Veritas InfoScale™ 7.3.1 What's New In This Release - AIX, Linux, Solaris

VERITAS™

Last updated: 2017-11-04

Legal Notice

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 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

This document covers the major new branding, licensing, and installation changes that are introduced in 7.3.1.

Changes related to installation and upgrades

The following changes are introduced to the installation and upgrading of Veritas Infoscale 7.3.1.

Change in upgrade path

You can upgrade to Veritas Infoscale 7.3.1 only if the base version of your currently installed product is 6.1 or later.

Changes related to Cluster Server agents

The following sections describe the changes that are introduced in the Cluster Server agents.

New High Availability agent for Amazon Web Services (AWS)

(Linux only)

Veritas has introduced the EBSVol agent for AWS.

This agent is bundled with the product.

AWS EBSVol agent

The EBSVol agent monitors the Amazon EBS volumes. It also attaches or detaches the Amazon EBS volumes to and from the EC2 instances respectively. This agent uses AWS CLI commands to determine the state of the Amazon EBS volumes.

The EBS volume resource does not depend on any other resources.

For more information, see *Cluster Server Bundled Agents Reference Guide - Linux*.

New High Availability agents for Azure environment

(Linux only)

Veritas has introduced the following high availability agents for Azure:

- AzureDisk agent
- AzureIP agent
- AzureDNSZone agent
- AzureAuth agent

These agents are bundled with the product.

AzureDisk agent

The AzureDisk agent brings online, takes offline, and monitors the managed and un-managed Azure data disks. It attaches the managed and un-managed data disks to a virtual machine of the same or different resource group. The AzureDisk agent uses Azure Python SDK to determine whether the Azure data disks are attached to the Azure virtual machines or not.

AzureIP agent

The AzureIP agent manages the networking resources in an Azure environment. The agent uses Azure Python APIs to associate IP resources in an Azure VM.

The agent does the following:

- Gets the NIC details, creates the configuration, and associates or disassociates the private IP address.
- Associates or disassociates the public IP address, with the private IP address.
- Manages the route table entries of overlay IP, for failing over across the subnets.

AzureDNSZone agent

The AzureDNSZone agent monitors and updates the host name to resource record mapping. This agent does the mapping for the Azure DNS domain when failing over to the nodes across the subnets or the regions.

AzureDNSZone agent provides DNS-based traffic routing and failover.

Use this agent, if the resource records need to be dynamically added and deleted from the domain during failover. This agent updates the new resource record mappings while failing over, and allows the clients to connect to the failed over instance of the application.

AzureAuth agent

AzureAuth agent authenticates the Azure subscription using the Service Principal credentials. This agent is a persistent resource that monitors the validity of Service Principal credentials.

For more information, see *Cluster Server Bundled Agents Reference Guide - Linux*.

Support for configuring applications for HA in Azure cloud using InfoScale Enterprise

(Linux only)

InfoScale Enterprise lets you configure applications for HA and disaster recovery (DR) in Azure cloud environment.

The following scenarios are supported:

- Failover within the same subnet using a private IP
- Failover across the subnets using an overlay IP
- Public access to the cluster nodes in an Azure environment using a public IP
- DR across Azure Regions or VNets, and from the on-premises cluster to Azure

For more information, see the following documents:

- *Cluster Server Bundled Agents Reference Guide - Linux*
- *Cluster Server Administrator's Guide - Linux*

- *Veritas InfoScale Disaster Recovery Implementation Guide - Linux*

Support for configuring LLT over TCP

(Linux only)

VCS now provides the option of using LLT over the Transmission Control Protocol (TCP) layer for clusters using wide-area networks and routers. TCP makes LLT packets routable and thus able to span longer distances more economically.

Use LLT over TCP in the following scenarios:

- When using CVM or FSS on AWS for improved IO Shipping performance
- When hardware, such as blade servers, do not support LLT over Ethernet

Also, when you configure LLT over TCP layer for clusters using wide-area networks and routers, you can use the LLT_TCP_CONNS parameter to control the number of TCP connections that can be established between peer nodes.

The default value of this parameter is 8 connections. However, you can extend this value to maximum 64 connections.

For more information, see *Cluster Server Configuration and Upgrade Guide – Linux*.

Start-only option for applications

(Linux only)

The Application agent now provides an option to start or stop an application without monitoring it continuously. Use the **StartOnly** attribute of this agent to indicate that VCS should only start or stop the application and not perform a monitor operation. When **StartOnly** is set, the agent uses the values of the **StartProgram** and the **StopProgram** attributes to determine whether to report the resource as online or offline.

