Veritas InfoScale 7.2 Installation and Upgrade Guide - Windows

VERITAS™

Veritas InfoScale Installation and Upgrade Guide

Document version: 7.2 Rev 0

Last updated: 2018-09-04

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo 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 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

Contents

Chapter 1 Preinstallation and planning 7

About the Veritas InfoScale product suite 7
About the Dynamic Multi-pathing for VMware component 11
Supported hardware and software 11
Disk space requirements 12
Installation requirements 12
Requirements for installing InfoScale Storage in a Microsoft Failover Cluster 15
Recommendations and best practices 16
Verifying the system configuration using the Windows Data Collector 16
About InfoScale licenses 17
Licensing notes 18
vxlicrep command 20

Chapter 2 Installing the Veritas InfoScale products 21

About installing the InfoScale products 21
About the co-existence of InfoScale products 23
Installing the server components using the installation wizard 24
Applying the selected installation and product options to multiple systems 30
Installing the server components using the command-line installer 30
Parameters for Setup.exe 32
Available product options and supported DMP DSMs 34
Registering the InfoScale Storage resource DLLs 37
Installing the client components 38

Chapter 3 Upgrading to InfoScale products 39

Preparing the systems for an upgrade 39
About the supported upgrade paths and the supported minimum product versions 39
General preparations 41
Recommendations and considerations for product upgrade 44

Performing the product upgrade 45
Upgrading SFW or SFW Basic in a non-clustered environment 46
Upgrading SFW or SFW Basic in a Windows Server Failover Cluster environment 50
Upgrading VCS 55
Upgrading SFW HA 58
Upgrading DMP 63
About transitioning between the InfoScale products 64

Chapter 4 Performing the post upgrade tasks 66

Deployment scenarios and applicable post upgrade tasks 66
Re-enabling Volume Replicator in a non-clustered environment 69
Re-enabling Volume Replicator in a Microsoft failover cluster environment 70
Reconnecting DMP DSM paths after the upgrade 71
Reconfiguring the Veritas InfoScale Messaging Service 71
Importing the configured rules 72
Upgrading the dynamic disk group version 72
Upgrading clusters for stronger security 74
Reinstalling the custom agents 75
Including custom resources 75

Chapter 5 Administering the InfoScale product installation 77

Adding or removing product options 77
Managing InfoScale licenses 79
Repairing an InfoScale product installation 82
About reinstalling InfoScale products 83

Chapter 6 Uninstalling the InfoScale products 84

About uninstalling the InfoScale products 84
Uninstalling the InfoScale products using the installation wizard 84

Chapter 7 Performing application upgrades in an InfoScale environment 87

Upgrading Microsoft SQL Server 87
Upgrading to later versions of SQL Server 88
Upgrading Oracle 90
Performing the post upgrade tasks 91

Associating the updated Oracle database with the listener 91
Configuring the Oracle database and listener to use the virtual IP address 92
Configuring Oracle and listener services 95
Modifying the ServiceName attribute for the netlsnr resource 96
Upgrading application service packs in an InfoScale environment 96
Upgrading the Exchange Server service packs 97
Upgrading the SQL Server service packs 98
Upgrading SharePoint Server 2010 to a service pack 104

Appendix A Services and ports 105

InfoScale ports and services 105

Appendix B Migrating from a third-party multi-pathing solution to DMP 108

Migrating from EMC PowerPath 108
Migrating from Hitachi Dynamic Link Manager (HDLM) 109
Uninstalling HDLM in a non-clustered environment 109
Uninstalling HDLM in a clustered (MSCS or VCS) environment 110
Configuring DMP for Active/Active load balancing in a cluster 111

Chapter 1

Preinstallation and planning

This chapter includes the following topics:

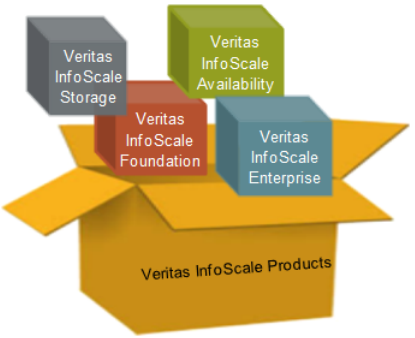
- About the Veritas InfoScale product suite
- Supported hardware and software
- Disk space requirements
- Installation requirements
- Requirements for installing InfoScale Storage in a Microsoft Failover Cluster
- Recommendations and best practices
- About InfoScale licenses

About the Veritas InfoScale product suite

Veritas InfoScale products address enterprise IT service continuity needs. They provide resiliency and software defined storage for critical services across your data center infrastructure.

The Veritas InfoScale product suite offers the following products:

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

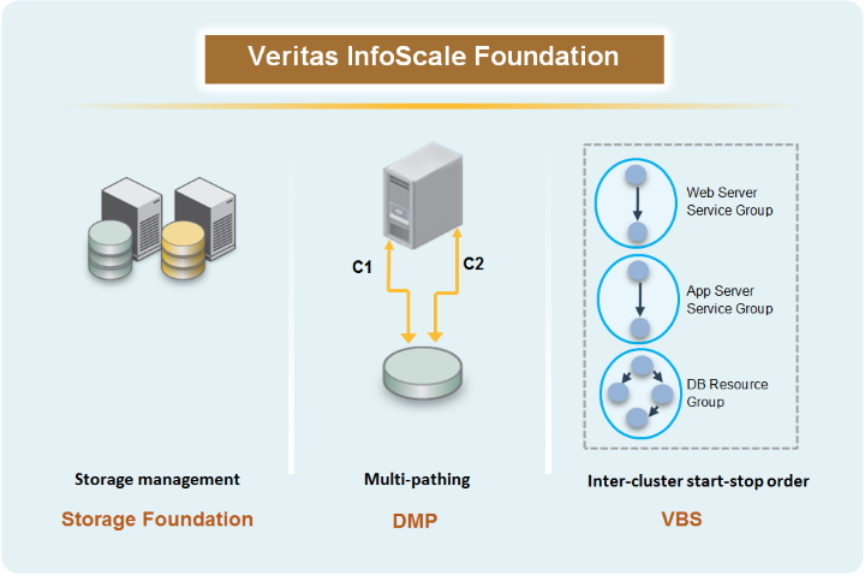


Veritas InfoScale Foundation

Veritas InfoScale Foundation simplifies the management of storage across the data center, with an efficient application-aware storage management solution. This product works across heterogeneous storage and server environments.

The following figure depicts the components that Veritas InfoScale Foundation offers.

Figure 1-1 Veritas InfoScale Foundation components

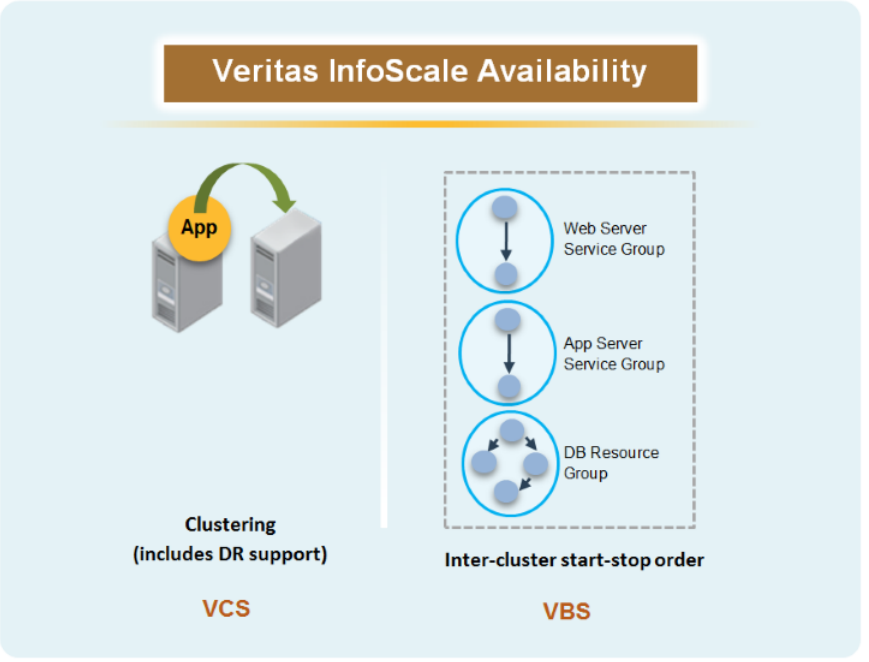


Veritas InfoScale Availability

Veritas InfoScale Availability is a comprehensive high availability and disaster recovery solution that protects critical business services from planned and unplanned downtime. The critical business services include individual databases, custom applications, and complex multi-tiered applications, which may span across physical and virtual environments and over any distance.

The following figure depicts the components that Veritas InfoScale Availability offers.

Figure 1-2 Veritas InfoScale Availability components

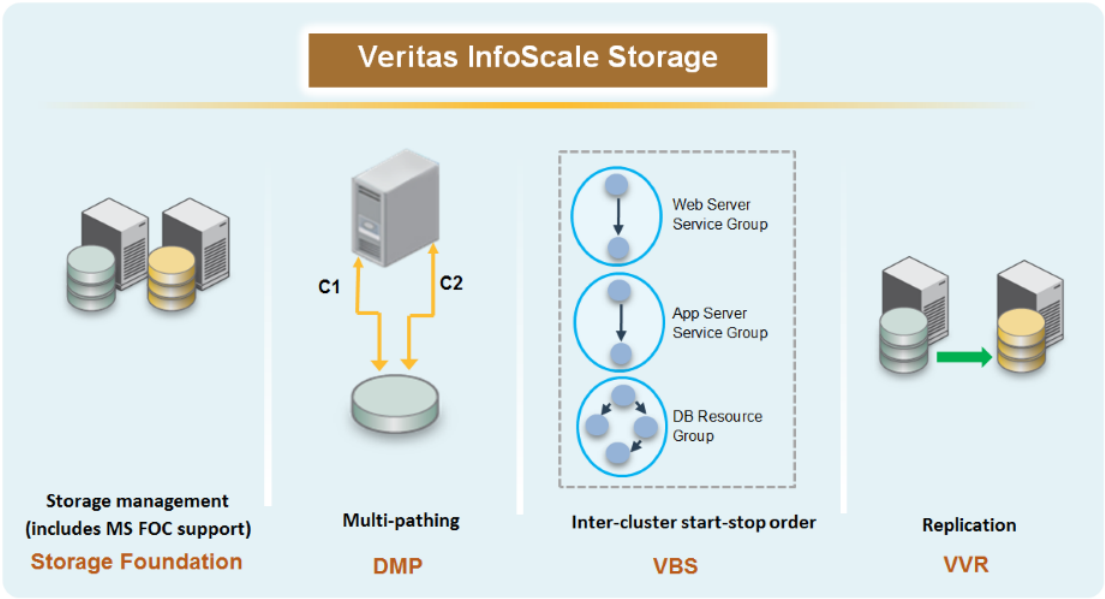


Veritas InfoScale Storage

Veritas InfoScale Storage provides a high-performance storage management solution that maximizes storage efficiency, data availability, operating system agility, and performance. This product works across heterogeneous server and storage environments.

The following figure depicts the components that Veritas InfoScale Storage offers.

Figure 1-3 Veritas InfoScale Storage components

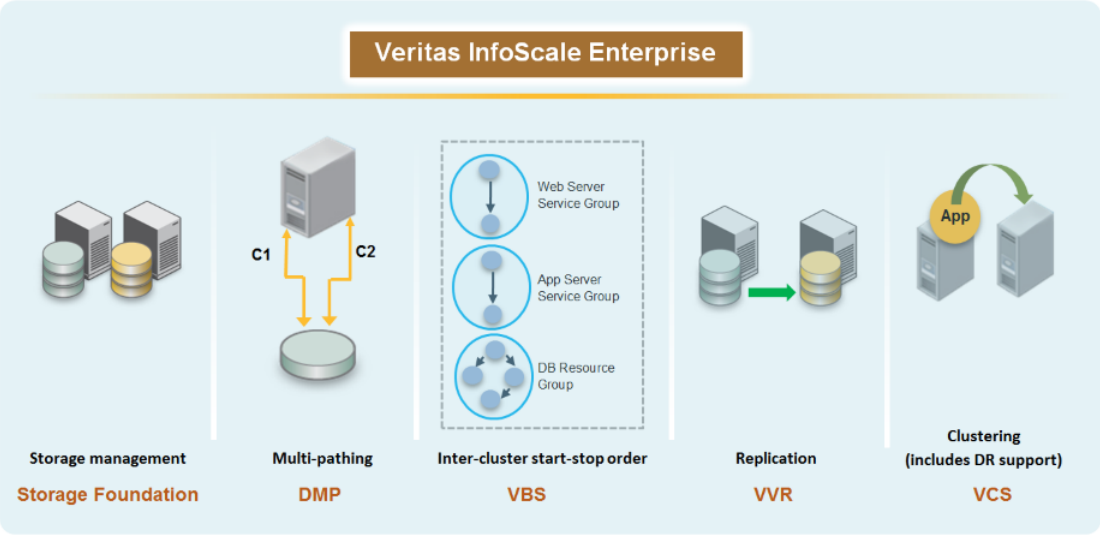


Veritas InfoScale Enterprise

Veritas InfoScale Enterprise provides a powerful combination of comprehensive storage management and application availability. This product helps you to increase performance, flexibility, and efficiency in your data center. With built-in application acceleration, Veritas InfoScale Enterprise lets you optimize data efficiently across heterogeneous storage or server environments and recover applications instantly from downtime. This product delivers unmatched performance and protection for business-critical applications across physical, virtual, or cloud deployments.

The following figure depicts the components that Veritas InfoScale Enterprise offers.

Figure 1-4 Veritas InfoScale Enterprise components



About the Dynamic Multi-pathing for VMware component

Dynamic Multi-Pathing for VMware (VxDMP) is a multi-pathing solution integrated with VMware vSphere Client infrastructure.

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*

Supported hardware and software

For the latest information on the supported hardware and software, see the compatibility list on the SORT Web site. In addition to the compatibility list, a Late Breaking News (LBN) is available for the latest updates, patches, and software issues regarding this release:

https://sort.veritas.com/documents

Disk space requirements

The following table summarizes the approximate disk space requirements for the InfoScale products.

Product	Required disk space
InfoScale Foundation	1200 MB
InfoScale Availability	832 MB
InfoScale Storage	1500 MB
InfoScale Enterprise	1589 MB

Installation requirements

Before you install an InfoScale product, ensure that all its installation requirements are met. The following table lists the requirements for each InfoScale product.

Table 1-1 Installation prerequisites

Pre-requisites	Veritas InfoScale Foundation	Veritas InfoScale Availability	Veritas InfoScale Storage	Veritas InfoScale Enterprise
Patches and hotfixes				
Ensure that the latest Microsoft Windows updates are installed on the systems. The latest updates are required for successfully installing Microsoft Visual C++ 2015 Redistributable-x64. The product installation fails if Microsoft Visual C++ 2015 Redistributable-x64 is not installed.	✓	X	✓	✓
Firewall and port settings				
Spyware monitoring and removal software is disabled	✓	✓	✓	✓
The ports and services that are used during installation for both, inbound and outbound communication are enabled For a list of required ports and services, See “InfoScale ports and services” on page 105.	✓	✓	✓	✓

Table 1-1 Installation prerequisites *(continued)*

Pre-requisites	Veritas InfoScale Foundation	Veritas InfoScale Availability	Veritas InfoScale Storage	Veritas InfoScale Enterprise
System requirements				
Three network adapters are available	X	✓	X	✓
Minimum one IP address that is not assigned by Dynamic Host Configuration Protocol (DHCP) is available	X	X	✓	✓
At least two IO paths from the server to the storage array for load balancing	✓	X	✓	✓
General requirements				
Microsoft Failover Cluster is configured (Applicable only if you plan to configure a Microsoft Failover Cluster)	X	X	✓	X
Microsoft .NET Framework 4.5 is installed (This requirement applies to both, server as well as client-only components.)	✓	✓	✓	✓
Computer Browser Service is enabled	✓	✓	✓	✓
Microsoft Windows Service (WMI) is activated	✓	✓	✓	✓
No parallel installations, live updates, or Microsoft Windows updates are in progress	✓	✓	✓	✓
Microsoft Visual C++ 2010 SP1 (x64) and the Microsoft Visual C++ 2010 SP1 (x86) re-distributable package is installed on the systems where you would want to launch the client consoles (VEA, Java Console, SCC etc)	✓	✓	✓	✓
Permission requirements				
The user account is included as a domain user and has local administrators privileges on all the systems	✓	✓	✓	✓
The user account has write permissions for the Active Directory objects corresponding to all the systems	✓	✓	✓	✓
The user account has administrative privileges on all the systems to load and unload device drivers	✓	✓	✓	✓
Network requirements				
All systems are in the same domain	✓	✓	✓	✓

Table 1-1 Installation prerequisites *(continued)*

Pre-requisites	Veritas InfoScale Foundation	Veritas InfoScale Availability	Veritas InfoScale Storage	Veritas InfoScale Enterprise
Remote systems are accessible over the network and the user account has local administrative privileges to all the systems	✓	✓	✓	✓
Trust relationship is set up for systems in different domains (Applicable only if you plan to configure a disaster recovery setup in which the systems at the primary site and the secondary site are in different domains)	X	✓	X	✓
For IPv4 network, the required IP addresses are available	✓	✓	✓	✓
For IPv6 network, the IP address configuration is "stateless automatic", the IP address format is "Global unicast" or "Unique localcast", and the prefix is advertised	✓	✓	✓	✓
The system is a part of Windows Active Directory domain	✓	✓	✓	✓
DNS Service is available	✓	✓	✓	✓
DMP DSM requirements				
The host has an HBA (host bus adapter) port for each path to the SAN switch	✓	X	✓	✓
The host has one SCSI or fiber cable per host bus adapter port	✓	X	✓	✓
In case of iSCSI, each host bus adapter port has a unique SCSI ID	✓	X	✓	✓
Only one path is connected during product installation	✓	X	✓	✓
The Windows Storport driver is installed	✓	X	✓	✓
Correct hardware drivers for the DMP DSMs are installed	✓	X	✓	✓
The MPIO server feature is enabled	✓	X	✓	✓
No other third-party DSMs are installed for the same array	✓	X	✓	✓

Requirements for installing InfoScale Storage in a Microsoft Failover Cluster

The following table lists the InfoScale requirements for installing InfoScale Storage, in an active Microsoft Failover Cluster.

