

Oracle® Solaris Cluster Data Services Planning and Administration Guide

License Restrictions Warranty/Consequential Damages Disclaimer

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

Warranty Disclaimer

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

Restricted Rights Notice

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Hazardous Applications Notice

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Trademark Notice

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group in the United States and other countries.

Third Party Content, Products, and Services Disclaimer

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Contents

Preface	9
1 Planning for Oracle Solaris Cluster Data Services	15
Configuration Guidelines for Oracle Solaris Cluster Data Services	16
Identifying Data Service Special Requirements	16
Determining the Location of the Application Binaries	16
Verifying the <code>nsswitch.conf</code> File Contents	17
Planning the Cluster File System Configuration	17
Enabling Oracle Solaris SMF Services to Run Under the Control of Oracle Solaris Cluster	18
Relationship Between Resource Groups and Device Groups	19
Understanding <code>HASStoragePlus</code>	19
Determining Whether Your Data Service Requires <code>HASStoragePlus</code>	20
Considerations for Installing and Configuring a Data Service	21
Overview of the Installation and Configuration Process	21
Installation and Configuration Task Flow	22
Example of Configuring a Failover Data Service	23
Tools for Data Service Resource Administration	24
Oracle Solaris Cluster Manager Graphical User Interface (GUI)	24
SPARC: The Oracle Solaris Cluster Module for the Sun Management Center GUI	25
<code>clsetup</code> Utility	25
Oracle Solaris Cluster Maintenance Commands	25
Summary by Task of Tools for Administering Data Service Resources	25
Standard Properties	26
Cluster Properties	27
Resource Type Properties	27
Resource Properties	28
Resource Group Properties	29

Resource Property Attributes	30
Node List Properties	31
Legal RGM Names	31
RGM Legal Names	31
RGM Values	33
2 Administering Data Service Resources	35
Overview of Tasks for Administering Data Service Resources	36
Configuring and Administering Oracle Solaris Cluster Data Services	39
Registering a Resource Type	40
▼ How to Register a Resource Type	40
Upgrading a Resource Type	42
▼ How to Install and Register an Upgrade of a Resource Type	42
▼ How to Migrate Existing Resources to a New Version of the Resource Type	43
Downgrading a Resource Type	47
▼ How to Downgrade a Resource to an Older Version of Its Resource Type	48
Creating a Resource Group	49
▼ How to Create a Failover Resource Group	49
▼ How to Create a Scalable Resource Group	51
Configuring Failover and Scalable Applications	53
▼ How to Configure a Failover Application Using the <code>ScalMountPoint</code> Resource	53
▼ How to Configure a Scalable Application Using the <code>ScalMountPoint</code> Resource	54
Tools for Adding Resources to Resource Groups	55
▼ How to Add a Logical Hostname Resource to a Resource Group by Using the <code>clsetup</code> Utility	56
▼ How to Add a Logical Hostname Resource to a Resource Group Using the Command-Line Interface	59
▼ How to Add a Shared Address Resource to a Resource Group by Using the <code>clsetup</code> Utility	61
▼ How to Add a Shared Address Resource to a Resource Group Using the Command-Line Interface	63
▼ How to Add a Failover Application Resource to a Resource Group	65
▼ How to Add a Scalable Application Resource to a Resource Group	68
Bringing Resource Groups Online	70
▼ How to Bring Resource Groups Online	71
Switching Resource Groups to Preferred Primaries	72

▼ How to Switch Resource Groups to Preferred Primaries	72
Enabling a Resource	73
▼ How to Enable a Resource	73
Quiescing Resource Groups	74
▼ How to Quiesce a Resource Group	74
▼ How to Quiesce a Resource Group Immediately	75
Suspending and Resuming the Automatic Recovery Actions of Resource Groups	75
Immediately Suspending Automatic Recovery by Killing Methods	76
▼ How to Suspend the Automatic Recovery Actions of a Resource Group	77
▼ How to Suspend the Automatic Recovery Actions of a Resource Group Immediately	77
▼ How to Resume the Automatic Recovery Actions of a Resource Group	77
Disabling and Enabling Resource Monitors	78
▼ How to Disable a Resource Fault Monitor	78
▼ How to Enable a Resource Fault Monitor	79
Removing Resource Types	80
▼ How to Remove a Resource Type	80
Removing Resource Groups	81
▼ How to Remove a Resource Group	81
Removing Resources	82
▼ How to Remove a Resource	82
Switching the Current Primary of a Resource Group	83
▼ How to Switch the Current Primary of a Resource Group	83
Disabling Resources and Moving Their Resource Group Into the UNMANAGED State	85
▼ How to Disable a Resource and Move Its Resource Group Into the UNMANAGED State	86
Displaying Resource Type, Resource Group, and Resource Configuration Information	88
Changing Resource Type, Resource Group, and Resource Properties	88
▼ How to Change Resource Type Properties	89
▼ How to Change Resource Group Properties	90
▼ How to Change Resource Properties	91
▼ How to Modify a Logical Hostname Resource or a Shared Address Resource	92
Clearing the STOP_FAILED Error Flag on Resources	93
▼ How to Clear the STOP_FAILED Error Flag on Resources	94
Clearing the Start_failed Resource State	95
▼ How to Clear a Start_failed Resource State by Switching Over a Resource Group	95
▼ How to Clear a Start_failed Resource State by Restarting a Resource Group	97
▼ How to Clear a Start_failed Resource State by Disabling and Enabling a Resource	99

Upgrading a Preregistered Resource Type	100
Information for Registering the New Resource Type Version	100
Information for Migrating Existing Instances of the Resource Type	101
Reregistering Preregistered Resource Types After Inadvertent Deletion	101
▼ How to Reregister Preregistered Resource Types After Inadvertent Deletion	102
Adding or Removing a Node to or From a Resource Group	102
Adding a Node to a Resource Group	103
Removing a Node From a Resource Group	106
Migrating the Application From a Global-Cluster Voting Node to a Global-Cluster Non-Voting Node	111
▼ How to Migrate the Application From a Global-Cluster Voting Node to a Global-Cluster Non-Voting Node	111
Synchronizing the Startups Between Resource Groups and Device Groups	114
Managed Entity Monitoring by HAStoragePlus	115
Additional Administrative Tasks to Configure HAStoragePlus Resources for a Zone Cluster	118
▼ How to Set Up the HAStoragePlus Resource Type for New Resources	118
▼ How to Set Up the HAStoragePlus Resource Type for Existing Resources	120
Configuring an HAStoragePlus Resource for Cluster File Systems	121
Sample Entries in /etc/vfstab for Cluster File Systems	121
▼ How to Set Up the HAStoragePlus Resource for Cluster File Systems	122
▼ How to Delete an HAStoragePlus Resource Type for Cluster File Systems	123
Enabling Highly Available Local File Systems	124
Configuration Requirements for Highly Available Local File Systems	125
Format of Device Names for Devices Without a Volume Manager	126
Sample Entries in /etc/vfstab for Highly Available Local File Systems	126
▼ How to Set Up the HAStoragePlus Resource Type by Using the clsetup Utility	127
▼ How to Set Up the HAStoragePlus Resource Type to Make File Systems Highly Available Other Than Solaris ZFS	130
▼ How to Set Up the HAStoragePlus Resource Type to Make a Local Solaris ZFS Highly Available	131
▼ How to Delete an HAStoragePlus Resource That Makes a Local Solaris ZFS Highly Available	134
Sharing a Failover File System Across Zone Clusters	134
Configuration Requirements for Sharing a Failover File System Directory to a Zone Cluster	134
▼ How to Set Up the HAStorage Plus Resource Type to Share a Failover File System Directory	

to a Zone Cluster	135
Upgrading From HASStorage to HASStoragePlus	138
▼ How to Upgrade From HASStorage to HASStoragePlus When Using Device Groups or CFS	138
▼ How to Upgrade From HASStorage With CFS to HASStoragePlus With Highly Available Local File System	139
Modifying Online the Resource for a Highly Available File System	140
▼ How to Add File Systems Other Than Solaris ZFS to an Online HASStoragePlus Resource	141
▼ How to Remove File Systems Other Than Solaris ZFS From an Online HASStoragePlus Resource	143
▼ How to Add a Solaris ZFS Storage Pool to an Online HASStoragePlus Resource	145
▼ How to Remove a Solaris ZFS Storage Pool From an Online HASStoragePlus Resource ..	146
▼ How to Recover From a Fault After Modifying the FileSystemMountPoints Property of an HASStoragePlus Resource	147
▼ How to Recover From a Fault After Modifying the Zpools Property of an HASStoragePlus Resource	148
Changing the Cluster File System to a Local File System in an HASStoragePlus Resource	149
▼ How to Change the Cluster File System to Local File System in an HASStoragePlus Resource	149
Upgrading the HASStoragePlus Resource Type	150
Information for Registering the New Resource Type Version	150
Information for Migrating Existing Instances of the Resource Type	151
Distributing Online Resource Groups Among Cluster Nodes	151
Resource Group Affinities	152
Enforcing Collocation of a Resource Group With Another Resource Group	153
Specifying a Preferred Collocation of a Resource Group With Another Resource Group ..	154
Distributing a Set of Resource Groups Evenly Among Cluster Nodes	155
Specifying That a Critical Service Has Precedence	156
Delegating the Failover or Switchover of a Resource Group	157
Combining Affinities Between Resource Groups	159
Zone Cluster Resource Group Affinities	159
Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources	160
▼ How to Replicate Configuration Data on a Cluster Without Configured Resource Groups, Resource Types, and Resources	161
▼ How to Upgrade Configuration Data on a Cluster With Configured Resource Groups, Resource Types, and Resources	162

Enabling Oracle Solaris SMF Services to Run With Oracle Solaris Cluster	163
▼ Encapsulating an SMF Service Into a Failover Proxy Resource Configuration	164
▼ Encapsulating an SMF Service Into a Multi-Master Proxy Resource Configuration	167
▼ Encapsulating an SMF Service Into a Scalable Proxy Resource Configuration	170
Tuning Fault Monitors for Oracle Solaris Cluster Data Services	173
Setting the Interval Between Fault Monitor Probes	174
Setting the Timeout for Fault Monitor Probes	174
Defining the Criteria for Persistent Faults	175
Specifying the Failover Behavior of a Resource	176
Denying Cluster Services For a Selected Non-Global Zone	176
▼ How to Deny Cluster Services For a Non-Global Zone	177
▼ How to Allow Cluster Services For a Non-Global Zone	177
 A Data Service Configuration Worksheets and Examples	179
Configuration Worksheets	179
Resource Types Worksheet	180
Network Resources Worksheet	182
Application Resources—Failover Worksheet	184
Application Resources—Scalable Worksheet	186
Resource Groups—Failover Worksheet	188
Resource Groups—Scalable Worksheet	190
 Index	193

Preface

Oracle Solaris Cluster Data Services Planning and Administration Guide explains how to install and configure Oracle Solaris Cluster data services.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, AMD64, and Intel 64. In this document, x86 refers to the larger family of 64-bit x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

This document is intended for system administrators with extensive knowledge of Oracle software and hardware. Do not use this document as a planning or presales guide. Before reading this document, you should have already determined your system requirements and purchased the appropriate equipment and software.

The instructions in this book assume knowledge of the Oracle Solaris Operating System and expertise with the volume-manager software that is used with Oracle Solaris Cluster software.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Oracle Solaris Cluster data services. The document does *not* contain comprehensive information about basic UNIX commands and procedures, such as shutting down the system, booting the system, and configuring devices. Information about basic UNIX commands and procedures is available from the following sources:

- Online documentation for the Oracle Solaris Operating System
- Oracle Solaris Operating System man pages
- Other software documentation that you received with your system

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Related Documentation

Information about related Oracle Solaris Cluster topics is available in the documentation that is listed in the following table. All Oracle Solaris Cluster documentation is available at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.

Topic	Documentation
Data service administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> Individual data service guides
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i>
Hardware administration	<i>Oracle Solaris Cluster 3.3 Hardware Administration Manual</i> Individual hardware administration guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function reference	<i>Oracle Solaris Cluster Reference Manual</i>

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your release of Oracle Solaris Cluster at <http://www.oracle.com/technetwork/indexes/documentation/index.html>.