For more information, see the *Cluster Server Bundled Agents Reference Guide - Linux*.

256-bit encryption for enhanced security

The Cluster Server component now provides 256-bit encryption for enhanced security. The `vcsencrypt` utility, which you use to generated encrypted VCS and agent passwords, now generates 256-bit encrypted passwords.

For more information, see the *Cluster Server Administrator's Guide*.

Stale key detection to avoid a false preexisting split brain condition

(Linux only)

It may happen that while the cluster nodes go offline, the network connectivity between the nodes and the coordination points is lost. If such an event occurs with a server-based fencing configuration, some stale keys may be left on one or more coordination point servers (CP servers). The presence of a stale key might create a preexisting split brain condition, after which manual intervention is required to bring the cluster online. The CP server functionality is enhanced to identify the false preexisting split brain condition that occurs due to the presence of such stale keys. The CP server informs the `vxfen` module about the stale keys, and the module performs its arbitration to continue the smooth operation of the cluster.

VCS stop timeout

An environment variable, VCS_STOP_TIMEOUT, is introduced, which prevents the VCS stop operation from being hung during system shutdown or restart.

Default value (seconds): 0

New attributes for Cluster Server agents

The following new attributes are introduced in this release:

(Linux only) Application agent

Attribute	Description
Name: **StartOnly** Type: boolean Dimension: scalar	Indicates whether the application must be monitored or not. If this attribute is set, the agent does not execute the script specified in the **MonitorProgram** attribute, but performs the following actions instead: ■ If the online function is executed and the return code of the **StartProgram** attribute is 0, it reports the resource as online. Otherwise, it reports the resource as offline. ■ If the offline function is executed and the return code of the **StopProgram** attribute is 0, it reports the resource as offline. Otherwise, it reports the resource as online. **Note:** If this attribute is set: You must set the **Critical** attribute to `0` so that VCS does not attempt to fail over or take any action if the application faults. You may increase the values of the **MonitorInterval** and the **OfflineMonitorInterval** attributes, because they do not have an impact. Default: `0` Example: `1`

(AIX and Solaris only) DiskGroup agent

Attribute	Description
Name: **DGOptions** Type: string Dimension: scalar	Specifies the options for the `vxdg import` command. The agent uses this attribute only while bringing a DiskGroup resource online. For more information, see the *vxdg(1m)* manual page. Example: `-o noautostart -o updateid`

Changes related to Veritas File System

The following changes are introduced to Veritas File System (VxFS) of Veritas InfoScale 7.3.1.

New option [-i] included in `fsadm` command to exclude the actively used files during file system reorganization

A new option [-i] is added to the `fsadm` command, to address the reorganization of actively used files.

The updated `fsadm` command is as follows:

```
fsadm [-t vxfs] [-e] [-d] [-E] [-D] [-H] [-i] [-r rawdev] mount_point
```

The [-i] option indicates that the files that are accessed in last 60 seconds must not be reorganized. During the file system reorganization, if an application is actively accessing a particular file or has accessed it in last 60 seconds, then the file system reorganization does not affect the read-write operation on that file and the file will not be reorganized.

Delayed allocation support extended to clustered file systems

The delayed allocation capability for extending writes on a file system was available for local mounts. This capability is now extended for clustered file systems (CFS mounts).

With this support, depending on an application I/O size, instead of allocating a single block for every write operation a clustered file system will now allocate multiple blocks in a single instance. The delayed allocation thus reduces the file system fragmentation.

When an application I/O is received, the delayed allocation capability enables the clustered file system to spilt the write operation in to the following sequence:

1. Reserve a disk space

   When an application I/O is received, the file system first reserves a disk space and the data is cached.

2. Allocate extents

   After a disk space is reserved, a scheduler allocates disk blocks at the background and the file system then combines multiple block allocation requests to allocate extents.

The delayed extent allocation thus helps to avoid file system fragmentation and keeps the extent contiguous even if several files grow at the same time.

The delayed allocation is not dependent on the file system disk layout version and is disabled by default. You can enable delayed allocation using the `vxtunefs` command. You can display the delayed allocation range in the file by using the `fsmap` command.

See the vxtunefs(1M) and fsmap(1M) manual pages.

Notes:

- Delayed allocation must be disabled in cases where the data must be immediately written to the disk. For example, direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O.
- Delayed allocation is not supported on memory-mapped files, BSD quotas, and shared mount points in a Cluster File System (CFS).
- When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.

