

Veritas InfoScale™ 7.2

What's New In This Release

- AIX, Linux, Solaris

Veritas InfoScale 7.2 What's new in this release

Last updated: 2016-11-03

Document version: 7.2 Rev 0

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third party software for which Veritas is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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

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

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

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

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

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

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

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

What's new in this release

This document includes the following topics:

- [About this document](#)
- [About the Veritas InfoScale product suite](#)
- [About the Dynamic Multi-Pathing for VMware component](#)
- [Components of the Veritas InfoScale product suite](#)
- [Support for up to 128 nodes in a cluster](#)
- [About Just In Time Availability](#)
- [Support for migrating applications from one cluster to another](#)
- [Automatically provision storage for Docker Containers](#)
- [Platforms supported](#)
- [Not supported in this release](#)
- [Changes related to Veritas Cluster Server](#)
- [Changes related to installation and upgrades](#)
- [Changes related to Veritas File System](#)
- [Changes related to Veritas Volume Manager](#)
- [Changes related to Replication](#)
- [Changes related to Operating System](#)

About this document

This document covers the major new branding, licensing, and installation changes that are introduced in 7.2.

Sections specify supported platforms. If no platform is mentioned, it means that the change is applicable across all platforms.

About the Veritas InfoScale product suite

The Veritas InfoScale product suite addresses enterprise IT service continuity needs. It draws on Veritas' long heritage of world-class availability and storage management solutions to help IT teams in realizing ever more reliable operations and better protected information across their physical, virtual, and cloud infrastructures. It provides resiliency and software defined storage for critical services across the datacenter infrastructure. It realizes better Return on Investment (ROI) and unlocks high performance by integrating next-generation storage technologies. The solution provides high availability and disaster recovery for complex multi-tiered applications across any distance. Management operations for Veritas InfoScale are enabled through a single, easy-to-use, web-based graphical interface, Veritas InfoScale Operations Manager.

The Veritas InfoScale product suite offers the following products:

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

About the Dynamic Multi-Pathing for VMware component

Dynamic Multi-Pathing for VMware 7.2 (VxDMP) is a multi-pathing solution integrated with VMware's vSphere infrastructure, which brings the established and proven enterprise-class functionality to VMware virtual environments.

In Veritas InfoScale 7.2, there are two installers. The Veritas InfoScale installer does not install the Dynamic Multi-Pathing for VMware component. To install the Dynamic Multi-Pathing for VMware component, you must use one of the following:

- Veritas_InfoScale_Dynamic_Multi-Pathing_7.2_VMware.zip
- Veritas_InfoScale_Dynamic_Multi-Pathing_7.2_VMware.iso

For more information about the Dynamic Multi-Pathing for VMware component, refer to the following guides:

- *Dynamic Multi-Pathing Installation Guide - VMware ESXi*
- *Dynamic Multi-Pathing Administrator's Guide - VMware ESXi*

Components of the Veritas InfoScale product suite

Each new InfoScale product consists of one or more components. Each component within a product offers a unique capability that you can configure for use in your environment.

[Table 1-1](#) lists the components of each Veritas InfoScale product.

Table 1-1 Veritas InfoScale product suite

Product	Description	Components
Veritas InfoScale™ Foundation	Veritas InfoScale™ Foundation delivers a comprehensive solution for heterogeneous online storage management while increasing storage utilization and enhancing storage I/O path availability.	Storage Foundation (SF) Standard (entry-level features)
Veritas InfoScale™ Storage	Veritas InfoScale™ Storage enables organizations to provision and manage storage independently of hardware types or locations while delivering predictable Quality-of-Service, higher performance, and better Return-on-Investment.	Storage Foundation (SF) Enterprise including Replication Storage Foundation Cluster File System (SFCFS)
Veritas InfoScale™ Availability	Veritas InfoScale™ Availability helps keep an organization's information and critical business services up and running on premise and across globally dispersed data centers.	Cluster Server (VCS) including HA/DR

Table 1-1 Veritas InfoScale product suite (*continued*)

Product	Description	Components
Veritas InfoScale™ Enterprise	Veritas InfoScale™ Enterprise addresses enterprise IT service continuity needs. It provides resiliency and software defined storage for critical services across your datacenter infrastructure.	Cluster Server (VCS) including HA/DR Storage Foundation (SF) Enterprise including Replication Storage Foundation and High Availability (SFHA) Storage Foundation Cluster File System High Availability (SFCFSHA) Storage Foundation for Oracle RAC (SF Oracle RAC) Storage Foundation for Sybase ASE CE (SFSYBASECE)

Support for up to 128 nodes in a cluster

(Linux only)

Veritas InfoScale now supports cluster configurations up to 128 nodes. The support is however limited to specific features for configurations beyond 64 nodes.