These requirements are applicable in addition to the InfoScale product-specific installation pre-requisites, and must be satisfied on all the Microsoft Failover Cluster nodes, where you plan to install InfoScale Storage.

Table 1-2 InfoScale requirements for Microsoft Failover Cluster

Pre-requisites	
✓	One CD-ROM drive is accessible to the node on which you are installing InfoScale Storage.
✓	The node has the required disk space available.
✓	The required storage hardware is configured to access the shared storage. SCSI or Fibre channel host bus adapters (HBAs) can be used.
✓	Three network adapters (two NICs exclusively for the private network and one for the public network) are available on each cluster node, and each private NIC is routed through a separate hub or switch to avoid single points of failure. To prevent lost heartbeats on the private networks, and to prevent the Microsoft cluster from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto negotiation options on the private network adapters.
✓	Static IP addresses are used for the public network and private network cards, and the DNS name resolution is configured for each node. This is required if you plan to configure replication using VVR.
✓	DNS and Active Directory Services are available and a reverse lookup zone exists in the DNS. Refer to the Microsoft documentation for instructions on creating a reverse lookup zone.
✓	Each cluster node is in the same Windows Server domain, and uses the same operating system version.
✓	The user account has administrative privileges on the cluster nodes.

Recommendations and best practices

- Do not install Veritas InfoScale Availability and Veritas InfoScale Enterprise on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- In case of IPv6, the following IP address formats are not supported:
  - Multicast
  - Anycast
  - Link local
  - Site local
- In case of IPv6, the Mixed mode configuration with "stateful" and "stateless" configurations are not allowed.
- Do not change the cable connection order after installing SFW.
  For example, if host bus adapter A is connected to port A on the array and host bus adapter B is connected to port B on the array, do not swap the connections between ports on the array (A to B and B to A) after installing the product.
- Verify the system configuration to gather information about the systems on which you plan to install the product.
  See "Verifying the system configuration using the Windows Data Collector" on page 16.

Verifying the system configuration using the Windows Data Collector

The Windows Data Collector enables you to gather information about the systems in your network. It thus helps you verify your system configuration before you begin with the product installation.

Installing the Windows Data Collector

You can download the Data Collector from the Veritas Services and Operations Readiness Tools (SORT) website at:

https://sort.veritas.com

Under the **SORT** tab, select **My SORT** and then on the Custom Reports widget, follow the instructions to download the Data Collector.

Running the verification reports

The Data Collector uses the gathered information to generate the reports that enable you to perform the following:

- Determine whether a system is ready for product install or upgrade.
- Analyze the configuration of your current Veritas products and make recommendations about availability, use, performance, and best practices.
- Get detailed information about your installed Veritas products, versions, and licenses.

The report contains a list of passed and failed checks and details about each of them. After the Data Collector completes the check, you can save a summary report as an HTML file and an XML file.

For more details on running a verification report, refer to the platform-specific README file located on the Custom Reports widget on the SORT Web site.

Ensure that the following Windows services are running on all the systems where you wish to run the Windows Data Collector report:

- Windows Management Instrumentation (WMI)
- IP Helper
- Remote Registry

About InfoScale licenses

A product-specific license is available for individual InfoScale product. You can procure the license from Veritas license certificate and portal.

https://my.veritas.com/

During installation, the product installer provides the following options to specify the license details.

- Keyless
- User Entered Key

You can use the keyless license for 60 days. If you install the product using the keyless option, a message is logged every day in the Event Viewer indicating that you must perform any one of the following tasks, within 60 days of product installation. Failing this, a non-compliance error is logged every four hours.

- Add the system as a managed host to a Veritas InfoScale Operations Manager Management Server.
  For more details, refer to the Veritas InfoScale Operations Manager documentation.
- Add an appropriate and valid license key on this system using the product installer from Windows Add or Remove Programs.

In case of an User Entered Key license, you must enter the license key that is procured from the Veritas license certificate and portal.

The product installer enables you to switch from a Keyless license to a User Entered Key license and vice-a-versa.

For more details about product co-existence and licensing, see:

See "About the co-existence of InfoScale products" on page 23.

Licensing notes

Review the following licensing notes before you install or upgrade the product.

- If you are installing the product for the first time, the Keyless option is selected by default.
- If you are upgrading the product, the wizard retrieves all the license keys that were installed earlier. You can continue to use the 6.0.x or 6.x license keys. However, note that the features available with the 6.0.x or 6.x license keys will only be enabled after the product upgrade. To enable all the available features, specify a 7.2 license key. Feature-based license keys are not available for 7.2.
  For example,
  If you upgrade from SFW Standard to InfoScale Storage and continue to use the SFW Standard license keys, the following features that are available with InfoScale Storage will not be enabled:
  - Snapshots
  - Support for Microsoft Failover Cluster
  - Cluster Volume Manager for Microsoft Failover Cluster
  - FastFailover
  - Fast file re-sync
  - Fast mirror re-sync
  - Cluster quorum
- During the product upgrade, you can choose to change the license option to Keyless. After changing the license option to Keyless, you can install all the available product options.
- You cannot install the 7.2 license key over the 6.0.x or the 6.x license key and vice-a-versa.
- If you use a keyless license option, you must configure Veritas InfoScale Operations Manager within two months of product installation. Failing this, a warning message for non-compliance is displayed periodically.

For more details on configuring Veritas InfoScale Operations Manager, refer to Veritas InfoScale Operations Manager product documentation.

- You can install new licenses or remove the existing licenses using the product installer.
  In case of InfoScale Foundation, Veritas InfoScale Storage, and InfoScale Enterprise if you remove the licenses, the vxsvc service fails to start and the service recovery options are changed to “Take No Action”. To start the service you must enter the license and then manually start the service or change the service recovery option to “Restart the Service”.
- You can use the license keys assigned for higher Stock Keeping Units (SKU) to install the lower SKUs.
  For example, if you have procured a license key that is assigned for InfoScale Enterprise, you can use the license key for installing any of the following products:
  - InfoScale Foundation
  - Veritas InfoScale Storage
  - InfoScale Availability

  The following table provides details about the license SKUs and the corresponding products that can be installed:

License SKU procured	**Products that can be installed**			
	InfoScale Foundation	Veritas InfoScale Storage	InfoScale Availability	InfoScale Enterprise
InfoScale Foundation	✓	X	X	X
Veritas InfoScale Storage	✓	✓	X	X
InfoScale Availability	X	X	✓	X
InfoScale Enterprise	✓	✓	✓	✓

Note: At any given point in time you can install only one product.

vxlicrep command

The `vxlicrep` command generates a report of the licenses in use on your system.

To use the vxlicrep command to display a license report

1. Open a command prompt window and navigate to the following path:

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

2. Enter the `vxlicrep` command without any options to generate a default report.
3. Enter the `vxlicrep` command with any of the following options to produce the type of report required:

-g	default report
-s	short report
-e	enhanced/detailed report
-l	print report for valid keys only
-k <key1, key2, ----	print report for input keys key1, key2, ----
-v	print version
-h	display this help

Chapter 2

Installing the Veritas InfoScale products

This chapter includes the following topics:

- About installing the InfoScale products
- About the co-existence of InfoScale products
- Installing the server components using the installation wizard
- Applying the selected installation and product options to multiple systems
- Installing the server components using the command-line installer
- Parameters for Setup.exe
- Available product options and supported DMP DSMs
- Registering the InfoScale Storage resource DLLs
- Installing the client components

About installing the InfoScale products

The InfoScale products have server and client components. When you install the server components, the following options are installed by default:

- Client components
  Includes the Veritas Enterprise Administrator (VEA) console and the Solutions Configuration Center (SCC).
  Note that the Cluster Manager (Java Console) is no longer packaged and installed with the product. You can download the Java Console at: https://www.veritas.com/content/trial/en/us/vcs-utilities.html

- Application and database agents
  These agents are used in a VCS cluster environment.
- VRTSvbs package
- Veritas InfoScale Operations Manager (Host Component)

You can choose to install the client components separately.

See "Installing the client components" on page 38.

You can install the server components using either the product installation wizard or the command line interface (CLI). The product installation wizard lets you install the product on multiple systems at a time. Using the CLI, you can install the product on a single system at a time.

The following figure depicts the high-level tasks that are involved in installing the server components, using the installation wizard:



For more information about installing the server components using the product installation wizard:

See "Installing the server components using the installation wizard" on page 24.

To install the server components using the CLI:

See "Installing the server components using the command-line installer" on page 30.

To install the client components:

See "Installing the client components" on page 38.

About the co-existence of InfoScale products

You can install an InfoScale product on a system where another InfoScale product is already installed.

The following table provides the supported co-existence scenarios.

Table 2-1 InfoScale products co-existence

Product installed	**Supported co-existence**			
	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
InfoScale Foundation	Not applicable	Supported	Not supported	Not supported
InfoScale Availability	Supported	Not applicable	Supported **Note:** Do not select Microsoft Failover Cluster option during the product installation	Not supported
InfoScale Storage	Not supported	Supported**	Not applicable	Not supported
InfoScale Enterprise	Not supported	Not supported	Not supported	Not applicable

Note: ** If you have configured Microsoft Failover Cluster, you must unconfigure it and remove the SFW component for the Failover Cluster, before installing InfoScale Availability.

To unconfigure Microsoft Failover Cluster, refer to Microsoft documentation.

To remove the SFW component for Failover Cluster, use Windows Add Remove Programs. On the System Selection panel of the product installer, clear the check box for Microsoft Failover Cluster.

Note that the following limitations apply in case of InfoScale Availability and InfoScale Storage co-existence:

- If Keyless licensing type is selected during the product installation, checks performed to monitor the number of days of product installation are based on the SFW component. As a result, if you do not enter a valid license key or do not add the host as a managed host within 60 days of InfoScale Storage installation, a non-compliance error is logged every 4 hrs in the Event Viewer.
- After InfoScale Storage is installed, the VCS configuration wizards discover only SFW dynamic disks for configuring storage resources in an application service group. If you have already configured NetApp or native disk (LDM) resources, you can continue to use them. However, you cannot add or remove any of these resources using the VCS configuration wizards. Then, you must manually modify the service group to add or delete the NetApp or LDM resources.

Installing the server components using the installation wizard

The product installation wizard allows you to install the product on multiple systems at a time.

Before you begin to install the product ensure that you have reviewed the installation prerequisites, licensing, and the product co-existence details.

See “About the Veritas InfoScale product suite” on page 7.

See “About InfoScale licenses” on page 17.

See “About the co-existence of InfoScale products” on page 23.

Note: If you plan to install InfoScale Storage in an active Microsoft Failover Cluster, ensure that you have reviewed the applicable pre-requisites, and use the "rolling-install" procedure to perform the product installation. To use the "rolling-install" procedure, install InfoScale Storage first on the inactive cluster node. Then move the cluster resources to the other node and install the product on the now inactive node.

Perform the following steps to install the server components

1. Download the installation package from the following location:

   https://sort.veritas.com

2. Allow the autorun feature to start the installation or double-click **Setup.exe**.

   The CD browser appears.

   Note: If you install the software using the product software disc, the CD browser displays the installation options for all the products that are specified earlier. However, if you download the installation package from the Veritas website, the CD browser displays the installation options only for the product to be installed.

3 Click the required product-specific tab and then click the link to install the components.

Note: The client components are installed by default along with the server components. However, the client components are not installed if the system is a server core machine.

In addition to the product-specific tabs, the CD browser also provides the following links:

Late Breaking News	Click to access the latest information about updates, patches, and software issues regarding this release.
Windows Data Collector	Click to verify that your configuration meets all pertinent software and hardware requirements.
SORT	Click to access the Veritas Services and Operations Readiness Tools (SORT) site. In addition to the product download you can also download the custom reports about your computer and Veritas enterprise products, a checklist providing configuration recommendations, and system and patch requirements to install or upgrade your software.
Browse Contents	Click to view the software disc contents.
Technical Support	Click to contact Veritas Technical Support.

4 On the Welcome panel, review the list of prerequisites and click **Next**.

5 On the License panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.

The **Participate in the Veritas Product Improvement Program by submitting system and usage information anonymously** check box is selected by default. The Product Improvement Program allows the wizard to collect installation, deployment, and usage data and submit it anonymously to Veritas. The collected information helps identify how customers deploy and use the product. If you do not want to participate in the Product Improvement Program, clear the selection of the check box.

6 On the System Selection panel, select the systems and the desired Installation and Product options:

You can select the systems in one of the following ways:

- In the System Name or IP text box, manually type the system name or its IP address and click **Add**.

  Note: The wizard does not support the Internet Protocol version 6. To add the systems having Internet protocol version 6, you must type the system name.

  The local host is populated by default.

- Alternatively, browse to select the systems.
  The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.

Once you add or select a system, the wizard performs certain validation checks, and notes the details in the Verification Details box. To review the details, select the desired system.

To select the installation and product options, perform the following tasks on each of the selected system.

Note: To apply the selection to multiple systems, select the system for which you have selected the installation and product options and then click **Apply to multiple systems**.

See "Applying the selected installation and product options to multiple systems" on page 30.

- By default the wizard uses %ProgramFiles%\Veritas as the installation directory. To customize the installation directory, click **Browse** and select the desired location. Click **OK**.
  Install the product at the same location on all the systems.
  If you are upgrading the product, the installation directory is selected by default.

  Note: The installation directory must contain only English characters, if:

  - You plan to configure the cluster for single sign-on authentication.

  - Your system runs a non-English locale operating system.

- Select the required license type from the **License key** drop-down list.

Note: The default license type is "Keyless".

If you select the "Keyless" license type, all the available product options are displayed and are selected by default.

If you select "User entered license key" as your license type, the License Details panel appears by default. On the License Details panel, enter the license key and then click **Add**. You can add multiple licenses for the various product options you want to use.

The wizard validates the entered license keys and displays the relevant error if the validation fails. After the validation is complete, click **OK**.

- From the list of product options, select the options to be installed.

  The options differ depending on the product you install.

  For the list of available options and details about the scenarios in which they can be used, refer to:

  See "Available product options and supported DMP DSMs" on page 34.

7 On the System Selection panel, click **Next**.

Note that the wizard fails to proceed with the installation, unless all the selected systems have passed the validation checks and are ready for installation. In case the validation checks have failed on any of the system, review the details and rectify the issue. Before you choose to proceed with the installation, select the system and click **Re-verify** to re-initiate the validation checks for this system.

8 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This selection reboots all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

9 On the Installation panel, review the progress of installation and click **Next** after the installation is complete.

If an installation is not successful on any of the systems, the status screen shows a failed installation.

Note: During the upgrade, the Installation panel displays a list of services and processes running on the systems. Select a system to view the services and processes running on it and review the list.

The wizard stops the product-specific services and discovers the processes running, if any, on the systems. These processes need to be stopped to proceed with the operation. Click **Next** to forcefully stop the processes and proceed with the operation. Alternatively, you can manually stop the processes. If the services or processes cannot be stopped, the operation fails. Rectify the error and then click **Retry** to validate the affected system again. Click **Retry All** to validate all the systems again.

In case you want to proceed with the upgrade without stopping a particular process, contact Veritas Technical Support.

---

10 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed on any of the system, refer to the log file for details. You may have to re-install the software.

11 On the Finish panel, click **Finish**.

If you had chosen to initiate the auto reboot, a confirmation message to reboot the local system appears. Click **Yes** to reboot immediately or **No** to reboot later.

In case you had not selected to initiate the auto reboot, ensure that you manually reboot these systems.

This completes the product installation. Check the SORT website for the applicable patches, agents, or the array-specific modules, if any, to be installed:

https://sort.veritas.com/

You can now proceed to configure the required components. Refer to the component-specific guides for more details about the configuration tasks.

Note: If you have installed InfoScale Storage with Microsoft Failover Cluster, but the cluster is not yet configured, you must register the InfoScale Storage resources, after configuring the Microsoft failover cluster software.

See "Registering the InfoScale Storage resource DLLs" on page 37.

However, if you have installed InfoScale Storage in an active Microsoft Failover Cluster, then you must remove the physical disk resources for all the basic disks. You must do so before configuring the SFW cluster disk groups. Failing this, a reservation conflict occurs.

Applying the selected installation and product options to multiple systems

To apply the selected installation and product options to multiple systems, perform the following steps:

1. Click on any one of the selected systems and select the desired installation and product options.
2. Click **Apply to multiple systems**.
3. On the Apply Installation Options panel, select the installation options to be applied and then select the desired systems. Click **OK**.

