracle Enterprise Manager for database control. (364982)

Unexplained errors with DR wizard and QR wizard

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)

VCS FD and DR wizards fail to configure application and hardware replication agent settings

This issue is observed in the following scenarios, when InfoScale Storage and InfoScale Availability co-exist on the same system (3873271):

- Configuring the application service group using the Disaster Recovery (DR) Wizard
- Configuring the fire drill service group using the Fire Drill (FD) Wizard

During the service group configuration, the wizards attempt to check if the Storage and Availability components are installed on all the systems at both the sites.

In case of a co-existence scenario, the wizards fail to validate the installation of Availability components and thus does not configure the application and the hardware replication settings.

Due to this issue, you cannot proceed with the service group configuration.

Workaround:

To resolve the issue, perform the following steps on all the systems at both the sites:

1. Navigate to the following path and create a registry key with any suitable name:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Veritas\
VPI\{F834E070-8D71-4c4b-B688- 06964B88F3E8}\Solutions\
```

1. Create a string value with display name as "DisplayTitle".
2. For the created string, set the value as "Veritas InfoScale Enterprise 9.1 for Windows".

Disaster recovery (DR) configuration issues

This section lists the known issues that you may encounter when working with DR configurations for the supported applications.

The Disaster Recovery Configuration Wizard or the Fire Drill Wizard cannot proceed when configuring an application in an EMC SRDF replication environment

When configuring an application for disaster recovery (DR) or for fire drill (FD) in an EMC SRDF replication environment, the wizard might encounter an issue and fail to proceed. (3412223, 3384096)

The wizard might display the following error:

```
V-52410-49479-116

0

Failed to discover 'SRDF' on node 'clusterNodeName'
```

This issue occurs when the wizard is unable to parse a SYMCLI output (XML file) that describes the replication configuration.

Workaround: To resolve this issue, ensure that the replication array is configured correctly.

Then, run the following EMC Symmetrix command on the system where you want to configure the application for DR or FD:

```
symrdf list -v pd -output xml >> XMLFilePath
```

If the generated XML is valid, it indicates that the array is configured correctly. Launch the wizard again to configure DR or FD.

If the generated XML is invalid or if you still encounter this error when running the wizard, you might want to configure DR or FD manually.

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 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: 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 *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.

Disaster Recovery (DR) Wizard fails to automatically set the correct storage replication option in case of SRDF

This is observed when you are setting up a disaster recovery cluster configuration with SRDF replication.

When you launch the DR wizard to configure the DR site, the wizard typically detects the underlying storage environment and automatically selects the appropriate replication option on the Replication Options panel.

However, it fails to detect SRDF replication and VVR replication option is selected by default. (3020038)

Workaround

You must manually choose SRDF replication option on the DR wizard panel and then proceed with the DR configuration.

Disaster Recovery (DR) Wizard reports an error during storage cloning operation in case of SRDF

This is observed when you are configuring disaster recovery in an SRDF replication environment using the DR wizard and you choose the temporary storage cloning option. (3019858, 3019876)

The DR wizard attempts to create the temporary disk group and volumes at the secondary site, but fails with the following error messages:

```
ERROR V-16-0-0 (1420:2800)
```

```
(:CVMPPlugin::GetErrorString - UMI message : :0)      Unable to reserve a  
majority of dynamic disk group members.
```

```
Failed to start SCSI reservation thread.
```

```
ERROR V-16-0-0 (1420:2800)
```

```
(:VMAPI: CVMPPlugin::__CommitCreateDGorAddDiskstoDG::
```

```
VxVmUpgradeToDynamicDiskEx failed:0)
```

```
VMAPI: VxVmUpgradeToDynamicDiskEx failed.
```

Workaround

Even if the DR wizard fails with these errors, the disk group creation operation is successful in the background. You must launch the DR wizard again and complete the wizard flow to perform the remaining operations.

Fire drill (FD) configuration issues

This section lists the known issues that you may encounter when working with FD configurations for the supported applications.

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 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 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. 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 VIOM 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 VIOM console or Java GUI cannot change the value of this attribute. (2735936)

Therefore, if you use the VIOM 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 *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.

Quick recovery (QR) configuration issues

This section lists the known issues that you may encounter when working with QR configurations for the supported applications.

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.

Internationalization and localization issues

This section lists the known issues related to using the InfoScale products 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.

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.

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 on a Windows Server

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.

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

If you install the InfoScale 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 InfoScale High Availability Console host fails. (2910613)

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

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 InfoScale High Availability Configuration wizard.(2906207)

While configuring application monitoring using the InfoScale 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 InfoScale High Availability tab and the InfoScale 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: 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".

Interoperability issues

This section lists the known issues related to the coexistence and usage of the InfoScale products with other software.

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.

VCS services do not start on systems where SEP 12.1 or later is installed

The following issues may occur if you install and configure VCS on systems where Symantec EndPoint Protection (SEP) 12.1 or a later version is installed.

1. In an IPv6 environment, the High Availability Engine (HAD) service may fail to start. (2439737, 2487369)

The following error may be displayed:

```
Failed to start the cluster. Error=FFFFFFFF.
```

```
Failed to start services on all the nodes.
```

1. 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: Configure the SEP Firewall policy to allow IPv6 traffic on the cluster nodes.

Edit the Firewall Rules table and enable the following settings:

1. Block IPv6: Set the Action to "Allow"
2. Block IPv6 over IPv4 (Teredo): Set the Action to "Allow"

3. Block IPv6 over IPv4 (ISATAP): Set the Action to "Allow"

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 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 the product.
- 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 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 `llttab.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
```

Miscellaneous issues

This section lists the known issues that common to the cluster management, storage management, replication, and multi-pathing components of the InfoScale products.