Related Third-Party Web Site References

Third-party URLs that are referenced in this document provide additional related information.

Note – Oracle is not responsible for the availability of third-party web sites mentioned in this document. Oracle does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Oracle will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation and Support

See the following web sites for additional resources:

- [Documentation \(http://www.oracle.com/technetwork/indexes/documentation/index.html\)](http://www.oracle.com/technetwork/indexes/documentation/index.html)
- [Support \(http://www.oracle.com/us/support/systems/index.html\)](http://www.oracle.com/us/support/systems/index.html)

Oracle Software Resources

[Oracle Technology Network \(http://www.oracle.com/technetwork/index.html\)](http://www.oracle.com/technetwork/index.html) offers a range of resources related to Oracle software:

- Discuss technical problems and solutions on the [Discussion Forums \(http://forums.oracle.com\)](http://forums.oracle.com).
- Get hands-on step-by-step tutorials with [Oracle By Example \(http://www.oracle.com/technetwork/tutorials/index.html\)](http://www.oracle.com/technetwork/tutorials/index.html).

Getting Help

If you have problems installing or using Oracle Solaris Cluster, contact your service provider and provide the following information:

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the Oracle Solaris Operating System (for example, Oracle Solaris 10)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 3.3)

Use the following commands to gather information about each node on your system for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Oracle Solaris Cluster release and package version information

Also have available the contents of the `/var/adm/messages` file.

Planning for Oracle Solaris Cluster Data Services

This chapter provides planning information and guidelines to install and configure Oracle Solaris Cluster data services. This chapter contains the following sections.

- “Configuration Guidelines for Oracle Solaris Cluster Data Services” on page 16
- “Relationship Between Resource Groups and Device Groups” on page 19
- “Understanding HASstoragePlus” on page 19
- “Considerations for Installing and Configuring a Data Service” on page 21
- “Overview of the Installation and Configuration Process” on page 21
- “Tools for Data Service Resource Administration” on page 24
- “Standard Properties” on page 26
- “Node List Properties” on page 31
- “Legal RGM Names” on page 31

For introductory information about data services, see *Oracle Solaris Cluster Concepts Guide*.

Oracle Solaris Cluster software can provide service only for those data services that are either supplied with the Oracle Solaris Cluster product or are created with the Oracle Solaris Cluster data services application programming interfaces (APIs).

If an Oracle Solaris Cluster data service is not provided for your application, consider developing a custom data service for the application. To develop a custom data service, use the Oracle Solaris Cluster data services APIs. For more information, see *Oracle Solaris Cluster Data Services Developer's Guide*.

Note – Oracle Solaris Cluster does not provide a data service for the `sendmail(1M)` subsystem. The `sendmail` subsystem can run on the individual cluster nodes, but the `sendmail` functionality is not highly available. This restriction applies to all the `sendmail` functionality, including the functionality of mail delivery and mail routing, queuing, and retry.

Configuration Guidelines for Oracle Solaris Cluster Data Services

This section provides configuration guidelines for Oracle Solaris Cluster data services.

Identifying Data Service Special Requirements

Identify requirements for all of the data services **before** you begin Oracle Solaris and Oracle Solaris Cluster installation. Failure to do so might result in installation errors that require that you completely reinstall the Oracle Solaris and Oracle Solaris Cluster software.

For example, the Oracle Data Guard option of Oracle Solaris Cluster Support for Oracle Real Application Clusters has special requirements for the hostnames that you use in the cluster. Oracle Solaris Cluster HA for SAP also has special requirements. You must accommodate these requirements before you install Oracle Solaris Cluster software because you cannot change hostnames after you install Oracle Solaris Cluster software.

Note – Some Oracle Solaris Cluster data services are not supported for use in x86 based clusters. For more information, see the release notes for your release of Oracle Solaris Cluster at (<http://www.oracle.com/technetwork/indexes/documentation/index.html>).

Determining the Location of the Application Binaries

You can install the application software and application configuration files on one of the following locations.

- **The local disks of each cluster node** – Placing the software and configuration files on the individual cluster nodes provides the following advantage. You can upgrade application software later without shutting down the service.

The disadvantage is that you then have several copies of the software and configuration files to maintain and administer.

- **The cluster file system** – If you put the application binaries on the cluster file system, you have only one copy to maintain and manage. However, you must shut down the data service in the entire cluster to upgrade the application software. If you can spare a short period of downtime for upgrades, place a single copy of the application and configuration files on the cluster file system.

For information about how to create cluster file systems, see “[Planning the Global Devices, Device Groups, and Cluster File Systems](#)” in *Oracle Solaris Cluster Software Installation Guide*.

- **Highly available local file system** – Using `HASStoragePlus`, you can integrate your local file system into the Oracle Solaris Cluster environment, making the local file system highly available. `HASStoragePlus` provides additional file system capabilities such as checks, mounts, and unmounts that enable Oracle Solaris Cluster to fail over local file systems. To fail over, the local file system must reside on global disk groups with affinity switchovers enabled.

For information about how to use the `HASStoragePlus` resource type, see [“Enabling Highly Available Local File Systems” on page 124](#).

Verifying the `nsswitch.conf` File Contents

The `nsswitch.conf` file is the configuration file for name-service lookups. This file determines the following information.

- The databases within the Oracle Solaris environment to use for name-service lookups
- The order in which the databases are to be consulted

Some data services require that you direct “group” lookups to “files” first. For these data services, change the “group” line in the `nsswitch.conf` file so that the “files” entry is listed first. See the documentation for the data service that you plan to configure to determine whether you need to change the “group” line.

For additional information about how to configure the `nsswitch.conf` file for the Oracle Solaris Cluster environment, see [“Planning the Oracle Solaris Cluster Environment” in *Oracle Solaris Cluster Software Installation Guide*](#).

Planning the Cluster File System Configuration

Depending on the data service, you might need to configure the cluster file system to meet Oracle Solaris Cluster requirements. To determine whether any special considerations apply, see the documentation for the data service that you plan to configure.

For information about how to create cluster file systems, see [“Planning the Global Devices, Device Groups, and Cluster File Systems” in *Oracle Solaris Cluster Software Installation Guide*](#).

The resource type `HASStoragePlus` enables you to use a highly available local file system in an Oracle Solaris Cluster environment that is configured for failover. For information about setting up the `HASStoragePlus` resource type, see [“Enabling Highly Available Local File Systems” on page 124](#).

Enabling Oracle Solaris SMF Services to Run Under the Control of Oracle Solaris Cluster

The Service Management Facility (SMF) enables you to automatically start and restart SMF services, during a node boot or service failure. This feature is similar to the Oracle Solaris Cluster Resource Group Manager (RGM), which facilitates high availability and scalability for cluster applications. SMF services and RGM features are complementary to each other.

Oracle Solaris Cluster includes three new SMF proxy resource types that can be used to enable SMF services to run with Oracle Solaris Cluster in a failover, multi-master, or scalable configuration. The SMF proxy resource types enables you to encapsulate a set of interrelated SMF services into a single resource, *SMF proxy resource* to be managed by Oracle Solaris Cluster. In this feature, SMF manages the availability of SMF services on a single node. Oracle Solaris Cluster provides cluster-wide high availability and scalability of the SMF services.

For information about how to encapsulate these services, see [“Enabling Oracle Solaris SMF Services to Run With Oracle Solaris Cluster” on page 163](#).

You might require Oracle Solaris Cluster to make highly available an application other than NFS or DNS that is integrated with the Solaris Service Management Facility (SMF). To ensure that Oracle Solaris Cluster can restart or fail over the application correctly after a failure, you must disable SMF service instances for the application as follows:

- For any application other than NFS or DNS, disable the SMF service instance on all potential primary nodes for the Oracle Solaris Cluster resource that represents the application.
- If multiple instances of the application share any component that you require Oracle Solaris Cluster to monitor, disable all service instances of the application. Examples of such components are daemons, file systems, and devices.

Note – If you do not disable the SMF service instances of the application, both the Solaris SMF and Oracle Solaris Cluster might attempt to control the startup and shutdown of the application. As a result, the behavior of the application might become unpredictable.

For more information, see the following documentation:

- [“How to Disable a Service Instance” in *System Administration Guide: Basic Administration*](#)
- [Oracle Solaris Cluster Data Service for Network File System \(NFS\) Guide](#)
- [Oracle Solaris Cluster Concepts Guide](#)

Relationship Between Resource Groups and Device Groups

Oracle Solaris Cluster uses the concept of **node lists** for device groups and resource groups. Node lists are ordered lists of primary nodes, which are potential masters of the disk device group or resource group. Oracle Solaris Cluster uses a **failback policy** to determine the behavior of Oracle Solaris Cluster in response to the following set of conditions:

- A node that has failed and left the cluster rejoins the cluster.
- The node that is rejoining the cluster appears earlier in the node list than the current primary node.

If failback is set to `True`, the device group or resource group is switched off the current primary and switched onto the rejoining node, making the rejoining node the new primary.

For example, assume that you have a disk device group, `disk-group-1`, that has nodes `phys-schost-1` and `phys-schost-2` in its node list, with the failback policy set to `Enabled`. Assume that you also have a failover resource group, `resource-group-1`, which uses `disk-group-1` to hold its application data. When you set up `resource-group-1`, also specify `phys-schost-1` and `phys-schost-2` for the resource group's node list, and set the failback policy to `True`.

To ensure high availability of a scalable resource group, make the scalable resource group's node list a superset of the node list for the disk device group. This setting ensures that the nodes that are directly connected to the disks are also nodes that can run the scalable resource group. The advantage is that, when at least one cluster node connected to the data is up, the scalable resource group runs on that same node, making the scalable services available also.

For more information about the relationship between device groups and resource groups, see [“Device Groups” in *Oracle Solaris Cluster Concepts Guide*](#).

For information about how to set up device groups, see the following documentation:

- [“Device Groups” in *Oracle Solaris Cluster Software Installation Guide*](#)

Understanding HASToragePlus

The HASToragePlus resource type can be used to configure the following options.

- Coordinate the boot order of disk devices and resource groups. Other resources in the resource group that contains the HASToragePlus resource are brought online *only* after the disk device resources become available.
- With `AffinityOn` set to `True`, enforce collocation of resource groups and device groups on the same node. This enforced collocation enhances the performance of disk-intensive data services.
- Monitor entities managed by the HASP resource, including global devices, file systems, and ZFS storage pools.

In addition, HASToragePlus is capable of mounting local and global file systems. For more information, see [“Planning the Cluster File System Configuration” on page 17](#).

Note – If the device group is switched to another node while the HASToragePlus resource is online, AffinityOn has no effect. The resource group does **not** migrate with the device group. However, if the resource group is switched to another node, the setting of AffinityOn to True causes the device group to follow the resource group to the new node.

See [“Synchronizing the Startups Between Resource Groups and Device Groups” on page 114](#) for information about the relationship between device groups and resource groups.

See [“Enabling Highly Available Local File Systems” on page 124](#) for procedures for mounting of file systems such as VxFS and Solaris ZFS (Zettabyte File System) in a local mode. The [SUNW.HASToragePlus\(5\)](#) man page provides additional details.

Determining Whether Your Data Service Requires HASToragePlus

The following types of data services require HASToragePlus:

- Data services with nodes that are not directly connected to storage
- Data services that are disk intensive

Data Services With Nodes That Are Not Directly Connected to Storage

Some nodes in the node list of a data service's resource group might not be directly connected to the storage. In this situation, you must coordinate the boot order between the storage and the data service. To meet this requirement, configure the resource group as follows:

- Configure HASToragePlus resources in the resource group.
- Set the dependency of the other data service resources to the HASToragePlus resource.

Data Services That Are Disk Intensive

Some data services, such as Oracle Solaris Cluster HA for Oracle and Oracle Solaris Cluster HA for NFS are disk intensive. If your data service is disk intensive, ensure that the resource groups and device groups are collocated on the same node. To meet this requirement, perform the following tasks.

- Adding an HASToragePlus resource to your data service resource group
- Switching the HASToragePlus resource online
- Setting the dependency of your data service resources to the HASToragePlus resource
- Setting AffinityOn to True

Note – The failback settings must be identical for both the resource group and device groups.