Note: The installation directory is selected by default on the systems where the product is being upgraded. The selected **Install Directory** option does not apply to these systems.

Installing the server components using the command-line installer

The command-line installer for the InfoScale products (`Setup.exe`) lets you perform a silent installation. You can perform a silent installation only on one system at a time.

Before you begin to install the product, ensure that you review the following topics and take the necessary actions:

- See "About the Veritas InfoScale product suite" on page 7.
- See "About InfoScale licenses" on page 17.

- See "About the co-existence of InfoScale products" on page 23.

Run all the commands in the command window in the **Run as administrator** mode.

To perform a silent installation of the server components

1. Open a command window and navigate to the folder where `Setup.exe` is located.
2. Use the following command syntax to install the product:

```
Setup.exe /s SOLUTIONS="1" INSTALL_MODE=InstallMode
  Telemetry=Telemetry [INSTALLDIR="InstallDirPath"]
  [REBOOT=RebootMode] [NODE="SysA"] [LICENSEKEY="LicenseKey"]
  [OPTIONS="a,b,c,..."] NoOptionDiscovery= NoOptionDiscovery
```

   The maximum length of the argument string is 2048 characters, and the syntax is not case sensitive.

   See "Parameters for Setup.exe" on page 32.

After the installation is complete, check the SORT website for the applicable patches, agents, or the array-specific modules, if any, to be installed:

https://sort.veritas.com/

About installing InfoScale Storage in a Microsoft failover cluster

If you plan to install InfoScale Storage in an active Microsoft failover cluster, ensure that the applicable prerequisites are met.

See "Requirements for installing InfoScale Storage in a Microsoft Failover Cluster" on page 15.

Then, use the rolling-installation method as follows:

1. Install InfoScale Storage on each inactive cluster node.
2. Move the cluster resources from the active node to any other inactive node.
3. Install InfoScale Storage on the now inactive node.

After you install InfoScale Storage in a Microsoft failover cluster environment, perform the following tasks:

- If the Failover Clustering featue is installed on the system but the cluster is not yet created, register the InfoScale Storage resources after the cluster is created. See "Registering the InfoScale Storage resource DLLs" on page 37.

- If the Microsoft failover cluster is already active, remove the physical disk resources for all the basic disks before configuring the SFW cluster disk groups. Otherwise, a disk reservation conflict occurs.

Parameters for Setup.exe

The following table describes the parameters that you can use with the `Setup.exe` command.

Parameter	Use
/s	Set for silent mode. If not set, boots the product installation wizard.
INSTALL_MODE	Set to indicate an installation or uninstallation. 1 = To install 5 = To uninstall Example: INSTALL_MODE=1 **Note:** The parameter, INSTALL_MODE=1 is used for both a new installation, as well as an upgrade. The installer switches to the correct mode (installation or upgrade) depending upon what has already been installed on the selected system.
SOLUTIONS	Set to the type of installation. 1 = Veritas InfoScale Storage 2 = Veritas InfoScale Enterprise 3 = Veritas InfoScale Availability 4 = Veritas InfoScale Foundation Example: SOLUTIONS=1
Telemetry	Set this parameter to participate in the Veritas Product Improvement Program by submitting system and usage information anonymously. The Product Improvement Program allows the product installer to collect installation, deployment, and usage data and submit it anonymously to Veritas. The collected information helps identify how customers deploy and use the product. If you do not want to participate in the Product Improvement Program, set this parameter to 0.

Parameter	Use
INSTALLDIR	Set the installation directory path. The path must start and end with a quotation mark. The default setting is SystemDrive: \Program files\Veritas Example: INSTALLDIR="C:\InstallationDirectory" This is an optional parameter. **Note:** If you plan to configure the cluster for single sign-on authentication and your system runs a non-English locale operating system, ensure that the installation directory contains only English characters.
Reboot	Set for the automatic reboot of the system at the completion of the installation. 0 = No reboot 1 = Reboot The default setting is 0 for no system reboot. Example: Reboot=1 **Note:** This is an optional parameter.
Node	Set the node name. Specify only one node at a time. The local node is the default setting when the node is unspecified. The machine name of the node must start and end with a quotation mark ("). Example: Node="PC177VM-3"
LICENSEKEY	Set the license key for the installation. Enter multiple keys by separating them with a comma (e.g. 123-345-567-789-123, 321-543-765-789-321, etc.) The license key must start and end with a quotation mark ("). LicenseKey has no default setting. Example: LICENSEKEY="123-234-123-234-345" **Note:** This parameter is applicable only if you plan to use the "User entered license key" as your license type. You need not specify this parameter for "Keyless" license type. For details about the InfoScale licenses, see: See "About InfoScale licenses" on page 17.

Parameter	Use
Options	Set the desired options, if any. The option must start and end with a quotation mark ("). Multiple options can be entered, using a comma as a separator. Options differ depending on your product and environment. There are no default settings. For details about the available product options and their usage, See "Available product options and supported DMP DSMs" on page 34.
NoOptionDiscovery	Set this parameter to uninstall the previously installed options during an upgrade. Default value is 0. If this parameter is set to 0, the setup discovers the previously installed options which are not specified in the OPTIONS parameter, and the setup exits. Rerun the setup and either include the previously installed options individually in the OPTIONS parameter or specify "Installed" in the OPTIONS parameter. If you set this parameter to 1 during an upgrade, the setup uninstalls the previously installed options which are not specified in the OPTIONS parameter.

Available product options and supported DMP DSMs

Microsoft Failover Cluster and DMP DSMs are available as the selectable product options. These product options are available depending on the product being installed.

The Microsoft Failover Cluster option installs the SFW component that is required if you plan to configure a Microsoft Failover Cluster.

The following table provides details about whether or not the options are applicable for the respective product:

Table 2-2 InfoScale product options

Product	Microsoft Failover Cluster	DMP DSM
Veritas InfoScale Foundation	X	✓

Table 2-2 InfoScale product options *(continued)*

Product	Microsoft Failover Cluster	DMP DSM
Veritas InfoScale Availability	X	X
Veritas InfoScale Storage	✓	✓
Veritas InfoScale Enterprise	X	✓

The following table provides a list of the available DMP DSMs.

Table 2-3 List of available DSMs

DMP device-specific modules (DSMs)	■ 3PARDATA (V3PARAA) ■ Compellent array (VCOMPLNT) ■ Dell EqualLogic array (VEQLOGIC) ■ Dell PowerVault MD3xxx (VDELLMD) ■ EMC Clarion (VEMCCLAR) ■ EMC Symmetrix/DMX (VEMCSYMM) ■ EMC VPLEX array (VEMCVPLX) ■ EMC XTREMEIO (VXTREMIO) ■ FUJITSU ETERNUS 2000 array (VFUJITSUAA) ■ Hitachi 95xx-AMS-WM (VHDSAP) ■ Hitachi TagmaStore/HP XP (VHDSAA) ■ HP 2000 array (VHPMSA2) ■ HP EVA-MSA (VHPEVA) ■ HUAWEI S5300/S2300 array (VHUAWEIAP) ■ IBM DS AP (VIBMAPDS) ■ IBM DS6000 (VIBMAP) ■ IBM DS4000/SUN 6000 (VENGAP) ■ IBM DS8000/ESS (VIBMAADS) ■ IBM FlashSystem (VIBFLASH) ■ IBM XiV Storage System (VXIV) ■ NETAPP (VNETAPP) ■ NEXSAN SATA/SAS Beast, E60/E18 array (VNEXSAN) ■ PILLAR (VPILLAR) ■ SUN Array - (VSUN) ■ VIOLIN V3000, V6000 (VVIOLIN) ■ VORACLE ■ NexentaStor (VNEXENTA) ■ NFINIDAT InfiniBox (VNFINIDAT)

For the latest information about the supported DSMs, refer to the hardware compatibility list (HCL) at:

https://sort.veritas.com/documents

Notes:

- 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.
- If you are upgrading the product using the product installer, do not clear the selection for the DSMs you want to remove. Clearing the default selection does not remove the installed DSMs.

  To remove the DSMs, perform any one of the following:

  - Before you begin to upgrade the cluster, remove the required DSM, using the Windows Add or Remove Programs. Reboot the node and then perform the upgrade.
  - Upgrade the cluster and then use the Windows Add or Remove Programs to remove the DSM.
  - Upgrade the cluster and then navigate to %ALLUSERSPROFILE%\Veritas\MPIO\.

    From the command prompt, run the following command:

    `instdsm.exe -u DSMName.inf`

  Note: If you clear the default selection during the upgrade, you cannot remove the DSM using the Windows Add or Remove Programs. To remove the DSM in this case, you must run the command that is mentioned earlier.

- If you are upgrading the product using CLI, you must specify the previously installed options in the OPTIONS parameter, else they will be uninstalled. To include the previously installed options in this parameter, either specify these options individually in the OPTIONS parameter or specify “Installed” in the OPTIONS parameter to upgrade all options (example: options=”Installed,flashsnap”).

- If you install InfoScale Availability on a system where you have already installed Veritas InfoScale Storage and have configured Microsoft Failover Cluster, you must first unconfigure the Microsoft Failover Cluster and remove the SFW component for Microsoft Failover Cluster.

  To unconfigure Microsoft Failover Cluster, refer to Microsoft documentation.

  To remove the SFW component for Microsoft Failover Cluster, use Windows Add Remove Programs.
- If you install InfoScale Storage on a system where you have installed InfoScale Availability, you cannot specify the Microsoft Failover Cluster option.

Registering the InfoScale Storage resource DLLs

You must perform this task only if you have installed InfoScale Storage with Microsoft failover cluster option, but Microsoft failover cluster is not yet configured in your environment.

- Using Windows Powershell cmdlets:
  - Import the FailoverClusters module
    Type the following cmdlet:
    ```
    Import-module failoverclusters
    ```
  - Register the Volume Manager Disk Group (VMDg) resource type
    Type the following cmdlet:
    ```
    Add-ClusterResourceType "Volume Manager Disk Group"
    C:\Windows\Cluster\vxres.dll -DisplayName "Volume Manager Disk
    Group"
    ```
  - Register the Replicated Volume Group (RVG) resource type
    Type the following cmdlet:
    ```
    Add-ClusterResourceType "Replicated Volume Group"
    C:\Windows\Cluster\mscsrvgresource.dll -DisplayName "Replicated
    Volume Group"
    ```
  - Register the Volume Manager Shared Volume resource type
    Type the following cmdlet:
    ```
    Add-ClusterResourceType "Volume Manager Shared Volume"
    C:\Windows\Cluster\vxvolres.dll -DisplayName "Volume Manager
    Shared Volume"
    ```

Installing the client components

To install the client components

1. Open the following link in a browser to download the client components.

   https://www.veritas.com/content/trial/en/us/vcs-utilities

2. Provide your contact information in the appropriate fields, and click **SUBMIT**.
3. Click **Download Now** corresponding to the client components you wish to install on your local system or a cluster node.

   Note: Client components cannot be installed on server core systems.

4. Double-click a downloaded file to launch the installer, and follow the instructions to complete the installation.

Chapter 3

Upgrading to InfoScale products

This chapter includes the following topics:

- Preparing the systems for an upgrade
- Performing the product upgrade
- About transitioning between the InfoScale products

Preparing the systems for an upgrade

Before you begin with the product upgrade, you must ensure that you perform the required pre-upgrade tasks. These tasks prepare the systems for the product upgrade.

The pre-upgrade tasks depend on your product deployment.

The following sections provide details on the tasks that must be performed for preparing the systems for a product upgrade.

About the supported upgrade paths and the supported minimum product versions

To upgrade to an InfoScale product, your systems must have a corresponding Veritas Storage and High Availability product. The following table lists the Veritas Storage and High Availability/InfoScale product versions from which you can upgrade to its 7.2 version:

Upgrade from	Upgrade to
6.0.2	7.2
6.x	
7.0	
7.0.x	
7.1	

Note: If your existing deployment setup employs Windows Server 2008 or Windows Server 2008 R2 Operating System, the you must first upgrade the OS to a supported OS and then upgrade the product.

If your current installation does not meet this minimum required level, you must manually apply the appropriate product upgrades before you proceed with the upgrade. You can get the intermediate versions of the products on the Veritas Support site:

https://www.veritas.com/support/en_US.html

For license keys, contact Veritas Sales. You can also uninstall the older versions of the product and install the new product.

The following table lists the supported upgrade paths.

Table 3-1 Supported upgrade paths

Upgrade from	Upgrade to			
	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
SFW Basic/InfoScale Foundation	✓	X	X	X
SFW/InfoScale Storage	X	X	✓	X
SFW HA/InfoScale Enterprise	X	X	X	✓
VCS/InfoScale Availability	X	✓	X	X
DMP	✓	X	X	X

Table 3-1 Supported upgrade paths *(continued)*

Upgrade from	Upgrade to			
VBS (Virtual Business Service 6.1 for 3rd party support)	✓	X	X	X

Note: While you plan to install the product, you may have to upgrade your Windows operating system to the supported minimum level. We recommend you to perform the Windows upgrade before you upgrade the product.

See "Supported hardware and software" on page 11.

General preparations

Perform the following general pre-upgrade checks or tasks on all the systems:

The following table lists the general pre-upgrade checks.

Table 3-2 General pre-upgrade checks

Tasks	Applicable for			
	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
Back up configuration and application data	✓	✓	✓	✓
Review licensing details	✓	✓	✓	✓
Review the installation requirements and perform the required pre-installation tasks	✓	✓	✓	✓
Ensure there are no parallel scheduled snapshots in progress	✓	✓	✓	✓
Ensure that the latest Microsoft Windows updates are installed on the systems. The latest updates are required for successfully installing Microsoft Visual C++ 2015 Redistributable-x64. The product installation fails if Microsoft Visual C++ 2015 Redistributable-x64 is not installed.	✓	X	✓	✓

Table 3-2 General pre-upgrade checks *(continued)*

Tasks	Applicable for			
Verify that the systems in your environment meet the requirements to upgrade the cluster. Run the Windows Data Collector that is available on the SORT website to perform these checks.	✓	✓	✓	✓
Ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress	✓	✓	✓	✓
If the systems have NetBackup version 6.0 or 6.5 installed and running, then shut down the OpsCenterServer service. Both, NetBackup and InfoScale products share the same AT broker and client.	✓	✓	✓	✓
Save and close the cluster configuration. This operation saves the latest configuration to disk and changes the configuration state to read-only mode To save the cluster configuration, perform one of the following tasks:	X	✓	X	✓
Take the backup of custom agent binaries During the product upgrade, a backup of the main.cf and other configuration files is taken. However, it does not take the backup of any agent binaries. During the upgrade, the contents of %VCS_home% folder are removed. As a result, all the binaries of all the enterprise agents and custom agents that were installed are removed. After the upgrade is complete, all the binaries of enterprise agents are installed again. However, the binaries of a custom agent are not installed again. The main.cf that is restored after the upgrade shows that the custom agent resources are configured. However, the binaries are not present in the %VCS_home% folder. You must manually install the custom agents after the upgrade is complete.	X	✓	X	✓

Table 3-2 General pre-upgrade checks *(continued)*

Tasks	Applicable for			
To perform parallel upgrade, take the service groups offline To take the service groups offline **1** From the command prompt, type: `C:\>hagrp -offline group_name -sys system_name` Where 'group_name' is the name of the service group and system_name is the node on which the group is online. **2** Repeat this command for all service groups that are online.	X	✓	X	✓
Close client applications If you are running any of the following client applications on the systems where you plan to upgrade the product, stop and exit all the application instances. ■ Cluster Manager (Java Console) ■ Solutions Configuration Center (SCC) ■ Veritas Enterprise Administrator (VEA)	✓	✓	✓	✓
Export the configured rules If you have configured any rules for event notification messages and actions, you must export them into XML format before you begin with the upgrade. To export the configured rules **1** From the VEA Control Panel perspective, select the action agent node in the tree view and double-click **Rule Manager** in the right pane. **2** In the Rule Manager window, select the rules you want to export and then click **Export**. **3** Save the rules at any temporary location in the XML format.	✓	X	✓	✓

Table 3-2 General pre-upgrade checks *(continued)*

Tasks	Applicable for			
Stop the Veritas High Availability Engine (HAD) on all the cluster nodes This task is applicable only if Veritas InfoScale Storage and InfoScale Availability co-exists in your environment. If you plan to upgrade Veritas InfoScale Storage in a co-existing environment, you must forcefully stop HAD before you begin the upgrade. To forcefully stop HAD, run the following command from any cluster node: `hastop -all -force`	X	X	✓	X

Recommendations and considerations for product upgrade

Note the following points before you begin with the upgrade:

- During the product upgrade if you also plan to upgrade the operating system, you must first upgrade the operating system and then upgrade the product.
- During the upgrade, verify the selected installation and product options. All the installed options are selected by default.
  If you do not want to include any of the installed options in the upgraded environment, you must uninstall those options before you proceed with the upgrade. Use the Windows Add or Remove Programs feature to uninstall the option.
  If you want to add any additional options, you must select the same.
- You can continue to use the 6.0.x or 6.x license keys. However, note that the features available with the 6.0.x or 6.x license keys will only be enabled after the product upgrade. To enable all the available features, specify a 7.2 license key. Feature-based license keys are not available for 7.2.
  For details, See "Licensing notes" on page 18.
- When you upgrade to InfoScale Foundation, InfoScale Storage, or InfoScale Enterprise the product installation wizard replaces the Disk Management Snap-in in the Windows Computer Management console and the Server Manager console with the Veritas Enterprise Administrator (VEA) GUI. To change this default, access the VEA GUI after the upgrade completes and proceed to restore the Disk Management Snap-in.
  For information about using the VEA GUI, see *Storage Foundation Administrator's Guide*.