If you want to configure an InfoScale feature or capability (that is not mentioned in the following list) on clusters comprising more than 64 nodes, contact Veritas Technical Support.

The supported features are as follows:

- Concatenated and striped volume layouts
- Thin provisioning and thin reclamation
- Veritas File Replicator (VFR)
- Volume grow and shrink operations
- Master switching and preference
- Command shipping
- Private region I/O shipping

- Split and join operations on disk groups
- Volume relayout
- Hot relocation
- SCSI-3 I/O fencing for avoiding network split brain
- Data disk fencing for preventing I/Os from nodes that are not part of cluster
- LLT over Ethernet
- Online Co-ordination Point Replacement (SCSI-3 to SCSI-3 mode)

The unsupported features are as follows:

- Mirrored volume layouts—mirror, concatenated mirror (concat-mirror), mirrored concatenated (mirror-concat), striped-mirror, mirror-stripe
- Fast mirror resynchronization (FMR)
- Volume snapshots
- Veritas Operations Manager (VOM) for managing the cluster and storage environment
- Veritas Volume Replicator (VVR)
- Flexible Storage Sharing (FSS)
- Campus cluster
- Public region I/O shipping
- CVM-DMP protocol support
- Disk cloning
- Auto-refresh feature of the Co-ordination Point agent
- Customized fencing
- LLT over UDP and RDMA
- AdaptiveHA
- Priority-based failover for service groups
- Multi-site management
- Virtual Business Services (VBS)
- Docker Containers
- Auto-Clear for service groups
- MonitorOnly for service groups

About Just In Time Availability

(Linux only)

The Just In Time Availability solution provides increased availability to the applications on a single node InfoScale Availability cluster in VMware virtual environments.

Using the Just In Time Availability solution, you can create plans for:

1. Planned Maintenance
2. Unplanned Recovery

Planned Maintenance

In the event of planned maintenance, the Just In Time Availability solution enables you to clone a virtual machine, bring it online, and failover the applications running on that virtual machine to the clone on the same ESX host. After the maintenance procedure is complete, you can failback the applications to original virtual machine. Besides failover and failback operations, you can delete a virtual machine clone, view the properties of the virtual machine and its clone, and so on.

Unplanned Recovery

When an application encounters an unexpected or unplanned failure on the original virtual machine on primary ESX, the Just In Time Availability solution enables you to recover the application and bring it online using the unplanned recovery feature.

With **Unplanned Recovery Policies**, the Just In Time Availability solution enables you to set up recovery policies as per your requirement to mitigate the unplanned failure that is encountered by an application. Just In Time Availability solution provides the following recovery policies for your selection. You may select one or all the recovery policies as per your need.

For more information see the InfoScale Solution Guide - Linux

Support for migrating applications from one cluster to another

The Application Migration add-on allows you to migrate applications that are under Cluster Server management from one cluster to another. The application migration operation is less complex and can be accomplished with minimal manual intervention. The application migration can be across operating systems, architectures, or virtualization technologies. In this release, you can migrate an application between different:

- Platforms—AIX, Linux, and Solaris
- Environments—Physical-to-physical, physical-to-virtual, virtual-to-virtual, and virtual-to-physical
- InfoScale versions

To migrate an application, you must create an application migration plan using the **Create Migration Plan** wizard. After you create a plan, you must execute the migration plan.

The add-on also allows you to:

- Pause and resume the operation for manual verification and correction, if required.
- Integrate custom scripts in the operation as per application requirements.
- Migrate application dependencies.
- Understand source cluster configuration and create target cluster configuration.
- Perform endian changes to the data as per architecture requirements.
- Rehearse the steps before the actual migration operation.

For more information, see the *Veritas InfoScale Operations Manager 7.2 Add-ons User's Guide*.