Some data services are not disk intensive. For example, HA for DNS, which reads all of its files at startup, is not disk intensive. If your data service is **not** disk intensive, configuring the `HAStoragePlus` resource type is optional.

Considerations for Installing and Configuring a Data Service

Use the information in this section to plan the installation and configuration of any data service. The information in this section encourages you to think about the impact your decisions have on the installation and configuration of any data service. For specific considerations for a data service, see the documentation for the data service.

- Retries within the I/O subsystem during disk failures might cause applications whose data services are disk intensive to experience delays. Disk-intensive data services are I/O intensive and have a large number of disks configured in the cluster. An I/O subsystem might require several minutes to retry and recover from a disk failure. This delay can cause Oracle Solaris Cluster to fail over the application to another node, even though the disk might have eventually recovered on its own. To avoid failover during these instances, consider increasing the default probe timeout of the data service. If you need more information or help with increasing data service timeouts, contact your local support engineer.
- For better performance, install and configure your data service on the cluster nodes with direct connection to the storage.
- Client applications that run on cluster nodes should not map to logical IP addresses of an HA data service. After a failover, these logical IP addresses might no longer exist, leaving the client without a connection.

Overview of the Installation and Configuration Process

Use the following procedures to install and configure a data service.

- Install the data service packages from the installation medium on which the packages are supplied.
- Install and configure the application to run in the cluster environment.
- Configure the resources and resource groups that the data service uses. When you configure a data service, specify the resource types, resources, and resource groups that the Resource Group Manager (RGM) is to manage. The documentation for the individual data services describes these procedures.

Before you install and configure data services, see [Oracle Solaris Cluster Software Installation Guide](#), which includes instructions for the following tasks:

- Installing the data service software packages
- Configuring IP network multipathing groups that the network resources use

Note – You can use Oracle Solaris Cluster Manager to install and configure the following data services: Oracle Solaris Cluster HA for Oracle, Oracle Solaris Cluster HA for Java System Web Server, HA for Java System Web Server, Oracle Solaris Cluster HA for Apache, HA for DNS, and Oracle Solaris Cluster HA for NFS. See the Oracle Solaris Cluster Manager online help for more information.

Installation and Configuration Task Flow

The following table summarizes the tasks for installing and configuring Oracle Solaris Cluster data services. The table also provides cross-references to detailed instructions for performing the tasks.

TABLE 1-1 Tasks for Installing and Configuring Oracle Solaris Cluster Data Services

Task	Instructions
Install the Oracle Solaris and Oracle Solaris Cluster software	Oracle Solaris Cluster Software Installation Guide
Set up IPMP groups	Oracle Solaris Cluster Software Installation Guide
Set up multihost disks	Oracle Solaris Cluster Software Installation Guide
Plan resources and resource groups	Appendix A, “Data Service Configuration Worksheets and Examples”
Decide the location for application binaries, and configure the <code>nsswitch.conf</code> file	“Determining the Location of the Application Binaries” on page 16 “Verifying the <code>nsswitch.conf</code> File Contents” on page 17
Install and configure the application software	The appropriate Oracle Solaris Cluster data services book
Install the data service software packages	Oracle Solaris Cluster Software Installation Guide or the appropriate Oracle Solaris Cluster data services book
Register and configure the data service	The appropriate Oracle Solaris Cluster data services book

Example of Configuring a Failover Data Service

This example summarizes how to set up the resource types, resources, and resource groups that a failover data service for the Oracle application requires. For complete instructions for configuring the data service for the Oracle application, see [Oracle Solaris Cluster Data Service for Oracle Guide](#).

The principal difference between this example and an example of a scalable data service is as follows: In addition to the failover resource group that contains the network resources, a scalable data service requires a separate resource group (*scalable resource group*) for the application resources.

The Oracle application has two components, a server and a listener. Oracle supplies the Oracle Solaris Cluster HA for Oracle data service, and therefore these components have already been mapped into Oracle Solaris Cluster resource types. Both of these resource types are associated with resources and resource groups.

Because this example is a failover data service, the example uses logical hostname network resources, which are the IP addresses that fail over from a primary node to a secondary node. Place the logical hostname resources into a failover resource group, and then place the Oracle server resources and listener resources into the same resource group. This ordering enables all of the resources to fail over as a group.

For Oracle Solaris Cluster HA for Oracle to run on the cluster, you must define the following objects.

- LogicalHostname resource type – This resource type is built in, and therefore you do not need to explicitly register the resource type.
- Oracle resource types – Oracle Solaris Cluster HA for Oracle defines two Oracle resource types—a database server and a listener.
- Logical hostname resources – These resources host the IP addresses that fail over in a node failure.
- Oracle resources – You must specify two resource instances for Oracle Solaris Cluster HA for Oracle — a server and a listener.
- Failover resource group – This container is composed of the Oracle server and listener and logical hostname resources that will fail over as a group.

Tools for Data Service Resource Administration

This section describes the tools that you can use to perform installation and configuration tasks.

Oracle Solaris Cluster Manager Graphical User Interface (GUI)

Oracle Solaris Cluster Manager is a web-based tool that enables you to perform the following tasks.

- Installing a cluster
- Administering a cluster
- Creating and configuring resources and resource groups
- Configuring data services with the Oracle Solaris Cluster software

Oracle Solaris Cluster Manager provides wizards to automate the configuration of Oracle Solaris Cluster data services for the following applications.

- Apache Web Server
- NFS
- Oracle
- Oracle Real Application Clusters
- SAP Web Application Server

Each wizard enables you to configure Oracle Solaris Cluster resources that the data service requires. The wizard does not automate the installation and configuration of the application software to run in an Oracle Solaris Cluster configuration. To install and configure application software to run in an Oracle Solaris Cluster configuration, use utilities of the application and Oracle Solaris Cluster maintenance commands. For more information, see your application documentation and the Oracle Solaris Cluster documentation set. Each wizard supports only a limited subset of configuration options for a data service. To configure options that a wizard does not support, use Oracle Solaris Cluster Manager or Oracle Solaris Cluster maintenance commands to configure the data service manually. For more information, see the Oracle Solaris Cluster documentation.

Oracle Solaris Cluster Manager provides wizards to automate the configuration of the following Oracle Solaris Cluster resources.

- Logical hostname resource
- Shared address resource
- Highly available storage resource

You can use a resource that you create by using a wizard with any data service regardless of how you configure the data service.

For instructions on using use Oracle Solaris Cluster Manager to install cluster software, see *Oracle Solaris Cluster Software Installation Guide*. Oracle Solaris Cluster Manager provides online help for most administrative tasks.

SPARC: The Oracle Solaris Cluster Module for the Sun Management Center GUI

The Oracle Solaris Cluster module enables you to monitor clusters and to perform some operations on resources and resource groups from the Sun Management Center GUI. See the *Oracle Solaris Cluster Software Installation Guide* for information about installation requirements and procedures for the Oracle Solaris Cluster module. Go to <http://www.oracle.com/technetwork/indexes/documentation/index.html> to access the Sun Management Center software documentation set, which provides additional information about Sun Management Center.

c1setup Utility

The `c1setup(1CL)` utility is a menu-driven interface that you can use for general Oracle Solaris Cluster administration. You can also use this utility to configure data service resources and resource groups. Select option 2 from the `c1setup` main menu to launch the Resource Group Manager submenu.

Oracle Solaris Cluster Maintenance Commands

You can use the Oracle Solaris Cluster maintenance commands to register and configure data service resources. See the procedure for how to register and configure your data service in the book for the data service. If, for example, you are using Oracle Solaris Cluster HA for Oracle, see “Registering and Configuring HA for Oracle” in *Oracle Solaris Cluster Data Service for Oracle Guide*.

For more information about how to use the commands to administer data service resources, see [Chapter 2, “Administering Data Service Resources.”](#)

Summary by Task of Tools for Administering Data Service Resources

The following table summarizes by task which tool you can use to administer data service resources. For more information about these tasks and for details about how to use the command line to complete related procedures, see [Chapter 2, “Administering Data Service Resources.”](#)

TABLE 1-2 Tools for Administering Data Service Resources

Task	Oracle Solaris Cluster Manager	SPARC: Sun Management Center	c1setup Utility
Register a resource type	+	—	+
Create a resource group	+	—	+
Add a resource to a resource group	+	—	+
Suspend the automatic recovery actions of a resource group	+	—	+
Resume the automatic recovery actions of a resource group	+	—	+
Bring a resource group online	+	+	+
Remove a resource group	+	+	+
Remove a resource	+	+	+
Switch the current primary of a resource group	+	—	+
Enable a resource	+	+	+
Disable a resource	+	+	+
Move a resource group to the unmanaged state	+	—	+
Display resource type, resource group, and resource configuration information	+	+	+
Change resource properties	+	—	+
Clear the STOP_FAILED error flag on resources	+	—	+
Clear the START_FAILED resource state for a resource	+	—	+
Add a node to a resource group	+	—	+

Standard Properties

You can specify the following standard cluster, resource type, resource, and resource group properties when configuring data services. Resource property attributes are also available to change system-defined properties and create extension properties.

This section lists the following properties:

- Cluster Properties
- Resource Type Properties
- Resource Properties
- Resource Group Properties
- Resource Property Attributes

Cluster Properties

Cluster properties are used to manage data services. For more information on the `Concentrate_load` property, see the [cluster\(1CL\)](#) man page.

Resource Type Properties

Resource type properties cannot be updated by administrative utilities with the exception of `Installed_nodes` and `RT_system`. `Installed_nodes` cannot be declared in the RTR file and can only be set by the cluster administrator. `RT_system` can be assigned an initial value in the RTR file, and can also be set by the cluster administrator.

For more detail on each resource type property, see the [rt_properties\(5\)](#) man page.

Following is a list of the resource type properties that are defined by Oracle Solaris Cluster:

- `API_version`
- `Boot`
- `Failover`
- `Fini`
- `Global_zone`
- `Init`
- `Init_nodes`
- `Installed_nodes`
- `Is_logical_hostname`
- `Is_shared_address`
- `Monitor_check`
- `Monitor_start`
- `Monitor_stop`
- `Pkglist`
- `Postnet_stop`
- `Prenet_start`
- `Proxy`
- `Resource_list`
- `Resource_type`
- `RT_basedir`

- RT_description
- RT_system
- RT_version
- Single_instance
- Start
- Stop
- Update
- Validate
- Vendor_ID

Resource Properties

This section describes the resource properties that are defined by the Oracle Solaris Cluster software.

For more detail on each resource property, see the [r_properties\(5\)](#) man page.

- Affinity_timeout
- Boot_timeout
- Cheap_probe_interval
- Conn_threshold
- Extension properties
- Failover_mode
- Fini_timeout
- Generic_affinity
- Global_zone_override
- Init_timeout
- IPsec_session_failover
- Load_balancing_policy
- Load_balancing_weights
- Monitor_check_timeout
- Monitor_start_timeout
- Monitor_stop_timeout
- Monitored_switch
- Network_resources_used (deprecated)
- Num_resource_restarts
- Num_rg_restarts
- On_off_switch
- Outgoing_Connection
- Port_list
- Postnet_stop_timeout
- Prenet_start_timeout
- R_description
- Resource_dependencies

- Resource_dependencies_offline_restart
- Resource_dependencies_restart
- Resource_dependencies_weak
- Resource_name
- Resource_project_name
- Resource_state
- Retry_count
- Retry_interval
- Round_robin
- Scalable
- Start_timeout
- Status
- Status_msg
- Stop_timeout
- Thorough_probe_interval
- Type
- Type_version
- UDP_affinity
- Update_timeout
- Validate_timeout
- Weak_affinity

Resource Group Properties

The following information describes the resource group properties that are defined by the Oracle Solaris Cluster software. For more information on each property, see the [rg_properties\(5\)](#) man page.