- If you upgrade SFW to InfoScale Storage or InfoScale Enterprise in a Volume Replicator environment, you must first perform the upgrade at the secondary site. After you have completed the upgrade and the post upgrade tasks at the secondary site, you must fail over the applications from the primary site to the secondary site and then proceed to upgrade the nodes at the primary site. Once the nodes at the primary site are upgraded, you must migrate the applications back.

  Since Volume Replicator 5.0.x and Volume Replicator 5.1 versions are interoperable, an upgrade can be performed on the secondary site while applications are kept running on the primary site.
- If you upgrade InfoScale Foundation, InfoScale Storage, or InfoScale Enterprise and your configuration has DMP configured, then it is recommended to reduce the number of paths to each array to one, before you begin the upgrade.

  If you do not have DMP DSMs in your existing environment, but plan to add this feature during the upgrade, add the HBA(host bus adapter) hardware before performing the upgrade. Connect no more than one path from the new HBA to the storage array before the upgrade and DMP DSMs installation. Select the DMP DSM option or the appropriate DMP DSMs while running the installer.
- Veritas InfoScale products do not provide Dynamic Multi-Pathing support for PROMISE arrays. If currently installed, Dynamic Multi-Pathing support for PROMISE will be removed after the upgrade is complete.

Performing the product upgrade

The product upgrade tasks are based on the type of configuration you have deployed.

Typically, the upgrade scenarios can be classified in the following categories:

- Upgrading SFW Basic or SFW in a non-clustered environment
- Upgrading SFW Basic or SFW in a Microsoft Failover Cluster environment
- Upgrading VCS
- Upgrading SFW HA
- Upgrading SFW HA in a setup that has replication configured
- Upgrading DMP

Note: If you had installed the SFW or VCS clients only, you cannot upgrade them to the current version. To install the latest version of clients-only, uninstall the earlier version and then install the current version.

See "Installing the client components" on page 38.

The following sections provide details about the tasks that are involved in each of these scenarios and the corresponding InfoScale product to upgrade to.

Note: Before you begin with the product upgrade, ensure that you have reviewed the supported upgrade paths, the installation prerequisites, and the licensing details.

See "Preparing the systems for an upgrade" on page 39.

Upgrading SFW or SFW Basic in a non-clustered environment

This section describes the tasks to be performed to upgrade SFW Basic or SFW in a non-clustered environment.

Supported upgrades:

- SFW Basic to InfoScale Foundation
- SFW to InfoScale Storage

The upgrade tasks depend on whether or not replication is configured in your environment.

The following figure depicts the tasks to be performed for upgrading SFW in a non-clustered environment, where replication is not configured.

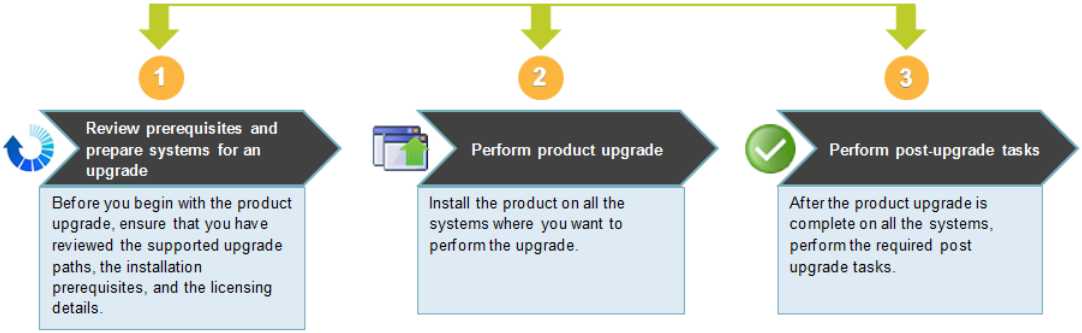
Note: If you have installed InfoScale Storage on a machine where InfoScale Availability co-exists, then before you begin to upgrade the product, you must run the BackupRestoreUtil.exe using Backup_HA.xml as an input parameter and then uninstall InfoScale Availability.

After the product upgrade is complete, you must install InfoScale Availability and then run the BackupRestoreUtil.exe using Restore_HA.xml as an input parameter.

To run the BackupRestoreUtil.exe, using the command prompt navigate to the following location:

InfoScale-7.1\Installer>BackupRestoreUtil.exe

Figure 3-1 SFW upgrade tasks in a non-clustered environment



References

Task	Reference
1	See "Preparing the systems for an upgrade" on page 39.
2	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
3	See "Deployment scenarios and applicable post upgrade tasks" on page 66.

The following figure depicts the tasks to be performed for upgrading SFW in a non-clustered environment, where replication is configured

---

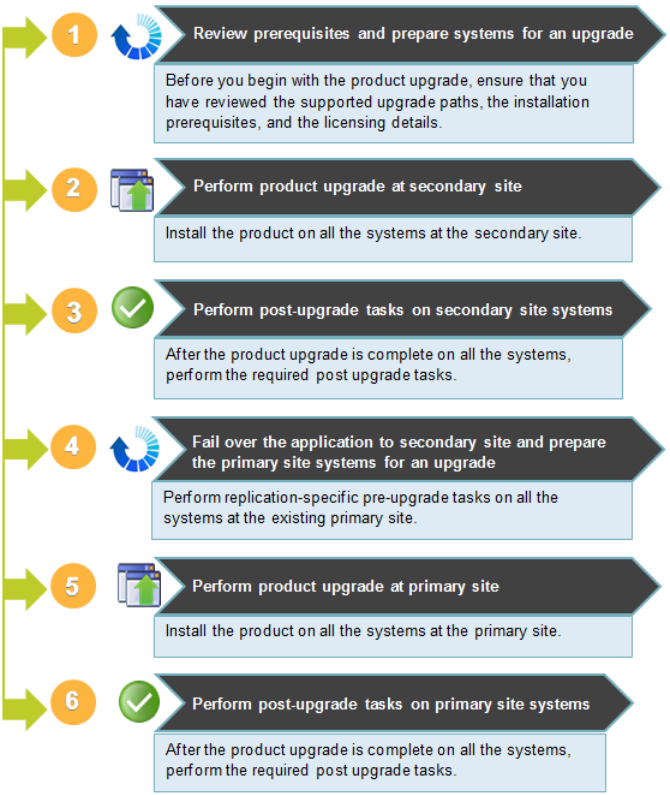
Note: If you have installed InfoScale Storage on a machine where InfoScale Availability co-exists, then before you begin to upgrade the product, you must run the BackupRestoreUtil.exe using Backup_HA.xml as an input parameter and then uninstall InfoScale Availability.

After the product upgrade is complete, you must install InfoScale Availability and then run the BackupRestoreUtil.exe using Restore_HA.xml as an input parameter.

To run the BackupRestoreUtil.exe, using the command prompt navigate to the following location:

InfoScale-7.1\Installer>BackupRestoreUtil.exe

Figure 3-2 SFW upgrade tasks in a non-clustered environment with replication configured



References

Task	Reference
1	See "Preparing the systems for an upgrade" on page 39.
2	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
3	See "Deployment scenarios and applicable post upgrade tasks" on page 66.
4	See "Preparing the primary site for upgrade in a non-clustered SFW environment" on page 49.

Task	Reference
5	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
6	See "Deployment scenarios and applicable post upgrade tasks" on page 66.

Preparing the primary site for upgrade in a non-clustered SFW environment

Perform the following procedure to prepare the systems on the primary site for the SFW upgrade in a Volume Replicator environment.

Note: Before you prepare the nodes on the primary site, ensure that you have upgraded SFW and performed the post-upgrade tasks on the secondary site.

To upgrade the cluster in a Volume Replicator environment, you must first stop the replicated volume group (RVG) to detach the replication links and disassociate the replication logs between the primary and secondary site.

To prepare the primary site

1. On the primary site, stop the application that uses Volume Replicator to replicate data between the sites.
2. From the command line, type:

   ```
   vxprint -lVP [-g diskgroup_name]
   ```

   This command lists the RLINK and RVG records.
3. Verify that the data on the Replicator Log is written to the secondary site by running the following command on the primary site:

   ```
   vxrlink [-g diskgroup_name] status rlink_to_secondary
   ```

   This command displays the replication status of the secondary site represented by the specified RLINK.

   Verify that the data volumes on the secondary site are consistent and up-to-date with the primary site before proceeding to the next step.
4. Migrate the primary RVG by performing one of the following procedures:
   - From the Veritas Enterprise Administrator (VEA) console, right-click the primary RVG and select the **Migrate** option. Select the required secondary host from the Secondary Name option list.

Click **OK** to migrate the primary role to the secondary. The primary and secondary roles will be interchanged.

- From the command line, type:

  ```
  vxrds [-g diskgroup_name] migrate local_rvg
  new_primary_hostname
  ```

  Where the secondary host is specified by the *new_primary_hostname* parameter.

5. Perform any necessary steps to start the applications on the new primary (old secondary).
6. If the existing replication settings are configured to use TCP, change the settings to use UDP. After both the primary and DR sites are upgraded, you can switch the replication settings back to TCP.

   You can now proceed to upgrade the nodes.

Upgrading SFW or SFW Basic in a Windows Server Failover Cluster environment

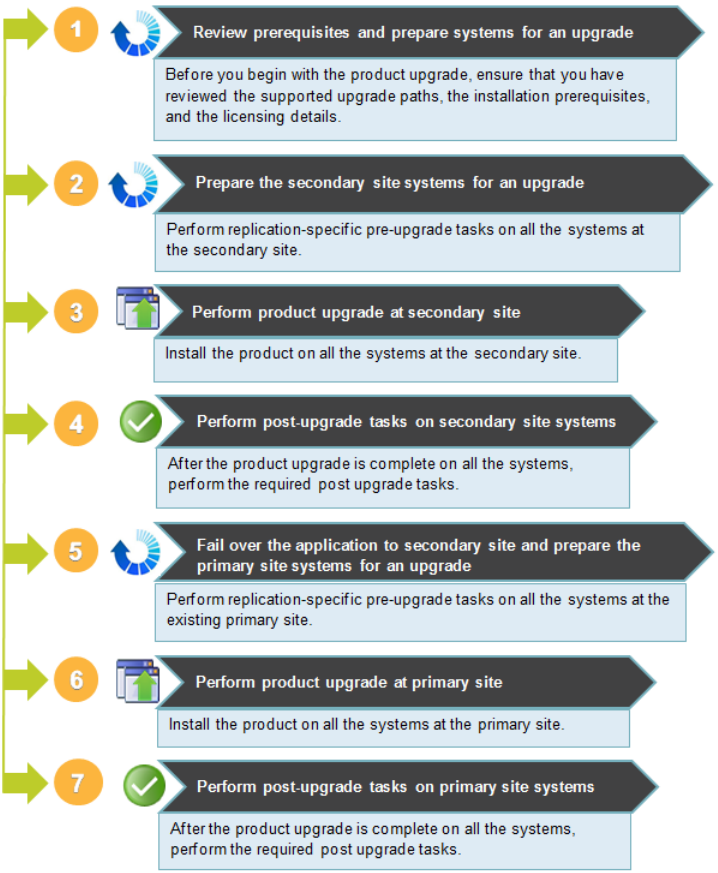
This section describes the tasks to be performed to upgrade SFW in a Windows Server Failover Cluster.

Supported upgrade paths:

- SFW to InfoScale Storage
- SFW Basic to InfoScale Foundation

Note: InfoScale Foundation does not support Windows Server Failover Cluster. To avail the Windows Server Failover Cluster capabilities, you must further transition to InfoScale Storage. See “About transitioning between the InfoScale products” on page 64.

Figure 3-3 SFW upgrade tasks in a Windows Server Failover Cluster



Note: After the upgrade is complete, the installer wizard performs certain cleanup tasks and a command prompt is displayed for certain time. Ignore the prompt and continue with the further tasks. The command prompt closes after the cleanup tasks are complete. Do not close the command prompt before the cleanup tasks are complete. The cleanup tasks are aborted if you close the prompt.

References

Task	References
1	See "Preparing the systems for an upgrade" on page 39.

Task	References
2	See "Preparing the secondary site for SFW upgrade in a Windows Server Failover Cluster environment" on page 52.
3	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
4	See "Deployment scenarios and applicable post upgrade tasks" on page 66.
5	See "Failing over application to secondary site" on page 53. See "Preparing the primary site for SFW upgrade in a Windows Server Failover Cluster environment" on page 53.
6	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
7	See "Deployment scenarios and applicable post upgrade tasks" on page 66.

Preparing the secondary site for SFW upgrade in a Windows Server Failover Cluster environment

Perform the following procedure to prepare the nodes on the secondary site for the SFW upgrade in a Windows Server Failover clustered Volume Replicator environment.

To prepare the secondary site

1. Take the RVGresource offline by performing one of the following procedures:
   - From the Cluster Administrator console, right-click the RVG resource and select the **Offline** option.
   - From the command line, type:
     ```
     [cluster resourcename] /Offline [:node name] [/Wait[:timeoutin seconds]]
     ```
2. Take the Disk Group resource offline on the secondary site by performing one of the following procedures:
   - From the Cluster Administrator console, right-click the Disk Group resource, and select the **Offline** option.
   - From the command prompt, type:
     ```
     [cluster resourcename] /Offline [:node name] [/Wait[:timeoutin seconds]]
     ```

Repeat step 2 to offline the IP resource and then the Network Name resource.

Warning: Taking the DG(Disk Group) resource offline pauses replication, and if applications continue to run on the primary for too long, this may cause a possible replicator log overflow.

3 If the existing replication settings are configured to use TCP, change the settings to use UDP. After both the primary and DR sites are upgraded, you can switch the replication settings back to TCP.

You can now proceed to upgrade SFW on all nodes of the secondary site.

Failing over application to secondary site

Perform the following steps to fail over the application to the secondary site.

To fail over the application

1 From the Failover Cluster console, navigate to the Cluster Group.

2 Right-click the Cluster Group and click Move Group.

This procedure moves the resources and the Resource Owner changes to a node on the secondary site.

Preparing the primary site for SFW upgrade in a Windows Server Failover Cluster environment

Once the secondary site is upgraded, the primary site can be prepared for upgrade by migrating the primary role to the DR site.

To upgrade the cluster in a Volume Replicator environment, you must first stop the replicated volume group (RVG) to detach the replication links and disassociate the replication logs between the primary and secondary site.

Perform the following procedure to prepare the primary site in a Windows Server Failover clustered Volume Replicator environment.

To prepare the primary site

1 Offline the Application resource on the primary site by performing one of the following procedures:

- From the Cluster Administrator console, right-click the Application resource and select the **Offline** option.
- From the command line, type:

```
[cluster resourcename] /Offline [:node name] [/Wait[:timeoutin
seconds]]
```

2 From the command line, type:

```
vxprint -lVP [-g diskgroup_name]
```

This command lists the RLINK and RVG records.

3 Verify that the data on the Replicator Log is written to the secondary site by running the following command on the primary:

```
vxrlink [-g diskgroup_name] status rlink_to_secondary
```

This command displays the replication status of the secondary represented by the specified RLINK.

Verify that the data volumes on the secondary site are consistent and up-to-date with the primary before proceeding to the next step.

4 Migrate the primary RVG by performing one of the following procedures:

- From the Veritas Enterprise Administrator (VEA) console, right-click the primary RVG and select the **Migrate** option. Select the required secondary host from the Secondary Name option list.
  Click **OK** to migrate the primary role to the secondary. The primary and secondary roles will be interchanged.
- From the command line, type:
  ```
  vxrds [-g diskgroup_name] migrate local_rvg
  new_primary_hostname
  ```
  Where the secondary host is specified by the *new_primary_hostname* parameter.

5 Bring online the Application resource on the new primary by performing one of the following procedures:

- From the Cluster Administrator console, right-click the Application resource and select the **Online** option.
- From the command line, type:
  ```
  [cluster resourcename] /Online [:node name] [/Wait[:timeoutin
  seconds]]
  ```

6 If bringing the service groups online does not start the application, perform any necessary tasks to start the application. Depending on the options available in your environment, these tasks may include mounting the databases or manually starting the application.

Upgrading VCS

This section describes the tasks to be performed to upgrade VCS.

Supported upgrade path: VCS to InfoScale Availability

The upgrade tasks depend on whether or not replication is configured in your environment.

The following figure lists the tasks to be performed to upgrade VCS, where replication is not configured.

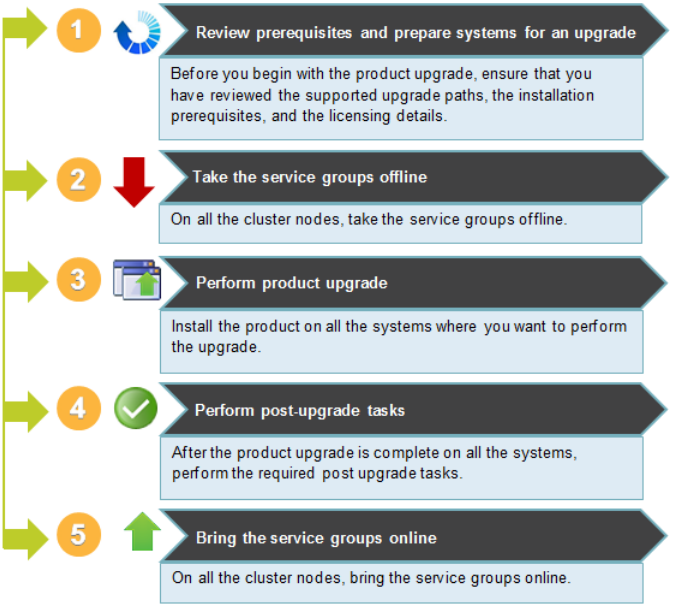
Note: If you have installed InfoScale Storage on a machine where InfoScale Availability co-exists, then before you begin to upgrade the product, you must run the BackupRestoreUtil.exe using Backup_HA.xml as an input parameter and then uninstall InfoScale Availability.