Automatically provision storage for Docker Containers

(Linux only)

Veritas InfoScale volume driver plugin for Docker extends the capability of Docker daemon to handle storage related operations such as creation of volumes or file systems, mount or unmount file systems, and other storage functions. With the volume plugin, docker containers can be started with storage attached to them automatically. It eases deployment of Docker containers. By using Veritas InfoScale storage, you can use all capabilities of Veritas InfoScale products. The Veritas driver supports Docker version 1.9 or later and it also integrates with docker volume CLI. The plugin driver also seamlessly works with Docker Swarm technology.

Platforms supported

For a detailed info of platforms supported see *Veritas InfoScale™ 7.2 Release Notes*.

Not supported in this release

(Solaris only)

The following features are not supported in this release:

- The Solaris 10 operating system is not supported.
- Partial upgrade and co-existence are not supported on the Solaris11x64 operating system.

Changes related to Veritas Cluster Server

The following section describes the changes introduced in Veritas Cluster Server (VCS) 7.2.

VCS supports up to 128 nodes in a cluster

(Linux only)

VCS supports up to 128 nodes in a cluster. The `numnids` static tunable parameter in the GAB configuration file can now accept values between 1-128. Likewise, the `maxnids` value in the LLT module can now accept values between 1-128. Any new values set for this tunable take effect only upon a system reboot.

systemD support for Oracle application service

(Linux only)

(RHEL 7 and SLES 12) systemD is a system and service manager for Linux operating systems. With systemD one of the improvements is that applications can be started as a unit service.

With `UseSystemD` attribute enabled in RHEL 7 or SLES 12, the Oracle resource comes online as a unit service in `system.slice` during application start. Without the `UseSystemD` attribute enabled, a typical online entry point starts the resource in `user.slice`. Starting the application unit service in `system.slice` avoids the possibility of an Oracle database crash. Further, configure application-specific environments by assigning key-value pairs to the `SystemAttrDList` attribute.

Vxfen debuglog buffer size changes

(Linux only)

The Vxfen debuglog buffer size is increased to support up to 128 nodes in a cluster.

- Default: 524288(512KB)

- Minimum: 65536(64KB)
- Maximum: 1048576(1MB)

RVGSharedPri agent supports multiple secondaries

After successful migration or takeover of a Secondary RVG, the RVGSharedPri agent automatically starts the replication from the new Primary to any additional Secondary(s) that exists in the Replicated Data Set (RDS).

Mount agent: Support for ext4 and XFS file systems

(Linux only)

Intelligent Monitoring Framework (IMF) for mounts is now supported on ext4 and XFS file systems.

New attributes in VMwareDisks agent

(Linux and Solaris)

The following section describes the attributes introduced in this release:

VMwareDisks agent

HAInfoDetails	<p>Determines whether or not vSphere HA is enabled. This attribute uses the vCenter Server hostname or IP address to determine the status of vSphere HA.</p> <p>The value must be specified in the format: Key=Value. Where:</p> <ul style="list-style-type: none"> ■ Key= vCenter Server hostname or IP address ■ Value=vCenter Server logon user credentials. This must be specified in the format: User name=Encrypted password <p>If you do not specify a value for this attribute, the agent considers the vSphere HA setting based on the IsVMHAEnabled attribute value.</p> <p>Type and dimension: string-association</p>
PanicVMOnESXLoss	<p>Set this attribute value to 1 (True) to trigger panic on the virtual machine when the ESX host loses network connectivity.</p> <p>Default: 0 (False)</p> <p>Type and dimension: boolean-scalar</p>
ForceRegister	<p>For internal use only.</p>

Changes related to installation and upgrades

The following changes are introduced to the installation and upgrading of Veritas InfoScale 7.2.

Partial upgrade and co-existence are not supported on Solaris 11

The Veritas InfoScale 7.2 release does not support partial upgrade and co-existence on the Solaris 11x64 operating system. Upgrade on the Solaris 11x64 operating system is only supported from VCS or InfoScale Availability.

Solaris 10 is not supported in 7.2 release

The Veritas InfoScale 7.2 release does not support the Solaris 10 operating system.

Changes related to Veritas File System

The following changes are introduced to Veritas File System (VxFS) of Veritas InfoScale 7.2.

Support to mount an inconsistent file system

In older DLVs an inconsistent file system would have stayed mounted and functional. However, if the same was unmounted for any reason, it was not possible to mount the same without a complete `fsck`. In 7.1, it is possible to attempt to mount an inconsistent file system, so that any immediate need of data availability is met. Also, a full `fsck` can be scheduled at a more convenient time. For more info see the `delayfsck` option in the `mount_vxfs (1M)` command.