Following is a list of resource group properties:

- Auto_start_on_new_cluster
- Desired_primaries
- Failback
- Global_resources_used
- Implicit_network_dependencies
- Load_factors
- Maximum_primaries
- Nodelist
- Pathprefix
- Pingpong_interval
- Preemption_mode
- Priority
- Resource_list
- RG_affinities

- RG_dependencies
- RG_description
- RG_is_frozen
- RG_mode
- RG_name
- RG_project_name
- RG_slm_cpu
- RG_slm_cpu_min
- RG_slm_type
- RG_slm_pset_type
- RG_state
- RG_system
- Suspend_automatic_recovery

Resource Property Attributes

This section describes the resource property attributes that you can use to change system-defined properties or to create extension properties. for more information on each property, see the [property_attributes\(5\)](#) man page.

Following is a list of resource properties:

- Array_maxsize
- Array_minsize
- Default
- Description
- Enumlist
- Extension properties
- Max
- Maxlength
- Min
- Minlength
- Per_node
- Property
- Tunable
- Type

Node List Properties

You can specify the following node list properties when configuring data services:

- `Installed_nodes` Property – See the [rt_properties\(5\)](#) man page for more information.
- `Nodelist` Property – See the [rg_properties\(5\)](#) man page for more information.
- `Auxnodelist` Property – See the [clressharedaddress\(1CL\)](#) man page for more information.

Legal RGM Names

This section lists the requirements for legal characters for Resource Group Manager (RGM) names and values.

This section covers the following topics:

- “RGM Legal Names” on page 31
- “RGM Values” on page 33

RGM Legal Names

RGM names fall into the following categories:

- Resource group names
- Resource type names
- Resource names
- Property names
- Enumeration literal names

Rules for Names Except Resource Type Names

Except for resource type names, all names must comply with these rules:

- Names must be in ASCII.
- Names must start with a letter.
- Names can contain uppercase and lowercase letters, digits, dashes (-), and underscores (_).
- The maximum number of characters that you can use in a name is 255.

Format of Resource Type Names

The format of the complete name of a resource type depends on the resource type, as follows:

- If the resource type's resource type registration (RTR) file contains the `#$upgrade` directive, the format is as follows:

vendor-id.base-rt-name:rt-version

- If the resource type's RTR file does *not* contain the `#$upgrade` directive, the format is as follows:

vendor-id.base-rt-name

A period separates *vendor-id* and *base-rt-name*. A colon separates *base-rt-name* and *rt-version*.

The variable elements in this format are as follows:

<i>vendor-id</i>	Specifies the vendor ID prefix, which is the value of the <code>Vendor_id</code> resource type property in the RTR file. If you are developing a resource type, choose a vendor ID prefix that uniquely identifies the vendor, such as your company's stock ticker symbol. For example, the vendor ID prefix of resource types that are developed by Oracle's Sun Microsystems, Inc. is <code>SUNW</code> .
<i>base-rt-name</i>	Specifies the base resource type name, which is the value of the <code>Resource_type</code> resource type property in the RTR file.
<i>rt-version</i>	Specifies the version suffix, which is the value of the <code>RT_version</code> resource type property in the RTR file. The version suffix is <i>only</i> part of the complete resource type name if the RTR file contains the <code>#\$upgrade</code> directive. The <code>#\$upgrade</code> directive was introduced in Release 3.1 of the Sun Cluster product.

Note – If only one version of a base resource type name is registered, you do not have to use the complete name in administrative commands. You can omit the vendor ID prefix, the version number suffix, or both.

For more information, see [“Resource Type Properties” on page 27](#).

EXAMPLE 1-1 Complete Name of a Resource Type With the `#$upgrade` Directive

This example shows the complete name of a resource type for which properties in the RTR file are set, as follows:

- `Vendor_id=SUNW`
- `Resource_type=sample`
- `RT_version=2.0`

The complete name of the resource type that is defined by this RTR file is as follows:

EXAMPLE 1-1 Complete Name of a Resource Type With the #Upgrade Directive (Continued)

```
SUNW.sample:2.0
```

EXAMPLE 1-2 Complete Name of a Resource Type Without the #Upgrade Directive

This example shows the complete name of a resource type for which properties in the RTR file are set, as follows:

- Vendor_id=SUNW
- Resource_type=nfs

The complete name of the resource type that is defined by this RTR file is as follows:

```
SUNW.nfs
```

RGM Values

RGM values fall into two categories: property values and description values. Both categories share the same rules:

- Values must be in ASCII.
- The maximum length of a value is 4 megabytes minus 1, that is, 4,194,303 bytes.
- Values cannot contain the following characters:
 - Null
 - Newline
 - Comma (,)
 - Semicolon (;)

Administering Data Service Resources

This chapter describes how to use the Oracle Solaris Cluster maintenance commands to manage resources, resource groups, and resource types within the cluster. To determine if you can use other tools to complete a procedure, see [“Tools for Data Service Resource Administration” on page 24](#).

For overview information about resource types, resource groups, and resources, see [Chapter 1](#), [“Planning for Oracle Solaris Cluster Data Services,”](#) and *Oracle Solaris Cluster Concepts Guide*.

This chapter contains the following sections.

- [“Overview of Tasks for Administering Data Service Resources” on page 36](#)
- [“Configuring and Administering Oracle Solaris Cluster Data Services” on page 39](#)
- [“Registering a Resource Type” on page 40](#)
- [“Upgrading a Resource Type” on page 42](#)
- [“Downgrading a Resource Type” on page 47](#)
- [“Creating a Resource Group” on page 49](#)
- [“Configuring Failover and Scalable Applications” on page 53](#)
- [“Tools for Adding Resources to Resource Groups” on page 55](#)
- [“Bringing Resource Groups Online” on page 70](#)
- [“Switching Resource Groups to Preferred Primaries” on page 72](#)
- [“Enabling a Resource” on page 73](#)
- [“Quiescing Resource Groups” on page 74](#)
- [“Suspending and Resuming the Automatic Recovery Actions of Resource Groups” on page 75](#)
- [“Disabling and Enabling Resource Monitors” on page 78](#)
- [“Removing Resource Types” on page 80](#)
- [“Removing Resource Groups” on page 81](#)
- [“Removing Resources” on page 82](#)
- [“Switching the Current Primary of a Resource Group” on page 83](#)
- [“Disabling Resources and Moving Their Resource Group Into the UNMANAGED State” on page 85](#)

- “Displaying Resource Type, Resource Group, and Resource Configuration Information” on page 88
- “Changing Resource Type, Resource Group, and Resource Properties” on page 88
- “Clearing the STOP_FAILED Error Flag on Resources” on page 93
- “Clearing the Start_failed Resource State” on page 95
- “Upgrading a Preregistered Resource Type” on page 100
- “Reregistering Preregistered Resource Types After Inadvertent Deletion” on page 101
- “Adding or Removing a Node to or From a Resource Group” on page 102
- “Migrating the Application From a Global-Cluster Voting Node to a Global-Cluster Non-Voting Node” on page 111
- “Synchronizing the Startups Between Resource Groups and Device Groups” on page 114
- “Configuring an HASToragePlus Resource for Cluster File Systems” on page 121
- “Enabling Highly Available Local File Systems” on page 124
- “Upgrading From HASTorage to HASToragePlus” on page 138
- “Modifying Online the Resource for a Highly Available File System” on page 140
- “Changing the Cluster File System to a Local File System in an HASToragePlus Resource” on page 149
- “Upgrading the HASToragePlus Resource Type” on page 150
- “Distributing Online Resource Groups Among Cluster Nodes” on page 151
- “Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources” on page 160
- “Enabling Oracle Solaris SMF Services to Run With Oracle Solaris Cluster” on page 163
- “Tuning Fault Monitors for Oracle Solaris Cluster Data Services” on page 173
- “Denying Cluster Services For a Selected Non-Global Zone” on page 176

Overview of Tasks for Administering Data Service Resources

The following table summarizes the tasks for installing and configuring Oracle Solaris Cluster data services. The table also provides cross-references to detailed instructions for performing the tasks.

TABLE 2-1 Tasks for Administering Data Service Resources

Task	Instructions
Register a resource type	“How to Register a Resource Type” on page 40
Upgrade a resource type	“How to Migrate Existing Resources to a New Version of the Resource Type” on page 43 “How to Install and Register an Upgrade of a Resource Type” on page 42
Downgrade a resource type	“How to Downgrade a Resource to an Older Version of Its Resource Type” on page 48

TABLE 2-1 Tasks for Administering Data Service Resources (Continued)

Task	Instructions
Create failover or scalable resource groups	“How to Create a Failover Resource Group” on page 49 “How to Create a Scalable Resource Group” on page 51
Add logical hostnames or shared addresses and data service resources to resource groups	“How to Add a Logical Hostname Resource to a Resource Group by Using the clsetup Utility” on page 56 “How to Add a Logical Hostname Resource to a Resource Group Using the Command-Line Interface” on page 59 “How to Add a Shared Address Resource to a Resource Group by Using the clsetup Utility” on page 61 “How to Add a Shared Address Resource to a Resource Group Using the Command-Line Interface” on page 63 “How to Add a Failover Application Resource to a Resource Group” on page 65 “How to Add a Scalable Application Resource to a Resource Group” on page 68
Enable resources and resource monitors, manage the resource group, and bring the resource group and its associated resources online	“How to Enable a Resource” on page 73 “How to Bring Resource Groups Online” on page 71 “How to Switch Resource Groups to Preferred Primaries” on page 72
Quiesce a resource group	“How to Quiesce a Resource Group” on page 74 “How to Quiesce a Resource Group Immediately” on page 75
Suspend and resume automatic recovery actions of a resource group	“How to Suspend the Automatic Recovery Actions of a Resource Group” on page 77 “How to Suspend the Automatic Recovery Actions of a Resource Group Immediately” on page 77 “How to Resume the Automatic Recovery Actions of a Resource Group” on page 77
Disable and enable resource monitors independent of the resource	“How to Disable a Resource Fault Monitor” on page 78 “How to Enable a Resource Fault Monitor” on page 79
Remove resource types from the cluster	“How to Remove a Resource Type” on page 80
Remove resource groups from the cluster	“How to Remove a Resource Group” on page 81
Remove resources from resource groups	“How to Remove a Resource” on page 82
Switch the primary for a resource group	“How to Switch the Current Primary of a Resource Group” on page 83

TABLE 2-1 Tasks for Administering Data Service Resources (Continued)

Task	Instructions
Disable resources and move their resource group into the UNMANAGED state	“How to Disable a Resource and Move Its Resource Group Into the UNMANAGED State” on page 86
Display resource type, resource group, and resource configuration information	“Displaying Resource Type, Resource Group, and Resource Configuration Information” on page 88
Change resource type, resource group, and resource properties	“How to Change Resource Type Properties” on page 89 “How to Change Resource Group Properties” on page 90 “How to Change Resource Properties” on page 91
Clear error flags for failed Resource Group Manager (RGM) processes	“How to Clear the STOP_FAILED Error Flag on Resources” on page 94
Clear the Start_failed resource state	“How to Clear a Start_failed Resource State by Switching Over a Resource Group” on page 95 “How to Clear a Start_failed Resource State by Restarting a Resource Group” on page 97 “How to Clear a Start_failed Resource State by Disabling and Enabling a Resource” on page 99
Reregister the built-in resource types LogicalHostname and SharedAddress	“How to Reregister Preregistered Resource Types After Inadvertent Deletion” on page 102
Update the network interface ID list for the network resources, and update the node list for the resource group	“Adding a Node to a Resource Group” on page 103
Remove a node from a resource group	“Removing a Node From a Resource Group” on page 106
Migrate an application from a global-cluster voting node to a global-cluster non-voting node	“How to Migrate the Application From a Global-Cluster Voting Node to a Global-Cluster Non-Voting Node” on page 111
Set up HASStoragePlus for resource groups to synchronize the startups between those resource groups and device groups	“How to Set Up the HASStoragePlus Resource Type for New Resources” on page 118 “How to Set Up the HASStoragePlus Resource Type for Existing Resources” on page 120 “How to Set Up the HASStoragePlus Resource for Cluster File Systems” on page 122 “How to Set Up the HASStoragePlus Resource Type by Using the clsetup Utility” on page 127
Set up the HASStoragePlus to make a local Solaris ZFS highly available	“How to Set Up the HASStoragePlus Resource Type to Make a Local Solaris ZFS Highly Available” on page 131

TABLE 2-1 Tasks for Administering Data Service Resources *(Continued)*

Task	Instructions
Modify online the resource for a highly available file system	“Modifying Online the Resource for a Highly Available File System” on page 140
Change the cluster file system to local file system in an HASToragePlus resource	“Changing the Cluster File System to a Local File System in an HASToragePlus Resource” on page 149
Upgrade the built-in resource types LogicalHostname and SharedAddress	“Upgrading a Resource Type” on page 42 “Upgrading a Preregistered Resource Type” on page 100
Upgrade the HASToragePlus resource type	“Upgrading a Resource Type” on page 42 “Upgrading the HASToragePlus Resource Type” on page 150
Distribute online resource groups among cluster nodes	“Distributing Online Resource Groups Among Cluster Nodes” on page 151
Replicate and upgrade configuration data for resource groups, resource types, and resources	“Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources” on page 160
Enable Solaris SMF services to run with Oracle Solaris Cluster	“Enabling Oracle Solaris SMF Services to Run With Oracle Solaris Cluster” on page 163
Tune fault monitors for Oracle Solaris Cluster data services	“Tuning Fault Monitors for Oracle Solaris Cluster Data Services” on page 173

Note – The procedures in this chapter describe how to use the Oracle Solaris Cluster maintenance commands to complete these tasks. Other tools also enable you to administer your resources. See [“Tools for Data Service Resource Administration” on page 24](#) for details about these options.