After the product upgrade is complete, you must install InfoScale Availability and then run the BackupRestoreUtil.exe using Restore_HA.xml as an input parameter.

To run the BackupRestoreUtil.exe, using the command prompt navigate to the following location:

InfoScale-7.1\Installer>BackupRestoreUtil.exe

Figure 3-4 VCS upgrade tasks

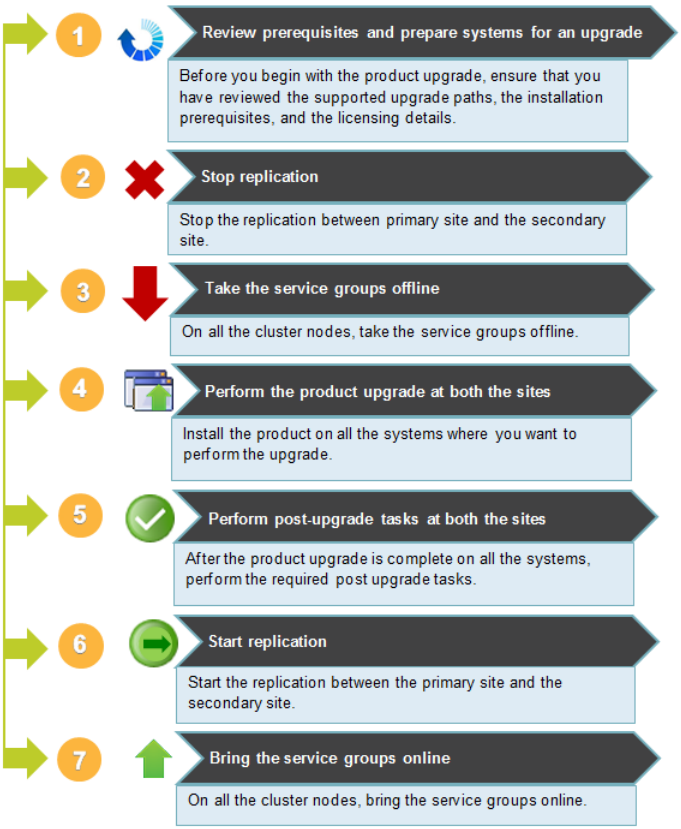


References

Task	References
1	See "Preparing the systems for an upgrade" on page 39.
2	Take the service groups offline on all nodes. Use the following command: `hagrp -offline service_group -sys system`
3	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
4	See "Deployment scenarios and applicable post upgrade tasks" on page 66.
5	Bring the service group online. Use the following command: `hagrp -online service_group -sys system`

The following figure lists the tasks to be performed to upgrade VCS, where replication is configured.

Figure 3-5 VCS upgrade tasks in a replicated cluster



References

Task	References
1	See "Preparing the systems for an upgrade" on page 39.
2	Refer to the 3rd party replication technology-specific documentation
3	Take the service groups offline on all nodes. Use the following command: `hagrp -offline service_group -sys system`
4	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.

Task	References
5	See "Deployment scenarios and applicable post upgrade tasks" on page 66.
6	Refer to the 3rd party replication technology-specific documentation
7	Bring the service group online. Use the following command: `hagrp -online service_group -sys system`

Upgrading SFW HA

This section describes the tasks to be performed while upgrading SFW HA.

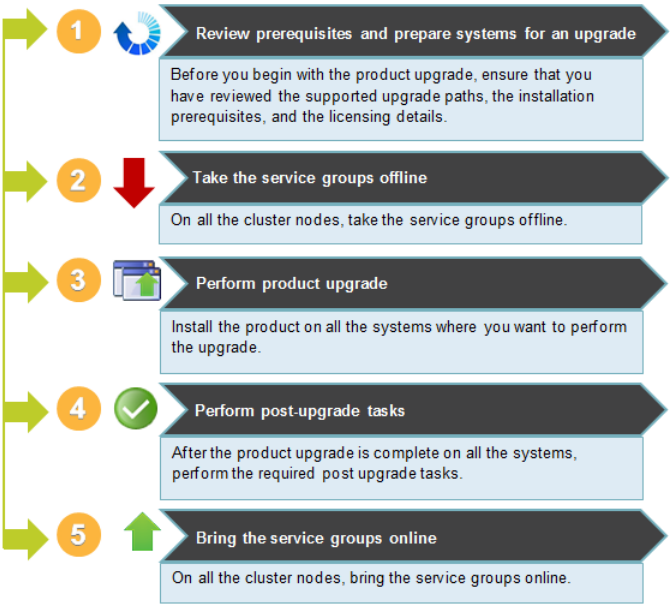
Supported upgrade path: SFW HA to InfoScale Enterprise

The upgrade tasks depend on whether or not replication is configured in your environment.

Note: After the upgrade is complete, the installer wizard performs certain cleanup tasks and a command prompt is displayed for certain time. Ignore the prompt and continue with the further tasks. The command prompt closes after the cleanup tasks are complete. Do not close the command prompt before the cleanup tasks are complete. The cleanup tasks are aborted if you close the prompt.

The following figure lists the tasks to be performed to upgrade VCS, where replication is not configured.

Figure 3-6 SFW HA upgrade tasks

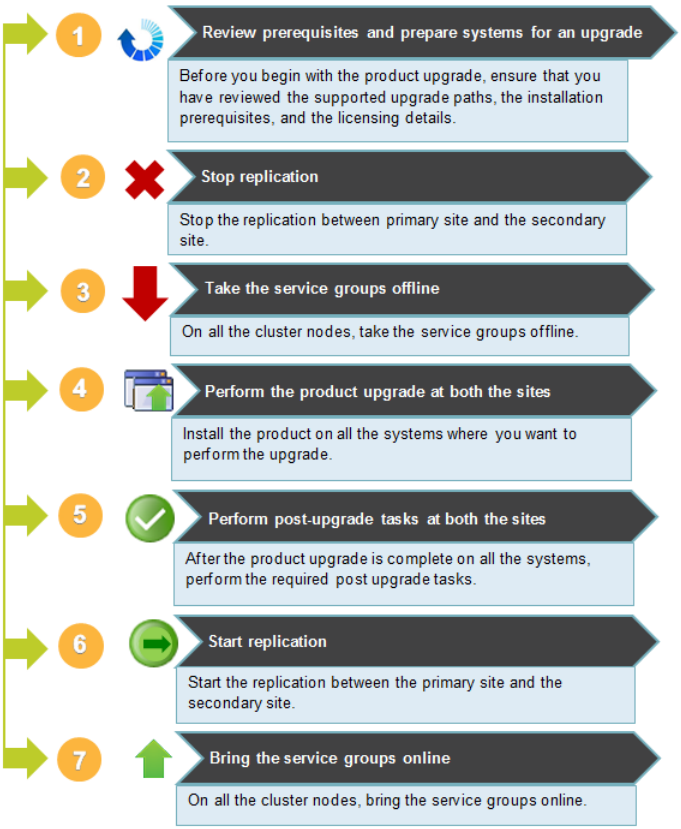


References:

Task	References
1	See "Preparing the systems for an upgrade" on page 39.
2	Take the service groups offline on all nodes. Use the following command: `hagrp -offline` *`service_group`* `-sys` *`system`*
3	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
4	See "Deployment scenarios and applicable post upgrade tasks" on page 66.
5	Bring the service groups online. Use the following command: `hagrp -online` *`service_group`* `-sys` *`system`*

The following figure lists the tasks to be performed to upgrade SFW HA, where replication is configured.

Figure 3-7 SFW HA upgrade tasks in a replicated cluster



References:

Task	Reference
1	See "Preparing the systems for an upgrade" on page 39.
2	See "Preparing the primary and secondary sites for upgrading SFW HA in a Volume Replicator environment" on page 61.
3	Take the service groups offline on all nodes. Use the following command: `hagrp -offline service_group -sys system`

Task	Reference
4	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
5	See "Deployment scenarios and applicable post upgrade tasks" on page 66.
6	See "Re-enabling Volume Replicator in a VCS cluster" on page 62.
7	Bring the service groups online. Use the following command: `hagrp -online service_group -sys system`

Preparing the primary and secondary sites for upgrading SFW HA in a Volume Replicator environment

To upgrade the SFW HA cluster in a Volume Replicator environment, you must first stop the replicated volume group (RVG) to detach the replication links and disassociate the replication logs between the primary and secondary site.

Perform the following steps from any one of the cluster nodes at the primary site

1. Bring the application that uses Volume Replicator to replicate data between its sites offline.
2. From the command prompt run the `vxprint -lVP` command.
3. Verify that the data on the Replicator Log is written to the secondary site and the RLINKSs are up-to-date.

   `vxrlink [-g diskgroup] status rlink_to_secondary`
4. Take the VvrRvg resource offline in the Volume Replicator replication service group.
5. Detach the RLINK to prevent Volume Replicator from replicating data to the secondary site. From the Veritas Enterprise Administrator console, right-click the secondary RVG and select the **Stop Replication** option to stop Volume Replicator from replicating to the secondary site.
6. Disassociate the Replicator Log from the RVG. From the Veritas Enterprise Administrator console, right-click the Replicator Log and select **Dissociate Replicator Log** option from the menu that appears.

Perform the following steps from any one of the cluster nodes at the secondary site

1. Take the VvrRvg resource offline in the Volume Replicator replication service group.
2. From the command prompt run the `vxprint -lVP` command.
3. Disassociate the Replicator Log from the RVG. From the Veritas Enterprise Administrator console, right-click the Replicator Log and select **Dissociate Replicator Log** option from the menu that appears.

Associating the replication logs and starting the replication

You must perform this task, if you have upgraded SFW HA in a Volume Replicator environment. Perform the task on any one of the upgraded node.

From the Veritas Enterprise Administrator, right-click the secondary RVG resource and select **Associate Replicator Log** option from the menu that appears. Also, select **Start Replication** option to enable Volume Replicator to begin the replication.

Re-enabling Volume Replicator in a VCS cluster

Follow the procedure below to enable the updated objects on the secondary site.

To enable the updated objects on the secondary site

1. Bring the Disk Group Resource online on the secondary site, by performing one of the following procedures:
   - From the Veritas InfoScale Operations Manager Management Server console, go to the **Availability** perspective, select the **Manage** tab, locate and expand the cluster and select the required service group. Then on the **Resources** tab, right-click on the desired Disk Group resource and click **Online**.
   - From the command line, type:

     ```
     hares -online resource_name -sys system_name
     ```

2. Bring the RVG service group online, by performing one of the following procedures:
   - From the Veritas InfoScale Operations Manager Management Server console, go to the **Availability** perspective and on the **Manage** tab, locate and expand the cluster and then right-click on the RVG service group and click **Online**.
   - From the command line, type:

```
hagrp -online group_name -sys system_name
```

For Volume Replicator environments with multiple secondary sites, any operations that need to be performed on a secondary site should be repeated on all secondary sites.

Upgrading DMP

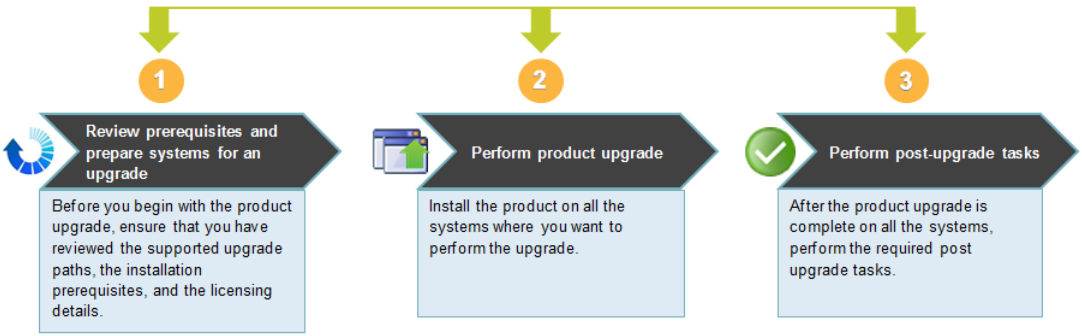
This section describes the tasks to be performed while upgrading DMP.

Supported upgrade path: DMP to InfoScale Foundation

Note: After the upgrade is complete, the installer wizard performs certain cleanup tasks and a command prompt is displayed for certain time. Ignore the prompt and continue with the further tasks. The command prompt closes after the cleanup tasks are complete. Do not close the command prompt before the cleanup tasks are complete. The cleanup tasks are aborted if you close the prompt.

The following figure lists the tasks to be performed to upgrade DMP, where replication is not configured.

Figure 3-8 DMP upgrade tasks



References:

Task	References
1	See “Preparing the systems for an upgrade” on page 39.

Task	References
2	See "Installing the server components using the installation wizard" on page 24. See "Installing the server components using the command-line installer" on page 30.
3	See "Deployment scenarios and applicable post upgrade tasks" on page 66.

About transitioning between the InfoScale products

You can transition between the products of the InfoScale family.

To transition to a product on a system where you already have an InfoScale product installed, run the product installer for the InfoScale product you want to transition to.

Note: The base product version and the version of the product to transition to must be same.

Supported paths:

The following table lists the supported paths to transition from one InfoScale product to another InfoScale product.

Table 3-3 Supported transition paths

Base product	Transition to			
	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
InfoScale Foundation	X	X	✓	✓
InfoScale Availability	X	X	X	✓

Table 3-3 Supported transition paths *(continued)*

Base product	Transition to			
InfoScale Storage	X	X	X	✓ **Note:** If you have configured Microsoft Failover Cluster, you must unconfigure it and remove the SFW component for the Failover Cluster, before transitioning to InfoScale Enterprise. To unconfigure Microsoft Failover Cluster, refer to Microsoft documentation. To remove the SFW component for Failover Cluster, use Windows Add Remove Programs.
InfoScale Enterprise	X	X	X	X

Chapter 4

Performing the post upgrade tasks

This chapter includes the following topics:

- Deployment scenarios and applicable post upgrade tasks
- Re-enabling Volume Replicator in a non-clustered environment
- Re-enabling Volume Replicator in a Microsoft failover cluster environment
- Reconnecting DMP DSM paths after the upgrade
- Reconfiguring the Veritas InfoScale Messaging Service
- Importing the configured rules
- Upgrading the dynamic disk group version
- Upgrading clusters for stronger security
- Reinstalling the custom agents
- Including custom resources

Deployment scenarios and applicable post upgrade tasks

The post-upgrade tasks are based on the type of configuration you have deployed.

The following table lists the typical product upgrade scenarios and the corresponding post-upgrade tasks.

Table 4-1 Typical product upgrade scenarios and the corresponding post-upgrade tasks

SFW Basic/ SFW (non-clustered environment)	SFW Basic/ SFW (Microsoft Failover Cluster environment)	VCS	SFW HA	SFW HA (Volume Replicator environment)	DMP
					
 Re-enable replication See "Re-enabling Volume Replicator in a non-clustered environment" on page 69.	 Re-enable replication See "Re-enabling Volume Replicator in a Microsoft failover cluster environment" on page 70.	---	---	Replication is re-enabled as part of the upgrade workflow	---
 Reconnect DSM paths, if applicable See "Reconnecting DMP DSM paths after the upgrade" on page 71.	 Reconnect DSM paths, if applicable See "Reconnecting DMP DSM paths after the upgrade" on page 71.	---	 Reconnect DSM paths, if applicable See "Reconnecting DMP DSM paths after the upgrade" on page 71.	 Reconnect DSM paths, if applicable See "Reconnecting DMP DSM paths after the upgrade" on page 71.	 Reconnect DSM paths See "Reconnecting DMP DSM paths after the upgrade" on page 71.

Table 4-1 Typical product upgrade scenarios and the corresponding post-upgrade tasks *(continued)*

SFW Basic/ SFW (non-clustered environment)	SFW Basic/ SFW (Microsoft Failover Cluster environment)	VCS	SFW HA	SFW HA (Volume Replicator environment)	DMP
 Re-configure services See "Reconfiguring the Veritas InfoScale Messaging Service" on page 71.	 Re-configure services See "Reconfiguring the Veritas InfoScale Messaging Service" on page 71.	---	 Re-configure services See "Reconfiguring the Veritas InfoScale Messaging Service" on page 71.	---	---
 Import rules See "Importing the configured rules" on page 72.	 Import rules See "Importing the configured rules" on page 72.	---	 Import rules See "Importing the configured rules" on page 72.	 Import rules See "Importing the configured rules" on page 72.	---
 Upgrade disk group See "Upgrading the dynamic disk group version" on page 72.	 Upgrade disk group See "Upgrading the dynamic disk group version" on page 72.	---	 Upgrade disk group See "Upgrading the dynamic disk group version" on page 72.	 Upgrade disk group See "Upgrading the dynamic disk group version" on page 72.	---

Table 4-1 Typical product upgrade scenarios and the corresponding post-upgrade tasks *(continued)*

SFW Basic/ SFW (non-clustered environment)	SFW Basic/ SFW (Microsoft Failover Cluster environment)	VCS	SFW HA	SFW HA (Volume Replicator environment)	DMP
---	---	 Re-install custom agents See "Reinstalling the custom agents" on page 75.	 Re-install custom agents See "Reinstalling the custom agents" on page 75.	 Re-install custom agents See "Reinstalling the custom agents" on page 75.	---
---	---	 Include custom resources See "Including custom resources" on page 75.	 Include custom resources See "Including custom resources" on page 75.	 Include custom resources See "Including custom resources" on page 75.	---

Re-enabling Volume Replicator in a non-clustered environment

After upgrading in a non-clustered environment where Volume Replicator replicates data from a primary site to a secondary site, you must re-enable Volume Replicator.