Support for migrating Oracle database from Oracle ASM to Veritas File System (VxFS)

(Linux and Solaris only) Veritas InfoScale supports real-time migration of:

- (Linux only) Standalone and RAC Oracle databases, hosted on Oracle ASM disks to VxFS systems.
- (Solaris only) Standalone Oracle ASM disks to VxFS system and is mounted on the VxVM disks.

The migration requires a source system on which the primary database is hosted on Oracle ASM disks, and target serves as standby during the migration. That target contains VxVM disks on which the Veritas File System is mounted. The target disks can be on the same host as the primary database or on a different host. In Linux environment, you can migrate multiple instances of database at a time in a RAC environment.

For more information, see *Veritas InfoScale Solution's Guide*.

Changes related to Veritas Volume Manager

The following changes are introduced to Veritas Volume Manager (VxVM) of Veritas InfoScale 7.3.1.

Support for InfoScale deployments in Azure cloud

(Linux only)

Veritas InfoScale can now be deployed in Microsoft Azure cloud to avail the following capabilities:

- Intelligent data movement to the cloud using SmartMove technology

- Storage reliability and performance using Flexible Storage Sharing (FSS) technology, where FSS leverages Azure's block storage to provide shared storage capability.
  FSS in cloud environment is supported only with LLT over UDP.
  For more information, see the *Storage Foundation Cluster File System High Availability Administrator's Guide*.
- High availability of data and disaster recovery using Volume Replicator (VVR) and CVR technologies
  See "Support for configuring volume replication in Azure cloud" on page 14.

NVMe device support for Solaris

For Solaris only

A single generic ASL is supported for all the NVMe devices. But, we do not guarantee full support to use generic ASL, to fetch all the properties of all NVMe devices. Currently only Samsung NVMe is tested for use with this generic ASL.

Following command output shows SAMSUNG NVMe device for InfoScale use:

```
# vxdisk -p list n04-t7-936090_nvme0_0

DISK                 : n04-t7-936090_nvme0_0
VID                  : NVMe
UDID                 : NVMe%5F144D108E%5FS2T7NAAH404821%5F3E17006194382500
SCSI_VERSION         : 0
PID                  : 144D108E
PHYS_CTLR_NAME       : /pci@304/pci@1/nvme@0
PGR_CAPABLE          : N
MEDIA_TYPE           : ssd
LUN_TYPE             : std
LUN_SNO_ORDER        : 0
LUN_SERIAL_NO        : 3E17006194382500
LIBNAME              : libvxnvme_sol.so
DMP_DEVICE           : n04-t7-936090_nvme0_0
DDL_DEVICE_ATTR      : ssd
CONN_TYPE            : local
CAB_SERIAL_NO        : S2T7NAAH404821
ATYPE                : A/A
ANAME                : NVMe
TRANSPORT            : SCSI
ENCLOSURE_NAME       : n04-t7-936090_nvme0
DMP_SINGLE_PATH      : /dev/rdsk/c4t1d0
LUN_SIZE             : 781404246
```

```
NUM_PATHS            : 1
STATE                : online
DISK_TYPE            : auto
FORMAT               : cdsdisk
DA_INFO              : format=cdsdisk,privoffset=208,pubslice=2,privslice=2
PRIV_OFF             : 208
PRIV_LEN             : 65536
PUB_OFF              : 65744
PUB_LEN              : 6251168128
PRIV_UDID            : NVMe%5F144D108E%5FS2T7NAAH404821%5F3E17006194382500
DISKID               : 1505772756.15.N04-T7-936090.engba.veritas.com
DISK_TIMESTAMP       : Day Month Date Time (AM/PM) Year
```

Support for Azure Blob storage connector

(Linux only)

Veritas InfoScale now supports the use of Blob storage connectors with VxVM. The Blob storage connector enables you to use cloud storage as a tier to manage your storage needs with agility and flexibility. You can build a hybrid storage environment that seamlessly integrates local on-premise storage with cloud storage. With the added support for Blob connector, Veritas InfoScale supports the ability to migrate data from on-premise storage to cloud storage.

For more information, see *Storage Foundation Cluster File System High Availability Administrator's Guide*.

Erasure coded volume enhancements

(Linux only)

Erasure coding feature supports in standalone and shared environments. Erasure coded volume can be used in the following use cases:

- Erasure Coded (EC) volume as a backing store for generic workloads such as transactional or FileStore.
- Erasure Coded (EC) volume is tuned for optimal performance for special workloads such as object store.