Configuring and Administering Oracle Solaris Cluster Data Services

Configuring an Oracle Solaris Cluster data service involves the following tasks.

- Registering a resource type
- Upgrading a resource type
- Creating resource groups
- Adding resources to resource groups
- Bringing resources online

Use the procedures in this chapter to update your data service configuration after the initial configuration. For example, to change resource type, resource group, and resource properties, go to [“Changing Resource Type, Resource Group, and Resource Properties” on page 88](#).

Registering a Resource Type

A resource type provides specification of common properties and callback methods that apply to all of the resources of the given type. You must register a resource type before you create a resource of that type. For details about resource types, see [Chapter 1, “Planning for Oracle Solaris Cluster Data Services.”](#)

An administrator can register a resource type for a zone cluster by specifying a resource type registration (RTR) file that resides inside the zone cluster. In other words, the file must be under the zone root path. The RTR file inside the zone cluster cannot have the `Global_zone` property set to `TRUE`. The RTR file inside the zone cluster cannot be of type `RTR_LOGICAL_HOSTNAME` or `RTR_SHARED_ADDRESS`.

Note – If you want to register a resource type in a zone cluster that uses Solaris Trusted Extensions, and you want to set the `Global_zone` resource-type property to `TRUE`, you must run the registration command only from the `/usr/cluster/lib/rgm/rtrreg` directory of the global cluster.

The administrator can also register a resource type for a zone cluster from the location `/usr/cluster/lib/rgm/rtrreg`. The administrator in the zone cluster cannot modify any RTR files in this directory. This enables registering system resource types for a zone cluster, even when the RTR file has one of the properties that cannot be set directly from the zone cluster. This process provides a secure way of delivering system resource types.

The resource types in the `/usr/cluster/lib/rgm/rtrreg` directory are for the exclusive use of the global cluster.

▼ How to Register a Resource Type

Note – Perform this procedure from any cluster node.

Before You Begin

Ensure that you have the name for the resource type that you plan to register. The resource type name is an abbreviation for the data service name. For information about resource type names of data services that are supplied with Oracle Solaris Cluster, see the release notes for your release of Oracle Solaris Cluster.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Register the resource type.

```
# clresourcetype register resource-type
```

resource-type Specifies name of the resource type to add. See the release notes for your release of Oracle Solaris Cluster to determine the predefined name to supply.
- 3 Verify that the resource type has been registered.

```
# clresourcetype show
```

Example 2-1 Registering a Resource Type

The following example registers the `SUNW.krb5` resource type, which represents the Java System Web Server application in an Oracle Solaris Cluster configuration.

```
# clresourcetype register SUNW.krb5
# clresourcetype show SUNW.krb5
```

Resource Type:	SUNW.krb5
RT_description:	HA-Kerberos KDC server for Sun Cluster
RT_version:	3.2
API_version:	6
RT_basedir:	/opt/SUNWsc krb5/bin
Single_instance:	False
Proxy:	False
Init_nodes:	All potential masters
Installed_nodes:	<All>
Failover:	True
Pkglist:	SUNWsc krb5
RT_system:	False

Next Steps After registering resource types, you can create resource groups and add resources to the resource group. For details, see [“Creating a Resource Group” on page 49](#).

See Also The following man pages:

- [clresourcetype\(1CL\)](#)
- [clresourcegroup\(1CL\)](#)
- [clresource\(1CL\)](#)

Upgrading a Resource Type

Upgrading a resource type enables you to use new features that are introduced in the new version of the resource type. A new version of a resource type might differ from a previous version in the following ways.

- Default settings of resource type properties might change.
- New extension properties of the resource type might be introduced.
- Existing extension properties of the resource type might be withdrawn.
- The set of standard properties that are declared for the resource type might change.
- The attributes of resource properties such as `min`, `max`, `arraymin`, `arraymax`, `default`, and `tunability` might change.
- The set of declared methods might differ.
- The implementation of methods or the fault monitor might change.

Upgrading a resource type involves the tasks that are explained in the following sections:

1. [“How to Install and Register an Upgrade of a Resource Type” on page 42](#)
2. [“How to Migrate Existing Resources to a New Version of the Resource Type” on page 43](#)

▼ How to Install and Register an Upgrade of a Resource Type

The instructions that follow explain how to use the `clresource(1CL)` command to perform this task. However, you are not restricted to using the `clresource` command for this task. Instead of the `clresource` command, you can use Oracle Solaris Cluster Manager or the Resource Group option of the `clsetup(1CL)` command to perform this task.

Before You Begin Consult the documentation for the resource type to determine what you must do before installing the upgrade package on a node. One action from the following list will be required:

- You must reboot the node in noncluster mode.
- You may leave the node running in cluster mode, but you must turn off monitoring of all instances of the resource type.
- You may leave the node running in cluster mode and leave monitoring turned on for all instances of the resource type.

If you must reboot the node in noncluster mode, prevent a loss of service by performing a rolling upgrade. In a rolling upgrade, you install the package on each node individually while leaving the remaining nodes running in cluster mode.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Install the package for the resource type upgrade on all cluster nodes where instances of the resource type are to be brought online.

- 3 Register the new version of the resource type.

To ensure that the correct version of the resource type is registered, you must specify the following information:

- The resource type name
- The resource type registration (RTR) file that defines the resource type

```
# clresourcetype register -f path-to-new-rtr-file resource-type-name
```

The format of the resource type name is as follows:

```
vendor-id.base-rt-name:rt-version
```

For an explanation of this format, see [“Format of Resource Type Names” on page 32](#).

- 4 Display the newly registered resource type.

```
# clresourcetype show resource-type-name
```

- 5 If necessary, set the `Installed_nodes` property to the nodes where the package for the resource type upgrade is installed.

You must perform this step if the package for the resource type upgrade is not installed on all cluster nodes.

The `nodelist` property of all resource groups that contain instances of the resource type must be a subset of the `Installed_nodes` property of the resource type.

```
# clresourcetype set -n installed-node-list resource-type
```

`-n installed-node-list` Specifies the names of nodes on which this resource type is installed.

▼ How to Migrate Existing Resources to a New Version of the Resource Type

The instructions that follow explain how to use the `clresource(1CL)` command to perform this task. However, you are not restricted to using the `clresource` command for this task. Instead of the `clresource` command, you can use Oracle Solaris Cluster Manager or the Resource Group option of the `clsetup(1CL)` command to perform this task.

Before You Begin Consult the instructions for upgrading the resource type to determine when you can migrate resources to a new version of the resource type.

- Any time
- Only when the resource is unmonitored
- Only when the resource is offline
- Only when the resource is disabled
- Only when the resource group is unmanaged

The instructions might state that you cannot upgrade your existing version of the resource. If you cannot migrate the resource, consider the following alternatives:

- Deleting the resource and replacing it with a new resource of the upgraded version
- Leaving the resource at the old version of the resource type

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**
- 2 **For each resource of the resource type that is to be migrated, change the state of the resource or its resource group to the appropriate state.**

- **If you can migrate the resource at any time, no action is required.**
- **If you can migrate the resource only when the resource is unmonitored, type the following command:**

```
# clresource unmonitor resource
```

- **If you can migrate the resource only when the resource is offline, type the following command:**

```
# clresource disable resource
```

Note – If other resources depend on the resource that you are migrating, this step fails. In this situation, consult the error message that is printed to determine the names of the dependent resources. Then repeat this step, specifying a comma-separated list that contains the resource that you are migrating and any dependent resources.

- **If you can migrate the resource only when the resource is disabled, type the following command:**

```
# clresource disable resource
```

Note – If other resources depend on the resource that you are migrating, this step fails. In this situation, consult the error message that is printed to determine the names of the dependent resources. Then repeat this step, specifying a comma-separated list that contains the resource that you are migrating and any dependent resources.

- **If you can migrate the resource only when the resource group is unmanaged, type the following commands:**

```
# clresource disable -g resource-group +
# clresourcegroup offline resource-group
# clresourcegroup unmanage resource-group
```

The replaceable items in these commands are as follows:

resource-group Specifies the resource group that is to be unmanaged

- 3 For each resource of the resource type that is to be migrated, change the `Type_version` property to the new version.**

If necessary, set other properties of the same resource to appropriate values in the same command. To set these properties, specify the `-p` option in the command.

To determine whether you are required to set other properties, consult the instructions for upgrading the resource type. You might be required to set other properties for the following reasons:

- An extension property has been introduced in the new version of the resource type.
- The default value of an existing property has been changed in the new version of the resource type.

```
# clresource set -p Type_version=new-version \
[-p extension-property=new-value] [-p standard-property=new-value] resource
```

Note – If the existing version of the resource type does not support upgrades to the new version, this step fails.

- 4 Restore the previous state of the resource or resource group by reversing the command that you typed in [Step 2](#).**

- **If you can migrate the resource at any time, no action is required.**

Note – After migrating a resource that can be migrated at any time, the resource probe might not display the correct resource type version. In this situation, disable and re-enable the resource's fault monitor to ensure that the resource probe displays the correct resource type version.

- If you can migrate the resource only when the resource is unmonitored, type the following command:

```
# clresource monitor resource
```

- If you can migrate the resource only when the resource is offline, type the following command:

```
# clresource enable resource
```

Note – If you disabled in [Step 2](#) other resources that depend on the resource that you are migrating, enable the dependent resources also.

- If you can migrate the resource only when the resource is disabled, type the following command:

```
# clresource enable resource
```

Note – If you disabled in [Step 2](#) other resources that depend on the resource that you are migrating, enable the dependent resources also.

- If you can migrate the resource only when the resource group is unmanaged, type the following commands:

```
# clresource enable -g resource-group +  
# clresourcegroup manage resource-group  
# clresourcegroup online resource-group
```

Example 2–2 Migrating a Resource That Can Be Migrated Only When Offline

This example shows the migration of a resource that can be migrated only when the resource is offline. The new resource type package contains methods that are located in new paths. Because the methods are not overwritten during the installation, the resource does not need to be disabled until after the upgraded resource type is installed.

The characteristics of the resource in this example are as follows:

- The new resource type version is 2.0.
- The resource name is myresource.
- The resource type name is myrt.
- The new RTR file is in /opt/XYZmyrt/etc/XYZ.myrt.
- No dependencies on the resource that is to be migrated exist.
- The resource that is to be migrated can be taken offline while leaving the containing resource group online.

This example assumes that the upgrade package is already installed on all cluster nodes according to the supplier's directions.

```
# clresourcetype register -f /opt/XYZmyrt/etc/XYZ.myrt myrt
# clresource disable myresource
# clresource set -p Type_version=2.0 myresource
# clresource enable myresource
```

Example 2-3 Migrating a Resource That Can Be Migrated Only When Unmonitored

This example shows the migration of a resource that can be migrated only when the resource is unmonitored. The new resource type package contains only the monitor and RTR file. Because the monitor is overwritten during installation, monitoring of the resource must be disabled before the upgrade package is installed.