In the procedure for preparing the primary site for upgrade, you migrated the primary role to the secondary site.

After both the primary and secondary sites have been upgraded, you may want to migrate the role of the primary back to the original primary site. To do this, you perform a migrate operation again as described in the following procedure.

To migrate the applications back to the original primary

1. On the current primary site, stop the application that uses Volume Replicator to replicate data between the sites.
2. From the command line, type:

   ```
   vxprint -lVP [-g diskgroup_name]
   ```

   This command lists the RLINK and RVG records.
3. Verify that the data on the Replicator Log is written to the secondary site by running the following command on the primary:

   ```
   vxrlink [-g diskgroup_name] status rlink_to_secondary
   ```

   This command displays the replication status of the secondary represented by the specified RLINK.

   Verify that the data volumes on the secondary site are consistent and up-to-date with the primary before proceeding to the next step.
4. To migrate the primary RVG perform one of the following procedures:
   - From the VEA, right-click the primary RVG and select the Migrate option. Select the required secondary host from the Secondary Name option list. Click OK to migrate the primary role to the secondary.
     The primary and secondary roles will be interchanged.
   - From the command line, type:

     ```
     vxrds [-g diskgroup_name] migrate local_rvg
     new_primary_hostname
     ```

     Where the secondary host is specified by the *new_primary_hostname* parameter.
5. Perform any necessary steps to start the applications on the new primary (old secondary).

Re-enabling Volume Replicator in a Microsoft failover cluster environment

In a Microsoft clustered environment after you have completed upgrading SFW on all the cluster nodes, re-enable Volume Replicator on the active cluster node.

Warning: A full synchronization is required if the procedures listed below are not performed in the given order.

To enable the updated objects on the secondary (DR) site

1. Bring online the Disk Group, IP, and Network Name resource in the Windows Server Failover Cluster resource group.
2. Bring online the RVG resource by performing one of the following procedures:
   - From the Cluster Administrator console, right-click the RVG resource and select the Online option on the secondary.
   - From the command line, type:
     ```
     [cluster resourcename] /online [:node name] [/wait[:timeoutin seconds]]
     ```

     Note: Refer to the appropriate Microsoft documentation for details on how to offline and online resources through the command line interface.

For Volume Replicator environments with multiple secondary sites, any operations that need to be performed on a secondary site should be repeated on all secondary sites.

Reconnecting DMP DSM paths after the upgrade

After you complete the upgrade for an existing DMP DSM environment or if you have added DMP DSMs during the upgrade, proceed to reconnect the DMP DSM paths:

To reconnect DMP DSM paths after the upgrade

1. Physically connect any additional paths that were disconnected before the upgrade.
2. Rescan the disks from the VEA console.

Reconfiguring the Veritas InfoScale Messaging Service

After you upgrade InfoScale Storage or InfoScale Enterprise, the Veritas InfoScale Messaging Service gets configured under a 'Local System account.'

You must re-configure the service under a domain user account having administrator privileges on all the cluster systems.

To reconfigure the user account for Veritas InfoScale Messaging Service

1. From Windows Computer Management or Windows Administrative Tools, access Services, and select Veritas InfoScale Messaging Service.
2. Right-click Veritas InfoScale Messaging Service and select **Properties** from the context menu.
3. On the Log On tab, select **This Account** and enter the domain user ID and password.

   The user account must have administrator privileges on all the cluster systems.
4. Confirm the password and click **Apply**, and then click **OK**.
5. In the Windows Services user interface, restart the Veritas InfoScale Messaging Service, for the changes to take effect.

Importing the configured rules

If you have exported the configured rules for event notification messages and actions, you must import them after the upgrade is complete.

To import the configured rules

1. From the VEA Control Panel perspective, select the server in the left pane.
2. Double-click Rule Manager in the right pane.
3. In the Rule Manager window, click **Import**.
4. Browse to the temporary location and select the XML file that you had saved.

Upgrading the dynamic disk group version

This task is applicable only if you have upgraded the product from a version prior to 6.1.

The dynamic disk group version is not upgraded automatically after you upgrade SFW or SFW HA. If a service group in the previous configuration contains a dynamic disk group, you must upgrade its version after you complete the SFW or SFW HA upgrade.

Please note that the following features and operations are introduced with the new disk group version and, therefore, will not be supported if the disk group version is not upgraded:

- The snapshot feature (prepare, snap shot, snap back, snap abort, and snap clear operations)
- The dirty region logging (DRL) feature (add a log and remove a log operations)

- Fast Resync and Fast File Resync features for volumes (add a log and remove a log operations)
- Import disk group (import a legacy disk group as a shared disk group operation)
- SSD caching

Note: When a legacy disk group with dirty region logging (DRL) logs in its volumes is upgraded, the DRL logs are deleted on upgrade. You need to add the DRL logs manually after the upgrade.

Note: Once a disk group version is upgraded, it cannot be changed back to an earlier disk group version and you can not import it on another server that is running any earlier versions of SFW.

Note: Microsoft Disk Management Disk Groups do not support upgrading the disk group version. You must first convert the disk group to an SFW dynamic disk group, and then upgrade the SFW dynamic disk group.

Note: After upgrading the disk group to the latest version and importing a disk group as a cluster-shared disk group (CSDG), you need to manually create the Volume Manager Shared Volume resource for each volume in a disk group.

For more information, see the *Storage Foundation Administrator's Guide*.

You can upgrade the dynamic disk group using the Veritas Enterprise Administrator (VEA) console or the CLI.

To upgrade the dynamic disk group version from the command line:

1. To upgrade the disk group using CLI, run the following command. You must run this command for each dynamic disk group separately.

   `vxdg -g`*DynamicDiskGroupName* [-T *version*] upgrade

   Where,

   DynamicDiskGroupName = Name of the dynamic disk group

   version = Target version of the dynamic disk group

2. After the disk group version is upgraded successfully, the disk group version should be 171. Run the following command to verify the upgraded version:

   `vxdg -g`*DynamicDiskGroupName* `dginfo`

To upgrade the dynamic disk group version using VEA

1. From the VEA console, right-click the disk group. You must perform these steps for each dynamic disk group separately.
2. Select **Upgrade Dynamic Disk Group Version**.

   A notification indicating that the disk group version has been upgraded is seen in the alert logs. After the disk group version is upgraded successfully, the disk group version should be 171. You can verify the disk group properties to confirm the upgraded version. After the disk group version is upgraded successfully, the disk group version should be 171.

Upgrading clusters for stronger security

If you have configured any secure clusters in your environment, upgrade them to use the 2048-bit key and SHA-256 signature certificates, which provide enhanced security.

Note: You do not need to perform this procedure if you have installed the Veritas InfoScale 7.0.1 patch and upgraded the existing clusters thereafter.

To upgrade a cluster using the wizard

1. Launch the Cluster Configuration Wizard from any node in the cluster that you want to upgrade.
2. Follow the wizard prompts to select the cluster for reconfiguration.
3. On the **Reconfigure Cluster Options** panel, select **Configure/Reconfigure Single Sign-on**.

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

To upgrade a cluster using the silent configuration utility

1. Back up your cluster configuration by creating a copy of the `main.cf` file.
2. Use the `VCWsilent` utility to perform the following operations sequentially:
   - Delete the cluster.
   - Create a new cluster.

   The syntax is as follows:

   ```
   VCWsilent XML_file_name_including_path
   ```

   For more information on the `VCWsilent` utility and the XML file formats to be used with it, see the *Cluster Server Administrator's Guide*.

3 Stop the cluster using the following command:

```
hastop -all -force
```

4 Manually restore your application configurations by copying only the relevant entries from the backed-up `main.cf` to newly created `main.cf`.

5 Start the cluster using the following command:

```
hastart
```

Reinstalling the custom agents

After performing the product upgrade, the installer does not upgrade the custom agents installed on the existing version of the product. You must re-install the custom agents, after the upgrade is complete. For more information, see the *Cluster Server Agent Developer's Guide*.

Including custom resources

The product installer does not upgrade custom resources. If a service group in the previous configuration contains custom resources, the wizard does not include the service group in the upgraded cluster.

To include a service group with custom resources in the upgraded cluster

1 Make sure that the agent binaries for the custom agent are available under `%VCS_HOME%\bin` where the variable *%VCS_HOME%* represents the VCS installation directory, typically `C:\Program Files\Veritas\cluster server`.

2 Stop the VCS engine (HAD) on all the nodes in the cluster.

From the command prompt, type:

```
C:\> hastop -all -force
```

3 During the installation, the installer copies previous configuration files to a backup location. Locate the backed up types.cf and main.cf files: `C:\Documents and Settings\All Users\Application Data\Veritas\cluster server\vpibackup`.

4 Copy the resource type definition for the custom resource from the backed up types.cf and add it to the types.cf file for the VCS cluster.

5 If resources for a custom resource type are dependent on resources for agents bundled with VCS, you must update the resource definition of the VCS bundled agent to include the new attributes or remove the deprecated attributes.

For information on the attribute values and descriptions, see the *Cluster Server Bundled Agents Reference Guide*.

6 Verify the configuration.

From the command prompt, type:

```
C:\> hacf -verify config_directory
```

The variable *config_directory* refers to the path of the directory containing the main.cf and types.cf.

7 Start the VCS engine (HAD) on the node where you changed the configuration. Type the following at the command prompt:

```
C:\> hastart
```

8 Start the VCS engine (HAD) on all the other cluster nodes.

Chapter 5

Administering the InfoScale product installation

This chapter includes the following topics:

- Adding or removing product options
- Managing InfoScale licenses
- Repairing an InfoScale product installation
- About reinstalling InfoScale products

Adding or removing product options

After you have installed the InfoScale products, you may want to add or remove the product options. The product installer wizard lets you to add or remove the installed options and the manage the product licenses.

This section provides details about adding or removing the product options. For details about managing product licenses, refer to:

See "Managing InfoScale licenses" on page 79.

Note the following points before you begin to add or remove the product options:

- You cannot add or remove the product options on a system that runs Server Core operating system. To add or remove the product options on these systems you must uninstall the product and then install it again.
- You can add or remove the product options only on the local system.

- You can add or remove the product options only if you have installed the server components.

If you are adding the DSMs to a deployment setup that involves Windows Server Failover Cluster or a VCS cluster, ensure that you move the resources to another node or take the resource offline. Then, install the required hardware drivers and perform the following steps:

To add or remove features

1. Open the Windows Control Panel and click **Programs and Features**.
2. Select the InfoScale product entry and click **Change**.
3. On the Mode Selection panel, select **Add or Remove** and then click **Next**.
4. On the System Selection panel, the wizard performs the verification checks and displays the available product options. To add or remove the options, select or clear the corresponding check boxes and then click **Next**.

   Note that the wizard proceeds only if the system passes the validation checks. In case the verification checks have failed, review the details and rectify the issue. Before you choose to proceed with the installation, click **Re-verify** to re-initiate the verification checks.
5. On the Pre-install Summary panel, review the summary and click **Next**.

   Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This will reboot all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.
6. On the Installation panel, review the progress of installation and click **Next** after the installation is complete.

   If an installation is not successful, the status screen shows a failed installation. Refer to the Post-install summary for more details. Rectify the issue and then proceed to re-install the component.

7 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed, refer to the log file for details.

8 On the Finish panel, click **Finish**.

If you had chosen to initiate the auto reboot, a confirmation message to reboot the local system appears. Click **Yes** to reboot immediately or **No** to reboot later.

In case you had not selected to initiate the auto reboot, ensure that you manually reboot these systems.

For adding the DSMs, if you had disconnected all but one path, you must reconnect the additional physical path now.

You can now proceed to configure the service groups for the newly added options.

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

Managing InfoScale licenses

After you have installed an InfoScale product, you may need to manage the product licenses to modify your license type.

You can manage your licenses by any of the following methods:

- Using the Windows Add or Remove programs
- Using the VMware vSphere Client

Note: This method any can be used only if you have installed InfoScale Availability or InfoScale Enterprise in a VMware environment, and have installed the Symantec High Availability Console to administer the virtual machines.

Note the following points before you begin to manage the licenses:

- You cannot manage licenses on a system that runs Server Core operating system. To manage licenses on these systems, you must uninstall the product and then install it again using the new licenses.
- You can manage the licenses only on the local system.
- You can manage the licenses only if you have installed the server components.

To manage licenses using the Windows Add or Remove programs

1 Open the Windows Control Panel and click **Programs and Features**.

2 Select the InfoScale product entry and click **Change**.

3 On the Mode Selection panel, select **Add or Remove** and then click **Next**.

4 On the System Selection panel, the wizard performs the verification checks and displays the applicable installation and product options. In case the verification checks have failed, review the details and rectify the issue. Before you choose to proceed with the installation click **Re-verify** to re-initiate the verification checks.

Note that the wizard proceeds only if the system passes the validation checks.

To manage the licenses, perform any of the following applicable task:

- To change the license type, select the required license type from the **License key** drop-down list.
  If you change your license type to "User entered license key", the License Details panel appears by default. Proceed through step 5 to add the license keys.
- To add or remove the licenses, click **Edit**.

5 On the License Details panel, enter the license key and then click **Add**.

The wizard validates the entered license key and displays the relevant error if the validation fails.

6 On the License Details panel, click **OK**.

The wizard displays the applicable installation and product options on the System Selection panel.

7 On the System Selection panel, select or clear the required product options and then click **Next**.

The wizard performs the verification checks and proceeds to the Pre-install Summary panel. In case the verification checks have failed, review the details and rectify the issue. Before you choose to proceed with the installation click **Re-verify** to re-initiate the verification checks.

Note that the wizard proceeds only if the verification checks are passed.

8 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This will reboot all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

9 On the Installation panel, review the progress of installation and click **Next** after the installation is complete.

If an installation is not successful, the status screen shows a failed installation. Refer to the Post-install summary for more details. Rectify the issue and then proceed to re-install the component.

10 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed, refer to the log file for details.

11 On the Finish panel, click **Finish**.

If you had chosen to initiate the auto reboot, a confirmation message to reboot the local system appears. Click **Yes** to reboot immediately or **No** to reboot later.

In case you had not selected to initiate the auto reboot, ensure that you manually reboot these systems.

Notes:

- If you make any changes to the InfoScale Foundation, InfoScale Storage, or the InfoScale Enterprise licenses, the changes take effect when the vxsvc service starts again. If you remove all the licenses, the vxsvc service fails to start and the service recovery options are changed to “Take No Action”. To start the service you must enter the licenses and then manually start the service and change the service recovery option to “Restart the Service”.
- If you make any changes to the InfoScale Availability licenses, you must restart Veritas High Availability Engine (HAD) service for the changes to take effect. If you remove all the licenses, HAD service fails to start. To start the HAD service, you must enter the required licenses and then run the `hastart` command or manually start the Veritas High Availability Engine service.

Note: The following method any can be used only if you have installed InfoScale Availability or InfoScale Enterprise in a VMware environment, and have installed the Symantec High Availability Console to administer the virtual machines.

To manage the licenses through the vSphere Client menu

1 Connect to the vCenter Server and navigate to **Home** > **Solutions and Applications** > **Symantec High Availability**

2 Click the **License Management** tab.

3 Select the desired virtual machines and click **Next**.

Only the machines that are running can be selected.

4 Select a virtual machine and click **Add License**.

Use the **CTRL** key to select multiple virtual machines.

5 On the Add License panel, enter the license key in the **New License Key** text box and click **Validate Key**.

The installer validates the license key. For successful validation the status reflects **New license applied**. In case of validation failure, the status is reflects **Validation Failed. Enter a valid license key**. Click **Ok**.

6 Click **Apply**.

The specified license keys take effect immediately.

To view the details of the existing licenses, select the individual virtual machine. The details are displayed in the **Existing License Details** table.

Repairing an InfoScale product installation

The product installer can repair an existing installation of the InfoScale products.

The **Repair** option restores the installation to its original state. This option fixes missing or corrupt files, shortcuts, and registry entries on the local system.

You can repair the installation only on the local system.

Note: Before you proceed to repair the installation, you must save your configuration to another system and fail over the service groups for your applications to another node.

To repair the installation

1 Open the Windows Control Panel and click **Programs and Features**.

2 Select the InfoScale product entry and click **Change**.

3 On the Mode Selection panel, select **Repair**. Click **Next**.

4 On the System Selection panel, installer performs the verification checks. Click **Next** once the status is "Ready for repair".

In case the verification checks have failed, review the details and rectify the issue. Before you choose to proceed with the installation, click **Re-verify** to re-initiate the verification checks.

Note: You cannot select the installation and product options.

---

5 On the Pre-install Summary panel, review the information and click **Next** to begin the repair process.

Note that if you are repairing the server installation, the **Automatically reboot systems after installer completes operation** check box is selected by default. This will reboot the system immediately after the repair operation is complete. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

6 On the Installation panel, review the list of services and processes running on the systems. Select a system to view the services and processes running on it.