Intent log version of Veritas File System incremented to 13

(Linux only)

The intent log version of the Veritas File System is now incremented to 13.

Migrate VxFS file system from 512-bytes sector size devices to 4K sector size devices

(Linux and Solaris)

With Veritas InfoScale 7.2, you can migrate VxFS file system from 512 bytes to 4K sector size devices.

Migration of VxFS file system from 512-bytes sector size to 4K sector size is supported only on Linux (RHEL and SLES) and Solaris 11 operating systems.

For more information see the InfoScale storage administration guides.

Technology Preview: Distributed SmartIO in Veritas InfoScale storage environments

(Linux only)

Distributed SmartIO is a new feature available as a technology preview in Veritas InfoScale for configuration and testing in non-production environments. It is primarily targeted for Oracle RAC or ODM.

With the advancements in the hardware technology - network interconnect such as infiniband, accessing and sharing the data using the network rather than the disk as a medium for data sharing is proving to be faster and cost efficient in storage environments. Data can be cached on faster but costlier SSD storage on few nodes in a cluster. High speed network interconnect can be used to fetch the data as required on any node in the cluster.

Considering these benefits, Veritas InfoScale has come up with a robust solution, **Distributed SmartIO**, which lets you share SSD resources between all the nodes in the cluster for caching frequently read data.

Supported operating systems

You can configure Distributed SmartIO on supported versions of Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) in this release.

For more information, see Veritas InfoScale™ 7.2 SmartIO for solid-state drives Solutions Guide - Linux.

Changes related to Veritas Volume Manager

(Linux and Solaris)

The following changes are introduced to Veritas Volume Manager (VxVM) of Veritas InfoScale 7.2.

Veritas InfoScale 7.2 support with 4K sector devices

(Linux and Solaris)

Veritas InfoScale 7.2, using Veritas Volume Manager and Veritas File System storage components provides a solution which supports the 4K sector devices (formatted with 4KB) in storage environment. Earlier, you were required to format

4K devices with 512-bytes. From Veritas InfoScale 7.2 release, you can directly use the 4K sector devices with Veritas InfoScale without any additional formatting.

You can use 4k sector devices with Veritas InfoScale 7.2 only on Linux (RHEL and SLES) and Solaris 11 operating systems.

For more information see the InfoScale storage administration guides.

Application isolation in CVM environments with disk group sub-clustering

(Linux only)

Veritas InfoScale supports application isolation in a CVM cluster through the creation of disk group sub-clusters. A disk group sub-cluster consists of a logical grouping of nodes that can selectively import or deport shared disk groups. The shared disk groups are not imported or deported on all nodes in the cluster as in the traditional CVM environment. This minimizes the impact of node failures or configuration changes on applications in the cluster.

You can enable the application isolation feature by setting the `CVMdGSubClust` attribute for the `CVMCluster` resource in the VCS configuration file. When the cluster restarts, the feature is enabled and shared disk groups are not auto-imported to all nodes in the cluster. The first node that imports the disk group forms a disk group sub-cluster and is elected as the disk group master for the sub-cluster. The remaining nodes in the cluster that import the shared disk group are treated as slaves. All disk group level operations run on the master node of the disk group sub-cluster. You can switch the master at any time for each disk group sub-cluster. A node can play the role of a master for a sub-cluster as well as that of a slave for another sub-cluster.

If a node loses connectivity to the SAN, the I/Os for that node are shipped to another node in the disk group sub-cluster just as in traditional CVM environments. If the disk group fails due to failed I/Os on all disks in the disk group, it is disabled and the nodes that share the disk group must deport and import the disk group again.

A node can belong to multiple disk group sub-clusters. Each disk group sub-cluster provides all the capabilities of a clustered Veritas Volume Manager environment, with the exception of some features.

The following CVM features are not available in a disk group sub-cluster:

- Rolling upgrade
- Campus cluster configurations in CVM
- Move and Join operations with different disk group sub-cluster masters (source and target disk group)
- Clustered Volume Replicator

- Clone devices

The application isolation feature is supported with CVM protocol version 160 and above. It is disabled, by default, both after installation and upgrade.

Hot-relocation in FSS environments

(Linux only)

In FSS environments, hot-relocation employs a policy-based mechanism for healing storage failures. Storage failures may include disk media failure or node failures that render storage inaccessible. This mechanism uses tunables to determine the amount of time that VxVM waits for the storage to come online before initiating hot-relocation. If the storage fails to come online within the specified time interval, VxVM relocates the failed disk.