The characteristics of the resource in this example are as follows:

- The new resource type version is 2.0.
- The resource name is myresource.
- The resource type name is myrt.
- The new RTR file is in /opt/XYZmyrt/etc/XYZ.myrt.

The following operations are performed in this example.

1. Before the upgrade package is installed, the following command is run to disable monitoring of the resource:

```
# clresource unmonitor myresource
```

2. The upgrade package is installed on all cluster nodes according to the supplier's directions.
3. To register the new version of the resource type, the following command is run:

```
# clresourcetype register -f /opt/XYZmyrt/etc/XYZ.myrt myrt
```

4. To change the Type_version property to the new version, the following command is run:

```
# clresource set -p Type_version=2.0 myresource
```

5. To enable monitoring of the resource after its migration, the following command is run:

```
# clresource monitor myresource
```

Downgrading a Resource Type

You can downgrade a resource to an older version of its resource type. The conditions for downgrading a resource to an older version of the resource type are more restrictive than the conditions for upgrading to a newer version of the resource type. The resource group that contains the resource must be unmanaged.

▼ How to Downgrade a Resource to an Older Version of Its Resource Type

The instructions that follow explain how to use the `clresource(1CL)` command to perform this task. However, you are not restricted to using the `clresource` command for this task. Instead of the `clresource` command, you can use Oracle Solaris Cluster Manager or the Resource Group option of the `clsetup(1CL)` command to perform this task.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.
- 2 Switch offline the resource group that contains the resource that you are downgrading.
`clresourcegroup offline resource-group`
- 3 Disable all resources in the resource group that contains the resource that you are downgrading.
`clresource disable -g resource-group +`
- 4 Unmanage the resource group that contains the resource that you are downgrading.
`clresourcegroup unmanage resource-group`
- 5 If necessary, reregister the old version of the resource type to which you are downgrading.
Perform this step only if the version to which you are downgrading is no longer registered. If the version to which you are downgrading is still registered, omit this step.
`clresourcetype register resource-type-name`
- 6 For the resource that you are downgrading, set the `Type_version` property to old version to which you are downgrading.
If necessary, edit other properties of the same resource to appropriate values in the same command.
`clresource set -p Type_version=old-version resource-todowngrade`
- 7 Enable all the resources that you disabled in [Step 3](#).
`# clresource enable -g resource-group +`
- 8 Bring to a managed state the resource group that contains the resource that you downgraded.
`# clresourcegroup manage resource-group`
- 9 Bring online the resource group that contains the resource that you downgraded.
`# clresourcegroup online resource-group`

Creating a Resource Group

A resource group contains a set of resources, all of which are brought online or offline together on a given node or set of nodes. You must create an empty resource group before you place resources into it. A resource group can be configured to run in global-cluster non-voting nodes.

Note – The global-cluster non voting nodes that are specified in the resource group's node list do not need to exist when the resource group is created. If the node specified in the node list is not detected by the RGM, a warning message is displayed but does not result in an error.

The two resource group types are **failover** and **scalable**. A failover resource group can be online on one node only at any time, while a scalable resource group can be online on multiple nodes simultaneously.

The following procedures explain how to use the `clresourcegroup(1CL)` command to create a resource group.

For conceptual information about resource groups, see [Chapter 1, “Planning for Oracle Solaris Cluster Data Services,”](#) and *Oracle Solaris Cluster Concepts Guide*.

▼ How to Create a Failover Resource Group

A failover resource group contains the following types of resources:

- Network address resources, which are instances of the built-in resource types `LogicalHostname` and `SharedAddress`
- Failover resources, which are the data service application resources for a failover data service

The network address resources and their dependent data service resources move between cluster nodes when data services fail over or are switched over.

Note – Perform this procedure from any cluster node.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Create the failover resource group.

```
# clresourcegroup create [-n node-zone-list] resource-group
```

`-n node-zone-list` Specifies a comma-separated, ordered list of nodes that can master this resource group. The format of each entry in the list is `node:zone`. In this

format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource group is created on all nodes in the cluster.

Note – To achieve highest availability, specify global-cluster non-voting nodes on different global-cluster voting nodes in a failover resource group's node list instead of different nodes on the same global-cluster voting node.

resource-group Specifies your choice of the name of the failover resource group to add. This name must begin with an ASCII character.

3 Verify that the resource group has been created.

```
# clresourcegroup show resource-group
```

Example 2-4 Creating a Failover Resource Group

This example shows the creation of the failover resource group resource-group-1. The global cluster voting nodes phys-schost-1 and phys-schost-2 can master this resource group.

```
# clresourcegroup create -n phys-schost1,phys-schost-2 resource-group-1
# clresourcegroup show -v resource-group-1
```

```
=== Resource Groups and Resources ===
```

Resource Group:	resource-group1
RG_description:	<NULL>
RG_mode:	Failover
RG_state:	Unmanaged
RG_project_name:	default
RG_affinities:	<NULL>
RG_SLM_type:	manual
Auto_start_on_new_cluster:	True
Failback:	False
NodeList:	phys-schost-1 phys-schost-2
Maximum_primaries:	1
Desired_primaries:	1
RG_dependencies:	<NULL>
Implicit_network_dependencies:	True
Global_resources_used:	<All>
Pingpong_interval:	3600
Pathprefix:	<NULL>
RG_System:	False
Suspend_automatic_recovery:	False

Next Steps After you create a failover resource group, you can add application resources to this resource group. See [“Tools for Adding Resources to Resource Groups” on page 55](#) for the procedure.

See Also The `clresourcegroup(1CL)` man page.

▼ How to Create a Scalable Resource Group

A scalable resource group is used with scalable services. The shared address feature is the Oracle Solaris Cluster networking facility that enables the multiple instances of a scalable service to appear as a single service. You must first create a failover resource group that contains the shared addresses on which the scalable resources depend. Next, create a scalable resource group, and add scalable resources to that group. The node list of a scalable resource group or of the shared address resource group must not contain more than one global-cluster non-voting node on the same node. Each instance of the scalable service must run on a different cluster node.

You can configure a scalable resource group to run in a global-cluster non-voting node as well. Do not configure a scalable resource to run in multiple global-cluster non-voting nodes on the same node.

Note – Perform this procedure from any cluster node.

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**
- 2 **Create the failover resource group that holds the shared addresses that the scalable resource is to use.**
- 3 **Create the scalable resource group.**

```
# clresourcegroup create -S [-p Maximum primaries=m] [-p Desired primaries=n] \
[-n node-zone-list] resource-group
```

- | | |
|------------------------|--|
| -S | Specifies that the resource group is to be multi-mastered. If the <code>-p Maximum primaries</code> and <code>-p Desired primaries</code> options are omitted, both properties are set to the number of nodes in the resource group's node list. |
| -p Maximum primaries=m | Specifies the maximum number of active primaries for this resource group. |
| -p Desired primaries=n | Specifies the number of active primaries on which the resource group should attempt to start. |

-n node-zone-list

Specifies a comma-separated, ordered list of nodes in which this resource group is to be available. The format of each entry in the list is *node:zone*. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource group is created on all nodes in the cluster.

The node list of the scalable resource can contain the same list or a subset of *nodename:zonename* pairs as the node list of the shared address resource

resource-group

Specifies your choice of the name of the scalable resource group to add. This name must begin with an ASCII character.

4 Verify that the scalable resource group has been created.

```
# clresourcegroup show resource-group
```

Example 2-5 Creating a Scalable Resource Group

This example shows the creation of the scalable resource group *resource-group-1*. This resource group is to be hosted in the global cluster of nodes *phys-schost-1* and *phys-schost-2*. The scalable resource group depends on the failover resource group *resource-group-2*, which contains the shared address resources.

```
# clresourcegroup create -S \
-p Maximum_primaries=2 \
-p Desired_primaries=2 \
-p RG_dependencies=resource-group-2 \
-n phys-schost-1, phys-schost-2 \
resource-group-1

# clresourcegroup show resource-group-1

=== Resource Groups and Resources ===

Resource Group:                resource-group-1
RG_description:                <NULL>
RG_mode:                      Scalable
RG_state:                      Unmanaged
RG_project_name:               default
RG_affinities:                 <NULL>
Auto_start_on_new_cluster:     True
Failback:                     False
Nodelist:                     phys-schost-1 phys-schost-2
Maximum_primaries:             2
Desired_primaries:             2
RG_dependencies:               resource-group2
```

Implicit_network_dependencies:	True
Global_resources_used:	<All>
Pingpong_interval:	3600
Pathprefix:	<NULL>
RG_System:	False
Suspend_automatic_recovery:	False

Next Steps After you have created a scalable resource group, you can add scalable application resources to the resource group. See [“How to Add a Scalable Application Resource to a Resource Group” on page 68](#) for details.

See Also The `clresourcegroup(1CL)` man page.

Configuring Failover and Scalable Applications

After the NAS device is installed and configured, you can use the `ScalMountPoint` resource to configure failover and scalable applications.

An instance of the `ScalMountPoint` resource type represents the mount point of one of the following types of file systems:

- QFS shared file systems
- File systems on a network-attached storage (NAS) device.

The NAS device and the file system must already be configured for use with Oracle Solaris Cluster.

The `ScalMountPoint` resource type is a scalable resource type. An instance of this resource type is online on each node in the node list of the resource group that contains the resource.

▼ How to Configure a Failover Application Using the ScalMountPoint Resource

Before You Begin This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC (role-based access control) authorization.

- 1 Create a scalable resource group containing the `ScalMountPoint` resource for the NAS NFS file system.

```
# clrg create -p RG_mode=Scalable \  
-p Desired_primaries=num_active_primary \  
-p Maximum_primaries=max_num_active_primary scalmp-rg  
  
# clrt register SUNW.ScalMountPoint  
  
# clrs create -g scalmp-rg -t SUNW.ScalMountPoint \  
-p TargetFileSystem=nas_device:path \  
-p FileSystemType=nas \  
-p MountPointDir=fs_mountpoint scalmp-rs  
  
# clrg online -eM scalmp-rg
```

- 2 Create a failover resource group that contains the failover application resource.

```
# clrg create -p rg_affinities=++scalmp-rg app-fo-rg
```

The failover application resource group must have a strong positive affinity upon the resource group created in [Step 1](#).

```
# clrt register app_resource_type  
  
# clrs create -g app-fo-rg -t app_resource_type \  
-p Resource_dependencies_offline_restart=scalmp-rs \  
...  
app-fo-rs
```

The failover application resource must have an offline restart dependency upon the `ScalMountPoint` resource created in [Step 1](#).

```
# clrg online -eM app-fo-rg
```

▼ How to Configure a Scalable Application Using the `ScalMountPoint` Resource

Before You Begin This procedure provides the long forms of the Oracle Solaris Cluster commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

To perform this procedure, become superuser or assume a role that provides `solaris.cluster.read` and `solaris.cluster.modify` RBAC (role-based access control) authorization.

- 1 Create a scalable resource group containing the `ScalMountPoint` resource for the NAS NFS file system.

```
# clrg create -p RG_mode=Scalable \
-p Desired_primaries=num_active_primary \
-p Maximum_primaries=max_num_active_primary scalmp-rg

# clrt register SUNW.ScalMountPoint

# clrs create -g scalmp-rg -t SUNW.ScalMountPoint \
-p TargetFileSystem=nas_device:path \
-p FileSystemType=nas \
-p MountPointDir=fs_mountpoint scalmp-rs

# clrg online -eM scalmp-rg
```

- 2 Create a scalable resource group that contains the application resource.

```
# clrg create -p RG_mode=Scalable \
-p Maximum_primaries=max_num_active_primary \
-p Desired_primaries=num_active_primary \
-p rg_affinities=++scalmp-rg app-rg
```

The application resource group must have a strong positive affinity upon the resource group created in [Step 1](#).

```
# clrt register app_resource_type

# clrs create -g app-rg -t app_resource_type \
...
-p Scalable=True \
-p resource_dependencies_offline_restart=scalmp-rs app-rs

# clrg online -eM app-rg
```

The application resource must have an offline restart dependency upon the `ScalMountPoint` resources created in [Step 1](#).