The wizard stops the product-specific services and discovers the processes running, if any, on the systems. These processes need to be stopped to proceed with the operation. Click **Next** to forcefully stop the processes and proceed with the operation. Alternatively, you can manually stop the processes.

If the services or processes cannot be stopped, the operation fails. Rectify the error and then click **Retry** to validate the affected system again. Click **Retry All** to validate all the systems again.

7 On the Post-install Summary panel, review the summary and click **Next**.

8 On the Finish panel, click **Finish**.

In case you had not selected to initiate the auto reboot, ensure that you manually reboot the node.

About reinstalling InfoScale products

If your product installation has failed due to some reason, you can choose to reinstall it without uninstalling the components that were installed during the failed attempt.

Note: You must reboot your system before you begin to reinstall the product.

To reinstall the product, rectify the cause of failure and then proceed with the installation.

See “Installing the server components using the installation wizard” on page 24.

If you choose to install the product using the product installer wizard, during the installation a confirmation message is displayed on the System Selection panel. Click **Yes** to proceed with the installation.

Chapter 6

Uninstalling the InfoScale products

This chapter includes the following topics:

- About uninstalling the InfoScale products
- Uninstalling the InfoScale products using the installation wizard

About uninstalling the InfoScale products

You can completely uninstall the product or uninstall the product options using the product installation wizard. To launch the product installer choose the Windows Add or Remove Programs feature.

If you have deployed a cluster configuration, you must unconfigure it before you begin to uninstall the product.

Also, if you are running NetBackup, you must stop the Symantec Private Branch Exchange (PBX) service.

Uninstalling the InfoScale products using the installation wizard

The Veritas product installer enables you to uninstall the product. You can simultaneously uninstall the product from multiple remote nodes. To uninstall the product from remote nodes, ensure that the product is installed on the local node.

Uninstalling the Server components, uninstalls the client components and the applicable high availability, replication and the database agents, if any.

To uninstall using the product installer

1. In the Windows Control Panel, select **Programs and Features**.
2. Select the InfoScale product entry and click **Uninstall**.
3. Review the information on the Welcome panel and then click **Next**.
4. On the System Selection panel, add the systems from which you want to uninstall the product.

   Note: By default the local system is selected for uninstallation. In case you are performing a remote uninstallation and do not want to uninstall the software from the local system, you must remove the system from the list.

   You can add the systems in one of the following ways:

   - In the System Name or IP text box, manually type the system name and click **Add**.

     Note: The wizard does not support the Internet Protocol version 6. To add the systems having Internet Protocol version 6, you must type the system name.

   - Alternatively, browse to select the systems.
     The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**. Once you add or select a system, wizard performs the verification checks, and notes the verification details.
5. Click **Next**.

   Note that the wizard fails to proceed with the uninstallation, unless all the selected systems have passed the verification checks and are ready for uninstallation. In case the verification checks have failed on any of the system, review the details and rectify the issue. Before you choose to proceed with the uninstallation click **Re-verify** to re-initiate the verification checks for this node.

6 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This selection reboots the remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

7 On the Uninstallation panel, review the list of services and processes running on the systems. Select a system to view the services and processes running on it.

The wizard stops the product-specific services and discovers the processes running, if any, on the systems. These processes need to be stopped to proceed with the operation. Click **Next** to forcefully stop the processes and proceed with the operation. Alternatively, you can manually stop the processes.

If the services or processes cannot be stopped, the operation fails. Rectify the error and then click **Retry** to validate the affected system again. Click **Retry All** to validate all the systems again.

8 On the Post-uninstall Summary panel, review the uninstallation results and click **Next**.

If the uninstallation has failed on any of the system, review its summary report and check the log file for details.

9 On the Finish panel, click **Finish**.

In case you had not selected to initiate the auto reboot for the remote systems, ensure that you manually reboot these systems.

Chapter 7

Performing application upgrades in an InfoScale environment

This chapter includes the following topics:

- Upgrading Microsoft SQL Server
- Upgrading Oracle
- Upgrading application service packs in an InfoScale environment

Upgrading Microsoft SQL Server

The following table lists the SQL Server upgrade scenarios that are supported in an InfoScale environment and the sections that provide the details:

Upgrade scenario	Refer to
■ Upgrading SQL Server 2008 or 2008 R2 to SQL Server 2012 ■ Upgrading SQL Server 2012 to SQL Server 2014	See "Upgrading to later versions of SQL Server" on page 88.
Upgrading SQL Server with its compatible service pack	See "Upgrading the SQL Server service packs" on page 98.

Note: If you plan to upgrade an InfoScale product and also your applications, you must first upgrade the InfoScale product and then upgrade the application.

See "Preparing the systems for an upgrade" on page 39.

Upgrading to later versions of SQL Server

This section describes how to upgrade the following application versions in an InfoScale environment:

- SQL Server 2008 or 2008 R2 to SQL Server 2012
- SQL Server 2012 to SQL Server 2014

Perform the following tasks before you begin the upgrade:

- Take a backup of the SQL databases.
- In case of a disaster recovery (DR) configuration:
  - Ensure that the databases on the primary site and the secondary site are synchronized and then stop the replication between the sites.
  - Make a note of the SQL virtual server name and all the IP addresses configured at the primary site and the secondary site. These details are needed later.

At a high level, the upgrade process involves the following tasks:

1. Upgrade SQL Server on the first cluster node.
2. Upgrade SQL Server on each additional failover node.
3. In a DR configuration, repeat the upgrade procedures on the nodes at the secondary site. First upgrade SQL Server on the first cluster node at the DR site, and then upgrade the application on the additional failover nodes.
4. Delete the existing SQL Server service group, including the service group at the DR site, if applicable.
5. Create a service group for the upgraded SQL Server version, using the SQL Server Configuration Wizard. In case of a DR setup, create a service group at the secondary site too.

Note: In a DR configuration, first upgrade SQL Server on the cluster nodes at the primary site and then proceed with the nodes at the secondary site. You must follow the same upgrade sequence at both sites—upgrade the first node and then the additional nodes—as described in this section.

Upgrading SQL Server on the first cluster node

This procedure assumes that a single SQL Server instance is configured in a multi-node cluster.

To upgrade SQL Server on the first cluster node

1. On the node where the application service group is online, take only the SQL Server resources offline and delete them. Leave the storage resources online.
2. If the resources are already offline, bring the storage resources online.
3. Take a backup of the existing SQL Server databases from the shared disk and store them in a temporary location.

   The backed-up directories are needed later, while you upgrade SQL Server on the additional failover nodes.
4. Launch the SQL Server installer for the appropriate version to install the application on the node. Select the option to upgrade the existing SQL Server instances when prompted to do so. The installer then automatically places the SQL Server data files in the appropriate location.

   Refer to the Microsoft SQL Server documentation for instructions.
5. Take the entire service group offline on the node.

   The upgrade steps on the first cluster node are now complete. Proceed to upgrading SQL Server on the additional failover nodes.

Upgrading SQL Server on additional failover nodes

Perform the following steps on each additional failover node that is part of the SQL Server service group.

To upgrade SQL Server on the additional nodes

1. Bring the storage resources online.
2. Rename the existing SQL Server data directories on the shared disks. These directories are updated when SQL Server is installed on the first node. You may also delete these directories.
3. Copy the backed-up SQL Server data directories from the temporary location to the shared disks.

   These directories are the same ones that you had backed up earlier while upgrading SQL Server on the first cluster node.

4 Launch the Microsoft SQL Server installer for the appropriate version to install the application on the node. Select the option to upgrade the existing SQL Server instances, when prompted to do so. The installer then automatically places the SQL Server data files at the appropriate location.

   Refer to the Microsoft SQL Server documentation for instructions.

5 Take the entire service group offline on the node.

Note: If there are no additional nodes for upgrade, you need not take the service group offline.

The upgrade steps on the additional cluster node are now complete. Delete the existing SQL Server service group and proceed to create the service group for the upgraded SQL Server version in the cluster.

Creating the new SQL Server service group

To configure a service group for the upgraded SQL Server version, run the SQL Server Configuration Wizard on the last upgraded node.

For details, refer to the application-specific implementation guide or solutions guide.

Note: In a disaster recovery (DR) configuration, repeat these steps on the first cluster node at the secondary site and then reconfigure the DR components.

To create the service group for the upgraded SQL Server version

1 Rename the RegRep directory, if present, on the shared disk.

2 Create the service group using the SQL Server Configuration Wizard.

   For details, refer to the application-specific implementation guide or solutions guide.

3 After creating the SQL Server service group, verify the configuration by switching the service group to another node in the cluster.

4 Delete the RegRep directory that you renamed in the first step.

Upgrading Oracle

This section describes the tasks necessary to upgrade Oracle in an InfoScale environment.

Note: If you plan to upgrade your applications while you upgrade an InfoScale product, you must upgrade the InfoScale product before you begin to upgrade the application.

See "Preparing the systems for an upgrade" on page 39.

Upgrading Oracle involves the following steps:

- Upgrading the Oracle binaries
- Bringing the Oracle service group online
- Stopping HAD using the `hastop -local -force` command
- Upgrading the Oracle database
- Performing the post upgrade tasks

For information about supported Oracle upgrade paths and the details about upgrading Oracle binaries, refer to the Oracle product documentation.

Performing the post upgrade tasks

Perform the following tasks to configure Oracle in an InfoScale environment:

- Associate the updated database with the listener for Oracle 10g and 11g.
  See "Associating the updated Oracle database with the listener" on page 91.
- Configure the database and listener to use the virtual IP address.
  See "Configuring the Oracle database and listener to use the virtual IP address" on page 92.
- Configure Oracle and listener services.
  See "Configuring Oracle and listener services" on page 95.
- Modify the ServiceName attribute for the Netlsnr resource.
  See "Modifying the ServiceName attribute for the netlsnr resource" on page 96.

Associating the updated Oracle database with the listener

To associate the database with the listener

1. Ensure that the initialization parameter file contains the following entries:
   - SERVICE_NAMES (the name of the database service)
   - INSTANCE_NAME (the name of the database instance)

These parameters are created during installation or database creation.

2 Use one of the following procedures to configure the new attribute listener_alias:

Run the following SQL command:

```
SQL> ALTER SYSTEM SET LOCAL_LISTENER='<listener_alias>' scope=spfile;
```

OR

Add the following entry to the initialization parameter file (pfile or spfile):

```
LOCAL_LISTENER = <listener_alias>
```

3 Define the parameter listener_alias. If your Oracle configuration uses the file `tnsnames.ora`, edit the file as instructed below. The default location of `tnsnames.ora` is `%ORACLE_HOME%\NETWORK\ADMIN`.

Add the following to `tnsnames.ora` file:

```
<listener_alias>=
(DESCRIPTION =
(ADDRESS=(Protocol=TCP)(HOST=virtual_IP_address)(Port=1521))
)
```

4 Stop and restart the database.

The listener_alias parameter gets appended by the default domain name that is specified in the file `sqlnet.ora`.

Configuring the Oracle database and listener to use the virtual IP address

All databases and listeners configured must use the same virtual IP. Update the Oracle files to set the virtual IP address.

Setting the virtual IP address involves the following tasks:

- Creating a virtual IP address.
- Verifying the initialization file settings.
- Updating the Oracle configuration files.

Use the following procedures to configure the Oracle database and listener.

To create a virtual IP address

1. Open the **Network Connections**.
2. Right-click the public network connection to be used and click **Properties**.
3. Select **Internet Protocol (TCP/IP)** and click **Properties**.
4. Click **Advanced**.
5. In the **IP Settings** tab, click **Add** to add a virtual IP address and subnet mask.

To verify the initialization file settings, if a PFILE is used

1. Open the Registry Editor.

   From the **Start** menu, choose **Run**. In the **Open** field, enter `regedit` and click **OK**.
2. Double-click the ORA_*SID*_PFILE registry key at HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\HOME_*ID*\.

   The variable *SID* represents the database instance.
3. Verify that the Value data field specifies the correct path at which the initialization file, `init.ora`, is located.

To verify the initialization file settings, if an SPFILE is used

1. Run sqlplus.exe.
2. Connect to the database.
3. Verify the following query returns the correct path of the SPFILE.

```
select value from v$parameter where name = 'spfile'
```

To update the Oracle configuration files

1. In the `listener.ora` and `tnsnames.ora` files, change the host name for all the TCP protocol address databases to the virtual IP address that you created.

   Replace

   `HOSTNAME=`*`machine_name`*

   with

   `HOSTNAME=`*`virtual_IP_address`*

2 In the initialization file, change the dispatchers parameter.

Oracle requires an initialization file, a PFILE or an SPFILE, to start database instances. Choose the appropriate reference depending on the initialization file you use.

See "Setting the dispatchers parameter in PFILE" on page 94.

See "Setting the dispatchers parameter in SPFILE" on page 94.

3 Restart the Oracle and listener services.

Setting the dispatchers parameter in PFILE

In the PFILE, set the host name for all TCP protocol address dispatchers to the virtual IP address that you created.

Edit the dispatchers parameter only for the host name and leave the rest of the configuration as it is. Set the value as:

```
dispatchers = '(ADDRESS = (Protocol=TCP)
(HOST=virtual_IP_address)
(any other previously existing entry))'
```

The variable *virtual_IP_address* represents the virtual IP address that you created.

For example:

```
dispatchers = '(ADDRESS = (Protocol=TCP)(HOST=10.210.100.110)
(SERVICE=Data1XDB)'
```

Setting the dispatchers parameter in SPFILE

Use the following steps to set the dispatchers parameter in SPFILE.

To set the dispatchers parameter in SPFILE

1 Convert the SPFILE to PFILE.

2 Modify the PFILE.

See "Setting the dispatchers parameter in PFILE" on page 94.

3 Convert the PFILE to SPFILE.

4 Save the SPFILE to the original location on the shared disk.

Refer to the Oracle documentation for specific information on converting a PFILE or an SPFILE.

Configuring Oracle and listener services

Configuring the Oracle and Listener services involves the following tasks:

- Making the Oracle and Netlsnr services manual.
- Configuring log on properties for Oracle services.

Use the following procedures to configure Oracle and listener services.

To make services manual

1. Open the **Services** applet from **Start > All Programs > Administrative Tools > Services** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
2. Double-click the Oracle service. In the SCM, the following appears:
   - Oracle services appear as OracleService*SID*, where *SID* represents the database instance.
   - Listener services appear as Oracle*Ora_Home*TNS*ListenerName*, where *Ora_Home* represents the Oracle home directory and *ListenerName* is the name of the listener set during the installation.
3. In the **Properties** window, click the **General** tab.
4. From the **Startup Type** drop-down list, select **Manual**.
5. Click **OK**.

To configure the log on properties for oracle services

1. Open the **Services** applet from **Start > All Programs > Administrative Tools > Services** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
2. Double-click the Oracle service. In the SCM, the names of Oracle services appear as OracleService*SID*, where *SID* represents the database instance.
3. In the **General** tab of the **Properties** window, click **Stop** to stop the service.
4. Click the **Log On** tab.
5. Choose **This Account**.
6. Enter the credentials of the user in whose context Oracle was installed.
7. Click the **General** tab and click **Start** to start the service with the new Log On properties. Click **OK**.

Modifying the ServiceName attribute for the netlsnr resource

Perform the following steps to modify the ServiceName attribute for the Netlsnr resource.

To modify the ServiceName attribute

1. Start the high availability daemon (HAD).

   Type the following on the command prompt:

   ```
   hastart
   ```

2. Take the Netlsnr resource offline.

   Type the following on the command prompt:

   ```
   hares -offline resource_name -sys system_name
   ```

3. Modify the ServiceName attribute for the Netlsnr resource.

   Type the following on the command prompt:

   ```
   hares -modify resource_name attribute_name attribute_value
   ```

   For example, to modify the ServiceName attribute of the Netlsnr resource, Netlsnr_res, type:

   ```
   hares -modify Netlsnr_res ServiceName  attribute_value
   ```

   where, *attribute_value* is the name of the listener service in Oracle 9i or 10g versions.

4. Bring the Netlsnr resource online.

   Type the following on the command prompt:

   ```
   hares -online resource_name -sys system_name
   ```

Upgrading application service packs in an InfoScale environment

This section describes the tasks to be performed if you plan to upgrade your application to its compatible service pack in an InfoScale environment.

See “Upgrading the Exchange Server service packs” on page 97.

See “Upgrading the SQL Server service packs” on page 98.

See “Upgrading SharePoint Server 2010 to a service pack” on page 104.

Note: If you plan to upgrade your applications while you upgrade your InfoScale product, you must upgrade the InfoScale product before you begin to upgrade the application.

Upgrading the Exchange Server service packs

This section describes how to upgrade Microsoft Exchange servers to their corresponding service packs. The outlined procedure is applicable only if you already have your Exchange setup in an InfoScale environment.

Exchange Server service pack upgrade scenarios	Refer to
Microsoft Exchange 2010 to 2010 SP1, SP2 or SP3	See "Upgrading Exchange Server 2010 to a service pack" on page 97.

Upgrading Exchange Server 2010 to a service pack

This section describes how to upgrade Exchange Server 2010 to any 2010 service pack. It is applicable only if you already have Exchange 2010 setup in an InfoScale environment.

Before you proceed to upgrade to a service pack, ensure that you have met all the necessary pre-requisites for installing the service pack on all the cluster nodes where you are upgrading Exchange.