VxVM uses the following tunables:

`storage_reloc_timeout` Specifies the time interval in minutes after which VxVM initiates hot-relocation when storage fails.

`node_reloc_timeout` Specifies the time interval in minutes after which VxVM initiates hot-relocation when a node fails.

The default value for the tunables is 30 minutes. You can modify the tunable value to suit your business needs. In the current implementation, VxVM does not differentiate between disk media and node failures. As a result, both tunables will have the same value. For example, if you set the value of the `storage_reloc_timeout` tunable to 15, then VxVM will set the value of the `node_reloc_timeout` tunable also to 15. Similarly, if you set the `node_reloc_timeout` tunable to a specific value, VxVM sets the same value for the `storage_reloc_timeout` tunable. You can use the `vxtune` command to view or update the tunable settings.

The hot-relocation process varies slightly for DAS environments as compared to shared environments. When a DAS disk fails, VxVM attempts to relocate the data volume along with its associated DCO volume (even though the DCO may not have failed) to another disk on the same node for performance reasons. During relocation, VxVM gives first preference to available spare disks, failing which VxVM looks for eligible free space.

Hot-relocation in FSS environments is supported only on new disk groups created with disk group version 230. Existing disk groups cannot be used for relocation.

For more information see the InfoScale storage administration guides.

Technology Preview: Erasure coding in Veritas InfoScale storage environments

(Linux only)

Erasure coding is a new feature available as a technology preview in Veritas InfoScale for configuration and testing in non-production environments. It is supported in DAS, SAN, FSS, and standalone environments.

Erasure coding offers a robust solution in redundancy and fault tolerance for critical storage archives. In erasure coding, data is broken into fragments, expanded and encoded with redundant data pieces and stored across different locations or storage media. When one or more disks fail, the data on failed disks is reconstructed using the parity information in the encoded disks and data in the surviving disks.

Erasure coding can be used to provide fault tolerance against disk failures in single node (DAS/SAN) or shared cluster (SAN) setups where all nodes share the same storage. In such environments, erasure coded volumes are configured across a set of independent disks.

In FSS distributed environments, where the storage is directly attached to each node, erasure coded volumes provide fault tolerance against node failures. You can create erasure coded volumes using storage from different nodes such that encoded data fragments are stored on different nodes for redundancy.

For more information see the InfoScale storage administration guides.

Changes related to Replication

The following changes are introduced to replication of Veritas InfoScale 7.2.

Pause and resume file replication jobs

(Linux only)

You can now pause and resume file replication jobs. Use the `vfradmin job pause` command to pause the current file replication job immediately without waiting for the running iteration to complete. It can also be used to pause replication jobs that are waiting for the next schedule. The command does not drop the replication job from the schedule.

Use the `vfradmin job resume` command to resume a paused file replication job. The command replicates the job from where it was paused.

Two new states—`full-sync-paused` and `paused` are introduced to indicate a paused replication job.

Shared extents no longer supported

(Linux only)

Shared extents awareness is no longer supported with file replication. As a result, options and operations related to shared extents are no longer supported.

Replication interval statistics now includes transfer rate

You can now view the transfer rate along with the list of files changed, file data synchronized, errors, and various time stamps for the most recent replication interval statistics.

Pattern lists in consistency groups

(Linux only)

You can now add pattern lists to consistency group definitions in addition to paths. Pattern lists are optional and can be used with both include and exclude lists. All the files and directories matching the specified pattern are replicated.

If an include or exclude path and an include or exclude pattern conflict, the include or exclude path will take precedence. For example, if the exclude pattern is `/mnt1/dir1/dir2` and the include pattern is `/mnt1/dir1/dir2/dir3/*.txt`, none of the files in `dir3` are replicated.

Patterns are formed using the following symbols:

*	Match zero or more characters, except /
?	Match exactly one character, except /
[]	Match any one of the characters specified within the brackets. Character ranges can be specified within the brackets using a hyphen character.
**	Match all files and zero or more directories or subdirectories

For more information on using pattern lists in consistency groups, see the *Veritas InfoScale Replication Administrator's Guide*.

Changes related to Operating System

The following changes are related to Operating System (OS).

AIX 6.1 and Solaris 10 are not supported.