Tools for Adding Resources to Resource Groups

A resource is an instantiation of a resource type. You must add resources to a resource group before the RGM can manage the resources. This section describes the following three resource types.

- Logical hostname resources
- Shared-address resources
- Data service (application) resources

Oracle Solaris Cluster provides the following tools for adding resources to resource groups:

- **Oracle Solaris Cluster Manager.** For more information, see the Oracle Solaris Cluster Manager online help.
- **The `clsetup(1CL)` utility.**
- **Oracle Solaris Cluster maintenance commands.**

You can use the wizards in the Oracle Solaris Cluster Manager, the `clsetup` utility, or the Oracle Solaris Cluster maintenance commands to add the logical hostname resources and shared-address resources to the resource group.

Oracle Solaris Cluster Manager and the `clsetup` utility enable you to add resources to the resource group interactively. Configuring these resources interactively reduces the possibility for configuration errors that might result from command syntax errors or omissions. Oracle Solaris Cluster Manager and the `clsetup` utility ensure that all required resources are created and that all required dependencies between resources are set.

Always add logical hostname resources and shared address resources to failover resource groups. Add data service resources for failover data services to failover resource groups. Failover resource groups contain both the logical hostname resources and the application resources for the data service. Scalable resource groups contain only the application resources for scalable services. The shared address resources on which the scalable service depends must reside in a separate failover resource group. You must specify dependencies between the scalable application resources and the shared address resources for the data service to scale across cluster nodes.

Note – The DEPRECATED flag marks the logical hostname or shared address resource as a deprecated address. These addresses are not suitable for outbound requests since they can migrate to a different cluster node due to a failover or switchover.

For more information about resources, see [Oracle Solaris Cluster Concepts Guide](#) and [Chapter 1, “Planning for Oracle Solaris Cluster Data Services.”](#)

▼ How to Add a Logical Hostname Resource to a Resource Group by Using the `clsetup` Utility

The following instructions explain how to add a logical hostname resource to a resource group by using the `clsetup` utility. Perform this procedure from one node only.

This procedure provides the long forms of the Oracle Solaris Cluster maintenance commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

Before You Begin Ensure that the following prerequisites are met:

- An entry for each logical hostname that is to be made available by the resource is added to the name service database.
- If you are using IP Networking Multipathing (IPMP) groups, the groups are configured on the nodes where the logical hostname resource can be brought online.
- Any global-cluster non-voting node that can master the resource is already configured on your cluster nodes.

Ensure that you have the following information:

- The hostnames that you plan to add to the resource group

1 Become superuser on any cluster node.

2 Start the `clsetup` utility.

```
# clsetup
```

The `clsetup` main menu is displayed.

3 Type the number that corresponds to the option for data services and press Return.

The Data Services menu is displayed.

4 Type the number that corresponds to the option for configuring the Logical Hostname resource and press Return.

The `clsetup` utility displays the list of prerequisites for performing this task.

5 Verify that the prerequisites are met, and press Return to continue.

The `clsetup` utility displays a list of the cluster nodes where the logical hostname resource can be brought online.

6 Select the nodes where the logical hostname resource can be brought online.

- To accept the default selection of all listed nodes in an arbitrary order, type `a` and press Return.
- To select a subset of the listed nodes, type a comma-separated or space-separated list of the numbers that correspond to the nodes. Then press Return.

- **To select all nodes in a particular order, type a comma-separated or space-separated ordered list of the numbers that correspond to the nodes and press Return.**

Ensure that the nodes are listed in the order in which the nodes are to appear in the logical hostname resource group's node list. The first node in the list is the primary node of this resource group.

- 7 To confirm your selection of nodes, type `d` and press Return.**

The `clsetup` utility displays a screen where you can specify the logical hostname that the resource is to make available.

- 8 Type the logical hostname that this resource is to make available and press Return.**

The `clsetup` utility displays the names of the Oracle Solaris Cluster objects that the utility will create.

- 9 If you require a different name for any Oracle Solaris Cluster object, change the name as follows.**

- a. Type the number that corresponds to the name that you are changing and press Return.**

The `clsetup` utility displays a screen where you can specify the new name.

- b. At the New Value prompt, type the new name and press Return.**

The `clsetup` utility returns you to the list of the names of the Oracle Solaris Cluster objects that the utility will create.

- 10 To confirm your selection of Oracle Solaris Cluster object names, type `d` and press Return.**

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

- 11 To create the configuration, type `c` and press Return.**

The `clsetup` utility displays a progress message to indicate that the utility is running commands to create the configuration. When configuration is complete, the `clsetup` utility displays the commands that the utility ran to create the configuration.

- 12 (Optional) Type `q` and press Return repeatedly until you quit the `clsetup` utility.**

If you prefer, you can leave the `clsetup` utility running while you perform other required tasks before using the utility again. If you choose to quit `clsetup`, the utility recognizes your existing logical hostname resource group when you restart the utility.

- 13 Verify that the logical hostname resource has been created.**

Use the `clresource(1CL)` utility for this purpose. By default, the `clsetup` utility assigns the name `node_name-rg` to the resource group.

```
# clresource show node_name-rg
```

▼ How to Add a Logical Hostname Resource to a Resource Group Using the Command-Line Interface

Note – When you add a logical hostname resource to a resource group, the extension properties of the resource are set to their default values. To specify a nondefault value, you must modify the resource after you add the resource to a resource group. For more information, see [“How to Modify a Logical Hostname Resource or a Shared Address Resource” on page 92.](#)

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the failover resource group to which you are adding the resource
- The hostnames that you plan to add to the resource group

1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

2 Add the logical hostname resource to the resource group.

```
# clreslogicalhostname create -g resource-group -h hostnamelist, ... [-N netiflist] resource
```

<code>-g resource-group</code>	Specifies the name of the resource group in which this resource resides.
<code>-h hostnamelist, ...</code>	Specifies a comma-separated list of UNIX hostnames (logical hostnames) by which clients communicate with services in the resource group. When a logical hostname resource is added to a resource group that runs in a global-cluster non-voting node, the corresponding IP addresses are configured in that node. These IP addresses are available only to applications that are running in that global-cluster non-voting node.

You must specify the fully qualified name with the `-h` option if you require a fully qualified hostname.

<code>-N netiflist</code>	Specifies an optional, comma-separated list that identifies the IPMP groups that are on each node. Each element in <i>netiflist</i> must be in the form of <i>netif@node</i> . <i>netif</i> can be given as an IPMP group name, such as <code>sc_ipmp0</code> . The node can be identified by the node name or node ID, such as <code>sc_ipmp0@1</code> or <code>sc_ipmp@phys-schost-1</code>
---------------------------	---

Note – Oracle Solaris Cluster does not support the use of the adapter name for `netif`.

resource Specifies an optional resource name of your choice. You cannot use the fully qualified name in the resource name.

3 Verify that the logical hostname resource has been added.

```
# clresource show resource
```

Example 2–6 Adding a Logical Hostname Resource to a Resource Group

This example shows the addition of logical hostname resource (`resource-1`) to a resource group (`resource-group-1`).

```
# clreslogicalhostname create -g resource-group-1 -h schost-1 resource-1
# clresource show resource-1
```

```
=== Resources ===
```

Resource:	resource-1
Type:	SUNW.LogicalHostname:2
Type_version:	2
Group:	resource-group-1
R_description:	
Resource_project_name:	default
Enabled{phats1}:	True
Enabled{phats2}:	True
Monitored{phats1}:	True
Monitored{phats2}:	True

Example 2–7 Adding Logical Hostname Resources That Identify IPMP Groups

This example shows the addition of the following logical host name resources to the resource group `nfs-fo-rg`:

- A resource that is named `cs23-rs`, which identifies the IPMP group `sc_ipmp0` on node 1 and node 2
- A resource that is named `cs24-rs`, which identifies the IPMP group `sc_ipmp1` on node 1 and node 2

```
# clreslogicalhostname create -g nfs-fo-rg -h cs23-rs -N sc_ipmp0@1,sc_ipmp0@2 cs23-rs
# clreslogicalhostname create -g nfs-fo-rg -h cs24-rs -N sc_ipmp1@1,sc_ipmp1@2 cs24-rs
```

Next Steps After you add logical hostname resources, see [“How to Bring Resource Groups Online” on page 71](#) to bring the resources online.

Troubleshooting Adding a resource causes the Oracle Solaris Cluster software to validate the resource. If the validation fails, the `clreslogicalhostname` command prints an error message and exits. To determine why the validation failed, check the `syslog` on each node for an error message. The message appears on the node that performed the validation, not necessarily the node on which you ran the `clreslogicalhostname` command.

See Also The `clreslogicalhostname(1CL)` man page.

▼ How to Add a Shared Address Resource to a Resource Group by Using the `clsetup` Utility

The following instructions explain how to add a shared address resource to a resource group by using the `clsetup` utility. Perform this procedure from any cluster node.

This procedure provides the long forms of the Oracle Solaris Cluster maintenance commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

Before You Begin Ensure that the following prerequisites are met:

- The shared address that is to be made available by the resource has an entry in a name service database.
- If you are using IP Networking Multipathing (IPMP) groups, the groups are configured on the nodes where the shared address resource can be brought online.
- Any global-cluster non-voting node that can master the resource is already configured on your cluster nodes.

Ensure that you have the following information:

- The hostnames that you plan to add to the resource group.

1 Become superuser on any cluster node.

2 Start the `clsetup` utility.

```
# clsetup
```

The `clsetup` main menu is displayed.

3 Type the number that corresponds to the option for data services and press Return.

The Data Services menu is displayed.

- 4 Type the number that corresponds to the option for configuring the shared address resource and press Return.**

The `clsetup` utility displays the list of prerequisites for performing this task.

- 5 Verify that the prerequisites are met, and press Return to continue.**

The `clsetup` utility displays a list of the cluster nodes where the shared address resource can be brought online.

- 6 Select the nodes where the shared address resource can be brought online.**

- To accept the default selection of all listed nodes in an arbitrary order, type `a` and press Return.
- To select a subset of the listed nodes, type a comma-separated or space-separated list of the numbers that correspond to the nodes. Then press Return.
- To select all nodes in a particular order, type a comma-separated or space-separated ordered list of the numbers that correspond to the nodes and press Return.

- 7 To confirm your selection of nodes, type `d` and press Return.**

The `clsetup` utility displays a screen where you can specify the shared address that the resource is to make available.

- 8 Type the shared address that this resource is to make available and press Return.**

The `clsetup` utility displays the names of the Oracle Solaris Cluster objects that the utility will create.

- 9 If you require a different name for any Oracle Solaris Cluster object, change the name as follows.**

- a. Type the number that corresponds to the name that you are changing and press Return.**

The `clsetup` utility displays a screen where you can specify the new name.

- b. At the New Value prompt, type the new name and press Return.**

The `clsetup` utility returns you to the list of the names of the Oracle Solaris Cluster objects that the utility will create.

- 10 To confirm your selection of Oracle Solaris Cluster object names, type `d` and press Return.**

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

11 To create the configuration, type `c` and Press Return.

The `clsetup` utility displays a progress message to indicate that the utility is running commands to create the configuration. When configuration is complete, the `clsetup` utility displays the commands that the utility ran to create the configuration.

12 (Optional) Type `q` and press Return repeatedly until you quit the `clsetup` utility.

If you prefer, you can leave the `clsetup` utility running while you perform other required tasks before using the utility again. If you choose to quit `clsetup`, the utility recognizes your existing shared address resource group when you restart the utility.

13 Verify that the shared address resource has been created.

Use the `clresource(1CL)` utility for this purpose. By default, the `clsetup` utility assigns the name `node_name-rg` to the resource group.

```
# clresource show node_name-rg
```

▼ How to Add a Shared Address Resource to a Resource Group Using the Command-Line Interface

Note – When you add a shared address resource to a resource group, the extension properties of the resource are set to their default values. To specify a nondefault value, you must modify the resource after you add the resource to a resource group. For more information, see [“How to Modify a Logical Hostname Resource or a Shared Address Resource” on page 92.](#)

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the resource group into which you are adding the resource. This group must be a failover resource group that you created previously.
- The hostnames that you plan to add to the resource group.