For details refer to Microsoft documentation.

To upgrade Exchange 2010 to a service pack

1. Bring the Exchange service group online.

   Type the following on the command prompt on a cluster node:

   ```
   hagrp -online service_group -sys system
   ```

2. Stop HAD on all the cluster nodes where you want to upgrade the Exchange installation. At the command prompt, type:

   ```
   hastop -local -force
   ```

3. Launch the Exchange 2010 service pack installer and install the service pack.

   You can install the service pack parallel on all the nodes, where you are upgrading Exchange. In case of disaster recovery, you can simultaneously upgrade both the sites.

Upgrading the SQL Server service packs

This section describes how to upgrade Microsoft SQL Server to its corresponding service packs. The outlined procedure is applicable only if you already have your SQL Server setup in an InfoScale environment.

SQL Server service pack upgrade scenarios	Refer to
Microsoft SQL Server 2012 or SQL Server 2014 to their latest service packs	See "Upgrading SQL Server 2012 or SQL Server 2014 to a service pack " on page 98.

Upgrading SQL Server 2012 or SQL Server 2014 to a service pack

This section describes the tasks to upgrade SQL Server 2012 or SQL Server 2014 in a InfoScale disaster recovery environment.

Before upgrading SQL Server

Consider the following points before you proceed with the upgrade:

- Ensure that you have installed and configured SQL Server.
- Ensure that the logged on user has administrative privileges to the SQL instance that you want to upgrade.
- Ensure that you have taken a recent backup of your system, user databases, and the SQL Server directories, from the shared storage.
- Refer to the Microsoft documentation for prerequisites related to SQL Server Service Pack installation.

Upgrading SQL Server

Consider a three-node disaster recovery cluster setup; Node A, Node B and Node C. Node A and Node B are on the primary site and Node C is on the secondary site. The SQL service group is online on Node A.

You can upgrade SQL Server 2012 or SQL Server 2014 to a service pack in any of the following ways:

- Upgrading SQL Server to a service pack one node at a time
- Upgrading SQL Server to a service pack on passive nodes and then on the active node

Upgrading SQL Server to a service pack one node at a time

In this procedure you upgrade SQL Server on the nodes at the primary site first and then on the nodes at the secondary site. This process involves moderate service group downtime, because you upgrade one cluster node at a time.

To upgrade SQL Server to a service pack, perform the following steps:

1. Stop the replication between the primary and the secondary site.

   If using Volume Replicator for replication, from the VEA Console, right-click the Secondary RVG and select **Stop Replication** from the menu that appears.

2. On Node A where the SQL Server service group is online, take the SQLServer, MSOlap, and SQLServer-Agent resources offline.

   Run the following command from the command prompt:

   ```
   hares -offline [-parentprop] resource -sys system
   ```

   Here, *resource* is the name of the SQL resource and *system* is the name of node where the SQL Server service group is online.

3. From Services.msc, ensure that all the SQL services and the SQL services for which VCS resources are configured are stopped.

4. Perform the following steps on the SQL Server service group on Node A (active node):

   - Bring the RegRep resource offline.
     Type the following on the command prompt:
     ```
     hares -offline [-parentprop] resource -sys system
     ```
   - Disable the RegRep resource.
     Type the following on the command prompt:
     ```
     hares -modify resource_name Enabled 0
     ```
   - Except the storage resources (MountV and VMDg), take all the resources offline.
   - Freeze the service group.
     Type the following on the command prompt:
     ```
     hagrp -freeze service_group [-persistent]
     ```

5. Install the Microsoft SQL Server Service Pack on Node A.

6. If a FILESTREAM resource is configured in the SQL Server service group, verify if a FILESTREAM share exists on the node and then delete it.

   Run the following commands from the command prompt:

   - `net share`

This command lists all the shares on the node.

- `net share share_name /delete`
  Here, *share_name* is the name of the FILESTREAM share.

7 Unfreeze the SQL Server service group.

Type the following on the command prompt:

```
hagrp -unfreeze service_group [-persistent]
```

8 Fail over the service group to Node B and perform the following steps on Node B, in the given order:

- Except the storage resources (MountV and VMDg), take all the resources offline.
- From Services.msc, ensure that all the SQL services and the SQL services for which VCS resources are configured are stopped.
- Freeze the service group.

9 Rename the SQL folders on the shared storage and copy the backed up SQL Server directories to the shared storage.

The SQL Server data files available on the shared storage are upgraded during the SQL upgrade on Node A. Before you begin to upgrade SQL on Node B, you must rename the folders containing the upgraded SQL data files and restore the initially backed up SQL Server directories. If you do not restore the initially backed up SQL Server directories, then the SQL upgrade on Node B may fail indicating that the SQL Server data files are already upgraded.

10 Install the Microsoft SQL Server Service Pack on Node B.

11 If a FILESTREAM resource is configured in the SQL Server service group, verify if a FILESTREAM share exists on the node and then delete it.

Run the following commands from the command prompt:

- `net share`
- `net share share_name /delete`
  Here, *share_name* is the name of the FILESTREAM share.

12 Unfreeze the service group on Node B and enable the RegRep resource.

Run the following commands from the command prompt:

- `hagrp -unfreeze service_group [-persistent]`
- `hares -modify resource_name Enabled 1`

13 Bring the service group online on Node B.

14 Start the replication between the primary and the secondary site.

15 Switch the service group to a node on the DR site (Node C).

Type the following on the command prompt:

```
hagrp -switch service_group -site site_name
```

16 Stop the replication between the primary and the secondary site again.

17 Perform the following steps on Node C, in the given order:

- Except the storage resources (MountV and VMDg), take all the resources offline.
- Disable the RegRep resource and freeze the service group.
- Rename the SQL folders on the shared storage and copy the backed up directories to the shared storage.
- Install the Microsoft SQL Server Service Pack on Node C.
- If a FILESTREAM resource is configured in the SQL Server service group, verify if a FILESTREAM share exists on the node and then delete it using the following commands:

  ```
  net share
  net share share_name /delete
  ```

- Unfreeze the service group and enable the RegRep resource.

18 Start replication between the primary and secondary site.

19 Switch the service group back to Node B (last upgraded node) on the primary site.

Note: You must bring the SQL service group online on Node B first. This is because the replication service group is online on Node B. You can then switch the SQL service group on any node on the primary site.

Upgrading SQL Server to a service pack on passive nodes and then on the active node

In this procedure you upgrade SQL Server on the passive nodes at both the sites first and then on the active node. This process involves less complexity and service group downtime.

Note: This procedure is applicable only if the SQL Server database and the SQL Server analysis (OLAP) files are installed on shared storage on the active node and at the default location (on the C: drive) on the passive nodes.

To upgrade SQL Server to a service pack

1. Freeze the SQL Server service group on Node A using the VCS Cluster Manager (Java Console).

   On the Service Groups tab, right-click the service group and then click **Freeze** > **Persistent**. Save the configuration and leave the group frozen until all the nodes are updated.

   Alternatively, run the following command on the command prompt:

   `hagrp -freeze service_group [-persistent]`

2. Back up the registry of the SQL Server database instance on the passive nodes, by following these steps sequentially:

   - Access the registry by using your preferred method to execute `regedit.exe`.
   - Locate the SQL Server instance registry key:
     `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\MSSQL11.instance`
   - Right-click on the registry key and click **Export**.
   - Save the registry key.

   Note: While upgrading SQL Server 2014 to a service pack, the *MSSQL11* value in the registry key changes to `MSSQL12` .

3. Stop the replication between the primary and the secondary sites.

   If VVR is used for replication, from the VEA Console right-click the Secondary RVG and select **Stop Replication** from the context menu.

4. If you find any registry keys from the following list on nodes B and C, update them to point to the root directory. Note that *instance* represents the SQL Server instance that you are trying to upgrade.

   - `HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.instance\MSSQLServer\BackupDirectory`
   - `HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.instance\MSSQLServer\DefaultLog`
   - `HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.instance\MSSQLServer\Parameters\SQLArg0`
   - `HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.instance\MSSQLServer\Parameters\SQLArg1`

- HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.*instance*\MSSQLServer\Parameters\SQLArg2
- HKLM\SOFTWARE\Microsoft\Microsoft SQL Server\MSAS11.*instance*\CPE\ErrorDumpDir
- HKLM\SOFTWARE\Microsoft\Microsoft SQL Server\MSAS11.*instance*\Setup\DataDir
- HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.*instance*\MSSQLServer\DefaultData
- HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.*instance*\SQLServerAgent\ErrorLogFile
- HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.*instance*\SQLServerAgent\WorkingDirectory
- HKLM\Software\Microsoft\Microsoft SQL Server\MSSQL11.*instance*\Setup\SQLDataRoot

For example, if the key `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\MSSQL11.instance\MSSQLServer\BackupDirectory` contains the value `F:\Test\MSSQL11.MSSQLSERVER\MSSQL\Backup`, change this value to point to the root directory `C:\Program Files\Microsoft SQL Server`.

Therefore, updated value becomes `C:\Program Files\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\Backup`.

5 Complete the SQL Server service pack upgrade on nodes B and C, and reboot the server.

Do not fail over until the service pack installation on Node A is complete.

6 On Node A, run the command `hastop -local -force` to stop the VCS services.

Doing so stops VCS, but leaves the storage and the SQL Server online. It also prevents VCS from taking any action on the resources when the installer restarts the service.

7 Complete the SQL Server service pack upgrade on Node A, and verify that the services are back in the Started state.

8 Run command `hastart` to start VCS services on Node A.

9 Start the replication between the primary and secondary sites.

10 Unfreeze the service group and test failover to nodes B and C.

Run the following command on the command prompt:

```
hagrp -unfreeze service_group [-persistent]
```

This completes the SQL Server upgrade to a service pack.

Upgrading SharePoint Server 2010 to a service pack

This section describes how to upgrade Microsoft SharePoint Server to SharePoint Server 2010 SP1 or SharePoint Server 2010 SP2. The outlined procedures are applicable only if you already have your SharePoint Server setup in an InfoScale environment.

Before you proceed to upgrade SharePoint Server, ensure that you have met all the necessary pre-requisites for installing SharePoint Server 2010 SP1 or SharePoint Server 2010 SP2 on all the cluster nodes where you are upgrading SharePoint Server.

For details, refer to Microsoft documentation.

To upgrade SharePoint Server 2010 to SharePoint Server 2010 SP1 or SharePoint Server 2010 SP2

1 Bring all the service groups corresponding to SharePoint application online.

2 Stop HAD on all the cluster nodes where you want to upgrade the SharePoint Server installation. At the command prompt, type:

```
hastop -all -force
```

3 Upgrade SharePoint Server 2010 to SharePoint Server 2010 SP1 or SharePoint Server 2010 SP2.

For details, refer to Microsoft documentation.

4 Start HAD. At the command prompt, type:

```
hastart -all
```

Appendix A

Services and ports

This appendix includes the following topics:

- InfoScale ports and services

InfoScale ports and services

If you have configured a firewall, then ensure that the firewall settings allow access to the services and ports used by the InfoScale products.

The following table displays the services and ports used by InfoScale products.

Ensure that you enable the ports and services for both, inbound and outbound communication.

Note: The port numbers marked with an asterisk are mandatory for configuring the InfoScale products.

Table A-1 InfoScale services and ports

Component Name/Port	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
vxsvc.exe 2148*, 3207/TCP/UDP Veritas Enterprise Administrator (VEA) Server	✓	X	✓	✓
CmdServer.exe 14150*/TCP Veritas Command Server	X	✓	✓	X

Table A-1 InfoScale services and ports *(continued)*

Component Name/Port	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
had.exe 14141*/TCP Veritas High Availability Engine	X	✓	X	✓
Plugin_Host.exe 7419*/TCP Veritas Plugin Host Service	X	✓	✓	✓
vcsauthserver.exe 14149/TCP/UDP VCS Authentication Service	X	✓	X	✓
vras.dll 8199/TCP Volume Replicator Administrative Service	X	X	✓	✓
vxrserver.exe 8989/TCP Volume Replicator Resync Utility	X	X	✓	✓
vxio.sys 4145/UDP Volume Replicator Connection Server	✓	X	✓	✓
VxSchedService.exe 4888/TCP Veritas Scheduler Service Use to launch the configured schedule.	✓	X	✓	✓
User configurable ports created at kernel level by vxio .sys file 49152-65535/TCP/UDP Volume Replicator Packets	✓	X	✓	✓

Table A-1 InfoScale services and ports *(continued)*

Component Name/Port	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
Notifier.exe 14144/TCP/UDP VCS Notification	X	✓	X	✓
wac.exe 14155/TCP/UDP VCS Global Cluster Option (GCO)	X	✓	X	✓
xprtld.exe 5634/HTTPS Veritas Storage Foundation Messaging Service	✓	X	✓	✓

Appendix B

Migrating from a third-party multi-pathing solution to DMP

This appendix includes the following topics:

- Migrating from EMC PowerPath
- Migrating from Hitachi Dynamic Link Manager (HDLM)
- Configuring DMP for Active/Active load balancing in a cluster

Migrating from EMC PowerPath

Migrating from EMC PowerPath involves removing the devices from EMC PowerPath (PP) control and enabling InfoScale Foundation on the devices.

The migration process requires you to:

- Stop the applications
- Stop VCS services, if using VCS

As a result, the migration process involves some downtime for the applications running on the systems.

To uninstall EMC PowerPath and install InfoScale Foundation, perform the following:

1. Disable/disconnect all but one path from the system to the storage.
2. Remove the EMC PowerPath (PP).

   Refer to the EMC instructions for removing PowerPath.
3. Reboot the system after the PowerPath has been removed.

4 Install InfoScale Foundation.

During installation, select MPIO device-specific modules (DSMs) for the attached storage (for example, EMC CLARiiON and Hitachi AMS).

5 Reboot the system.

Migrating from Hitachi Dynamic Link Manager (HDLM)

Migrating from HDLM involves removing the devices from HDLM control and enabling InfoScale Foundation on the devices.

The migration process requires you to:

- Back up all the data on the host and on the on the management target device, where HDLM is installed
- Stop applications
- Stop VCS services, if using VCS
- Reboot one or more hosts after uninstalling HDLM

After HDLM is uninstalled, sometimes the following files listed below would not be deleted. The following files will be deleted when you restart the host:

- HDLM-installation-folder\DLMTools\perfhdlm\provhdlm.dll
- HDLM-installation-folder\lib\libdlm.dll
- HDLM-installation-folder\lib\hdlmhcc60.dll

The default installation folder for HDLM is Windows-installation-drive:

C: \Program Files\HITACHI\DynamicLinkManager.

Refer to HDLM documentation for details.

Uninstalling HDLM in a non-clustered environment

To uninstall HDLM in a non-clustered environment, perform the following steps. Refer to HDLM documentation for details:

1 Log on to Windows as a member of the Administrators group.

2 Stop all the processes and services that use the HDLM management-target paths.

Stop any processes or application services, such as a DBMS, that are using the HDLM management-target paths.

3 If the host and the storage subsystem are connected via multiple paths, reconfigure it so that only one path connects the host to the storage subsystem.

After uninstalling HDLM, if you start the host in a multi-path configuration, the disk contents might become corrupted.

4 Start the uninstallation procedure.

5 When uninstallation finishes, a dialog box appears prompting you to restart the host.

Click **OK** to restart the host.

6 Now, install the Dynamic Multi-Pathing (DMP) .

See "Installing the server components using the installation wizard" on page 24.

Uninstalling HDLM in a clustered (MSCS or VCS) environment

To uninstall HDLM in an MSCS or VCS environment perform the following steps. Refer to HDLM documentation for details.

Uninstalling HDLM in a clustered (MSCS or VCS) environment

1 Log on to Windows as a member of the Administrators group.

2 Stop all the processes and services that use the HDLM management-target paths.

Stop any processes or application services, such as a DBMS, that are using the HDLM management-target paths.

3 Stop MSCS or VCS on all the hosts that make up the cluster.

When MSCS is used, follow this procedure:

Choose **Administrative Tools** and then **Services**. In the list of services, right-click **Cluster Service**, and then from the **Action** menu choose **Stop** to stop the service.

A message prompting you to restart the system might be displayed. If this happens, choose No.

4 If a host and a storage subsystem are connected via multiple paths, reconfigure it so that only one path connects the host to the storage subsystem.

Uninstalling HDLM in a multi-path configuration, might cause the disk contents to become corrupted when the host restarts. Make sure that you uninstall HDLM from a single path configuration only.

5 Start the uninstallation procedure.

6 When uninstallation finishes, a dialog box appears prompting you to restart the host.

Click **OK** to restart the host.

7 Install InfoScale Foundation.

See "Installing the server components using the installation wizard" on page 24.

Configuring DMP for Active/Active load balancing in a cluster

SCSI-3 is required for configuring Active/Active (A/A) load balancing in a clustered environment. SCSI-3 is enabled by default when DMP is installed in a clustered environment.

If the disk resources have already been created before setting SCSI-3 support at array level, then they are reserved using SCSI-2 and A/A load balancing policies will not work on those disks.

To use A/A load balancing, enable SCSI-3 reservation for all disk under an array using the vxdmpadm setarray command. This ensures that the disks under the selected array will be reserved using SCSI-3 even if the cluster application issues SCSI-2 reservation for these disks.

Syntax for `vxdmpadm setarray` command:

`vxdmpadm setarrayscsi3 scsi3support=1` *`Harddisk name`*.

Refer to the *Dynamic Multi-Pathing Administrator's Guide* for details.