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, Arctera recommends that you perform the following steps in the given order.

To modify the private network adapter properties

1. Stop the 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
```

1. Stop the VCS LLT service on the node.

Type the following on the command prompt:

```
net stop llt
```

1. Modify the network properties of the network adapters on the node, as desired.
2. Start the High Availability Engine (HAD) on the node.

Type the following on the command prompt:

```
hastart
```

1. Repeat step 1 to step 4 on all the other passive nodes in the cluster.
2. 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.
3. Repeat step 1 to step 4 on the active node.
4. 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.

Cross grade from Foundation to Enterprise and Foundation to Storage fails (4193597)

Workaround:

Uninstall InfoScale Foundation and install Enterprise.

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.

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

51030	UNABLE TO FIND A SUITABLE REMOTE FAILOVER TARGET FOR GLOBAL GROUP %S.
	ADMINISTRATIVE ACTION IS REQUIRED.
	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

51030	UNABLE TO FIND A SUITABLE REMOTE FAILOVER TARGET FOR GLOBAL GROUP %S.
	ADMINISTRATIVE ACTION IS REQUIRED.
	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.

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 VCSHigh 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 llc
```

1. On a node on which VCS was upgraded successfully, open the file `llchosts.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 `llchosts.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 `llchosts.txt` file should look like this:

```
\# This is program generated file, please do not edit

0 N1

1 N2
```

The `llchosts.txt` file is typically located at `C:\Program Files\VERITAS\comms\llc`.

Here C:\ is the drive on which VCS is installed.

1. 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.

1. From the Windows Services snap-in, change the startup type of the High Availability Engine (HAD) service to Manual.
2. Repeat step 2, step 3, and step 4 on all the nodes on which VCS was upgraded successfully.
3. 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`.

1. 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\vp` 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 Arctera 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: Arctera 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.

Live migration of a VM, which is part of a VCS cluster where LLT is configured over Ethernet, from one Hyper-V host to another may result in inconsistent HAD state

Consider the following scenario:

- Two or more Hyper-V hosts are set up, between which live migration can be performed.
- Two or more virtual machines (for example, VM1 and VM2) are configured as nodes of a VCS cluster on one of the hosts.
- In the VCS cluster, LLT is configured over Ethernet.

Perform live migration of one of the virtual machines (say VM1, which may be the active node).

After live migration, the HAD state is reported as follows:

- VM1 shows the HAD state as RUNNING for both the nodes.
- VM2 shows the HAD state as FAULTED for VM1, but RUNNING for VM2.

Events such as the following may be seen in the System Event Viewer for VM1 (one event for each LLT link):

```
LLT ERROR V-14-1-10085 LLT protocol is unbinding from link adapterID
```

This issue does not occur when LLT is configured over UDP. (3053241, 3056450)

Workaround: To avoid this issue, you might want to configure LLT over UDP.

If you choose to configure LLT over Ethernet and if you encounter this issue, perform the following steps after live migration:

1. Forcibly stop the High Availability Engine (HAD) service on the migrated node using the following command:

```
taskkill /f /im had.exe
```

- a. If the HAD service starts again, stop it using the following command:

```
hastop -local
```

2. Verify that the HAD service is in the stopped state.

3. Run the following commands sequentially on the migrated node:

```
net stop vcscomm
```

```
net stop gab
```

```
net stop llt
```

1. On the migrated node, restart the VCS cluster using the following command:

```
hastart
```

- a. Verify that all the cluster nodes report a consistent state for the HAD service using the following command:

```
hasys -state
```

If a NIC that is configured for LLT protocol is disabled, LLT does not notify clients

This issue occurs if a NIC configured for LLT is disabled using the Windows Network and Sharing Center. (3396648)

In a VCS cluster, GAB maintains cluster membership by receiving input on the status of the heartbeat from each node by LLT. When a system no longer receives heartbeats from a peer, it marks the peer as DOWN and excludes the peer from the cluster.

If the NIC is disabled using the Windows Network and Sharing Center, the LLT links are disabled and LLT does not send notification to clients. As a result, GAB does not mark the peer node as DOWN and considers the peer as active.

Workaround: To resolve this issue, restart the VCS LLT service.

Perform the following steps, using the command line

1. Stop the VCS LLT service on the node.

Type the following on the command prompt:

```
net stop llt
```

1. For VCS, restart the VCS cluster using the following command:

```
hastart -all
```

For SFW, type the following command:

```
net start vcscmm
```

Recommendation: Arctera recommends that you do not disable the NIC configured for LLT protocol.

Fibre Channel adapter issues

This section lists the known issues related to using specific Fibre Channel host bus adapters with the InfoScale products.

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.

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.

Storage agent issues and limitations in VMware virtual environments

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 InfoScale 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)