The following enhancements have been added to erasure coded volumes.

- Customized failure domain
  To use customized failure domain user can specify the constraints on column, and define the failure domain as per the need and available configuration. Customizing is done by setting tags on the devices which are used for allocation.
- Using Stripe Group and Stripe Confined Group during volume creation.

When you create an EC volume, specify the Stripe Group and Stripe Confined Group, and if nothing is specified the Stripe Group value is the host.

For more information, see *Storage Foundation Cluster File System High Availability Administrator's Guide*.

Changes related to replication

The following changes are introduced to replication of Veritas InfoScale 7.3.1.

Support for configuring volume replication in Azure cloud

(Linux only)

Veritas InfoScale lets you to configure replication between two data centers to support various HA and DR scenarios for applications in the Azure cloud.

InfoScale supports the following scenarios for setting up replication in an Azure cloud:

- On-premise to cloud
- Across user-defined sites in the same region (includes both, within a single VNet and across VNet scenarios)
- Across regions
- Across multiple sites and regions (Campus cluster)

You can use any of the following Veritas' technologies to set up replication:

- Use Volume Replicator (VVR) to configure standalone replication between the attached Azure block storage devices.
- Use Cluster Volume Replication (CVR) to configure replication between multi-node clusters.
- Use site tagging to use FFS to create mirrored volumes across user-defined sites. FSS enables you to create shared-nothing clusters by sharing Azure block storage over the network.

For more information, see the following documents

- *Storage Foundation Cluster File System High Availability 7.3.1 Administrator's Guide - Linux*
- *Veritas InfoScale 7.3.1 Disaster Recovery Implementation Guide - Linux*

Support for setting up replication across clouds

Veritas replication technology can now be used to set up replication across AWS and Azure cloud.

With Veritas Volume Replicator (VVR), you can leverage one of the cloud setup as an on-premise data center and the other cloud as a DR site to replicate data across the two clouds.

For details about setting up replication across cloud, refer to *Veritas InfoScale Disaster Recovery Implementation Guide - Linux*.

Veritas Volume Replicator Performance Improvements

Previously TCP was one of the replication protocols used, where one data structure was used to transport the data. This data structure had eight connections which used the bandwidth. These eight connections can be increased and tuned to 16 or 32 to use more bandwidth. However, with one data structure the bandwidth utilization would have very little effect on latency between the Primary host and Secondary host.

Now in this 7.3.1 release, to improve the bandwidth utilization, TCP connections are tuned to multiple data structures to achieve better performance. Every data structure has one transport connection which is set as default, and available all the time. Currently there are 16 data structures available and enabled which are intended to use in parallel. Hence all the data is transited in parallel through these data structure connections, and distributes the network load. If more bandwidth is required, you can tune the existing data structure in the multiple of 16. See*Veritas InfoScale™ Replication Administrator's Guide* for more information.

Support for replication of encrypted volumes

To enable the volume encryption for RG volumes, you need to perform few settings on primary and secondary volumes. To facilitate the process execute the `vxsetupencryption` utility for each disk group which is to be replicated. The `vxsetupencryption` utility sets consistent encryption keys on primary as well secondary volumes.

Changes related to operating systems

The following changes are related to the support for operating systems (OSs).

Support for Root Disk Encapsulation (RDE) on All Linux Distributions is Deprecated

Veritas InfoScale no longer supports Root Disk Encapsulation (RDE) on any of the Linux Distributions, from 7.3.1 release onwards.

- If you have already deployed InfoScale product in a Linux environment, and now plan to upgrade the product version to InfoScale 7.3.1. You must first un-encapsulate the root disk and then upgrade the InfoScale Product. To remove the encapsulation of root-disk from the system, you need to unencapsulate the root disk.

  To unencapsulate the root disk see, *Storage Foundation Administrators Guide-Linux* and *Storage Foundation Cluster File System High Availability Administrators Guide*

Changes related to Dynamic Multipathing

The following changes are introduced to Dynamic Multipathing 7.3.1 of Veritas InfoScale7.3.1.

Support for dynamic multipathing in the KVM guest virtualized machine

For Linux only

DMP in the KVM guest virtualized machine provides:

- SCSI-3 PR I/O fencing support for Linux KVM hosts and guests with Layered DMP configuration in KVM environment.
- From 7.3.1 release onwards, Linux KVM host and their guests can use SCSI-3 PR fencing with virtual devices backed by DMP devices in KVM host.