1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**2 Add the shared address resource to the resource group.**

```
# clressharedaddress create -g resource-group -h hostnamelist, ... \
[-X auxnodelist] [-N netiflist] resource
```

- g *resource-group*** Specifies the resource group name. In the node list of a shared address resource, do not specify more than one global-cluster non-voting node on the same global-cluster voting node. Specify the same list of *nodename:zonename* pairs as the node list of the scalable resource group.
- h *hostnamelist, ...*** Specifies a comma-separated list of shared address hostnames.
- X *auxnodelist*** Specifies a comma-separated list of node names or IDs that identify the cluster nodes that can host the shared address but never serve as primary if failover occurs. These nodes are mutually exclusive, with the nodes identified as potential masters in the resource group's node list. If no auxiliary node list is explicitly specified, the list defaults to the list of all cluster node names that are not included in the node list of the resource group that contains the shared address resource.

Note – To ensure that a scalable service runs in all global-cluster non-voting nodes that were created to master the service, the complete list of nodes must be included in the node list of the shared address resource group or the `auxnodelist` of the shared address resource. If all the nodes are listed in the node list, the `auxnodelist` can be omitted.

- N *netiflist*** Specifies an optional, comma-separated list that identifies the IPMP groups that are on each node. Each element in *netiflist* must be in the form of *netif@node*. *netif* can be given as an IPMP group name, such as `sc_ipmp0`. The node can be identified by the node name or node ID, such as `sc_ipmp0@1` or `sc_ipmp@phys-schost-1`.

Note – Oracle Solaris Cluster does not support the use of the adapter name for *netif*.

- resource*** Specifies an optional resource name of your choice.

3 Verify that the shared address resource has been added and validated.

```
# clresource show resource
```

Example 2–8 Adding a Shared Address Resource to a Resource Group

This example shows the addition of a shared address resource (`resource-1`) to a resource group (`resource-group-1`).


```
# clressharedaddress create -g resource-group-1 -h schost-1 resource-1
# clresource show resource-1

=== Resources ===

Resource:                resource-1
Type:                    SUNW.SharedAddress:2
Type_version:            2
Group:                   resource-group-1
R_description:
Resource_project_name:   default
Enabled{phats1}:         False
Enabled{phats2}:         False
Monitored{phats1}:       True
Monitored{phats2}:       True
```

Next Steps After you add a shared address resource, use the procedure “[How to Bring Resource Groups Online](#)” on page 71 to enable the resource.

Troubleshooting Adding a resource causes the Oracle Solaris Cluster software to validate the resource. If the validation fails, the `clressharedaddress` command prints an error message and exits. To determine why the validation failed, check the `syslog` on each node for an error message. The message appears on the node that performed the validation, not necessarily the node on which you ran the `clressharedaddress` command.

See Also The `clressharedaddress(1CL)` man page.

▼ How to Add a Failover Application Resource to a Resource Group

A failover application resource is an application resource that uses logical hostnames that you previously created in a failover resource group.

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the failover resource group to which you are adding the resource
- The name of the resource type for the resource
- The logical hostname resources that the application resource uses, which are the logical hostnames that you previously included in the same resource group

Note – This procedure also applies to proxy resources.

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**

- 2 **Add a failover application resource to the resource group.**

```
# clresource create -g resource-group -t resource-type \  
[-p "extension-property[{node-specifier}]"=value, ...] [-p standard-property=value, ...] resource  
-g resource-group
```

Specifies the name of a failover resource group. This resource group must already exist.

```
-t resource-type
```

Specifies the name of the resource type for the resource.

```
-p "extension-property[{node-specifier}]"=value, ...
```

Specifies a comma-separated list of extension properties that you are setting for the resource. The extension properties that you can set depend on the resource type. To determine which extension properties to set, see the documentation for the resource type.

node-specifier is an *optional* qualifier to the `-p` and `-x` options. This qualifier indicates that the extension property or properties on *only* the specified node or nodes are to be set when the resource is created. The specified extension properties on other nodes in the cluster are not set. If you do not include *node-specifier*, the specified extension properties on all nodes in the cluster are set. You can specify a node name or a node identifier for *node-specifier*. Examples of the syntax of *node-specifier* include the following:

```
-p "myprop{phys-schost-1}"
```

The braces ({}) indicate that you are setting the specified extension property on only node `phys-schost-1`. For most shells, the double quotation marks (") are required.

You can also use the following syntax to set an extension property in two different global-cluster voting nodes on two different nodes:

```
-x "myprop{phys-schost-1:zoneA,phys-schost-2:zoneB}"
```

Note – The extension property that you specify with *node-specifier* must be declared in the RTR file as a per-node property.

```
-p standard-property=value, ...
```

Specifies a comma-separated list of standard properties that you are setting for the resource. The standard properties that you can set depend on the resource type. To determine which standard properties to set, see the following man pages: `cluster(1CL)`, `rt_properties(5)`, `rg_properties(5)`, `r_properties(5)`, and `property_attributes(5)`.

resource

Specifies your choice of the name of the resource to add.

The resource is created in the enabled state.

3 Verify that the failover application resource has been added and validated.

```
# clresource show resource
```

Example 2–9 Adding a Failover Application Resource to a Resource Group

This example shows the addition of a resource (resource-1) to a resource group (resource-group-1). The resource depends on logical hostname resources (schost-1, schost-2), which must reside in the same failover resource groups that you defined previously.

```
# clresource create -g resource-group-1 -t resource-type-1 \
-p Resource_dependencies=schost-1,schost2 resource-1 \
# clresource show resource-1
```

```
=== Resources ===
```

Resource:	resource-1
Type:	resource-type-1
Type_version:	
Group:	resource-group-1
R_description:	
Resource_project_name:	default
Enabled{phats1}:	False
Enabled{phats2}:	False
Monitored{phats1}:	True
Monitored{phats2}:	True

Next Steps After you add a failover application resource, use the procedure [“How to Bring Resource Groups Online” on page 71](#) to enable the resource.

Troubleshooting Adding a resource causes the Oracle Solaris Cluster software to validate the resource. If the validation fails, the `clresource` command prints an error message and exits. To determine why the validation failed, check the `syslog` on each node for an error message. The message appears on the node that performed the validation, not necessarily the node on which you ran the `clresource` command.

See Also The `clresource(1CL)` man page.

▼ How to Add a Scalable Application Resource to a Resource Group

A scalable application resource is an application resource that uses the network load balancing features of Oracle Solaris Cluster software. The scalable application resource is in a multi-mastered resource group and has a dependency on one or more shared-address resources. The shared-address resources are in a failover resource group.

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the scalable resource group to which you are adding the resource
- The name of the resource type for the resource
- The shared address resources that the scalable service resource uses, which are the shared addresses that you previously included in a failover resource group

Note – This procedure also applies to proxy resources.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Add a scalable application resource to the resource group.

```
# clresource create -S -g resource-group -t resource-type \
-p Resource_dependencies=network-resource[,network-resource...] \
-p Scalable=True
[-p "extension-property[{node-specifier}]"=value, ...] [-p standard-property=value, ...] resource
-S
```

Specifies that the resource group is to be multi-mastered. If the `-p Maximum primaries` and `-p Desired primaries` options are omitted, both properties are set to the number of nodes in the resource group's node list.

`-g resource-group`
Specifies the name of a scalable service resource group that you previously created.

`-t resource-type`
Specifies the name of the resource type for this resource.

`-p Resource_dependencies= network-resource[,network-resource...]`
Specifies the list of network resources (shared addresses) on which this resource depends.

`-p Scalable=True`
Specifies that this resource uses the network load balancing feature of Oracle Solaris Cluster software.

-p "*extension-property*[*{node-specifier}*]"=*value*, ...

Specifies a comma-separated list of extension properties that you are setting for the resource. The extension properties that you can set depend on the resource type. To determine which extension properties to set, see the documentation for the resource type.

node-specifier is an *optional* qualifier to the **-p** and **-x** options. This qualifier indicates that the extension property or properties on *only* the specified node or nodes are to be set when the resource is created. The specified extension properties on other nodes in the cluster are not set. If you do not include *node-specifier*, the specified extension properties on all nodes in the cluster are set. You can specify a node name or a node identifier for *node-specifier*. Examples of the syntax of *node-specifier* include the following:

-p "myprop{phys-schost-1}"

The braces ({}) indicate that you are setting the specified extension property on only node phys-schost-1. For most shells, the double quotation marks (") are required.

You can also use the following syntax to set an extension property in two different global-cluster voting nodes on two different nodes:

-x "myprop{phys-schost-1:zoneA,phys-schost-2:zoneB}"

Note – The extension property that you specify with *node-specifier* must be declared in the RTR file as a per-node property.

-p *standard-property*=*value*, ...

Specifies a comma-separated list of standard properties that you are setting for the resource. The standard properties that you can set depend on the resource type. For scalable services, you typically set the `Port_list`, `Load_balancing_weights`, and `Load_balancing_policy` properties. To determine which standard properties to set, see the following man pages: [cluster\(1CL\)](#), [rt_properties\(5\)](#), [rg_properties\(5\)](#), [r_properties\(5\)](#), and [property_attributes\(5\)](#).

resource

Specifies your choice of the name of the resource to add.

The resource is created in the enabled state.

3 Verify that the scalable application resource has been added and validated.

```
# clresource show resource
```

Example 2–10 Adding a Scalable Application Resource to a Resource Group

This example shows the addition of a resource (resource-1) to a resource group (resource-group-1). Note that resource-group-1 depends on the failover resource group that contains the network addresses that are in use (schost-1 and schost-2 in the following

example). The resource depends on shared address resources (schost-1, schost-2), which must reside in one or more failover resource groups that you defined previously.

```
# clresource create -S -g resource-group-1 -t resource-type-1 \  
-p Resource_dependencies=schost-1,schost-2 resource-1 \  
-p Scalable=True  
# clresource show resource-1  
  
=== Resources ===  
  
Resource:                resource-1  
Type:                    resource-type-1  
Type_version:  
Group:                   resource-group-1  
R_description:  
Resource_project_name:   default  
Enabled{phats1}:         False  
Enabled{phats2}:         False  
Monitored{phats1}:      True  
Monitored{phats2}:      True
```

Next Steps After you add a scalable application resource, follow the procedure [“How to Bring Resource Groups Online” on page 71](#) to enable the resource.

Troubleshooting Adding a resource causes the Oracle Solaris Cluster software to validate the resource. If the validation fails, the `clresource` command prints an error message and exits. To determine why the validation failed, check the `syslog` on each node for an error message. The message appears on the node that performed the validation, not necessarily the node on which you ran the `clresource` command.

See Also The `clresource(1CL)` man page.

Bringing Resource Groups Online

To enable resources to begin providing HA services, you must perform the following operations:

- Bring resource groups online
- Enable the resources in their resource groups
- Enable the resource monitors
- Make the resource groups managed

You can perform these tasks individually or by using a single command.

After you bring a resource group online, it is configured and ready for use. If a resource or node fails, the RGM switches the resource group online on alternate nodes to maintain availability of the resource group.

▼ How to Bring Resource Groups Online

Perform this task from any cluster node.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.
- 2 Type the command to bring online the resource groups.
 - If you have intentionally disabled a resource or a fault monitor that must remain disabled, type the following command:


```
# clresourcegroup online rg-list
```

rg-list Specifies a comma-separated list of the names of the resource groups to bring online. The resource groups must exist. The list may contain one resource group name or more than one resource group name.
 - If you require the resources and their fault monitors to be enabled when the resource groups are brought online, type the following command:


```
# clresourcegroup online -emM rg-list
```

rg-list Specifies a comma-separated list of the names of the resource groups to bring online. The resource groups must exist. The list can contain one resource group name or more than one resource group name.

Note – If any resource group that you are bringing online declares a strong affinity for other resource groups, this operation might fail. For more information, see [“Distributing Online Resource Groups Among Cluster Nodes” on page 151](#).

- 3 Verify that each resource group that you specified in [Step 2](#) is online.

The output from this command indicates on which nodes each resource group is online.

```
# clresourcegroup status
```

Example 2–11 Bringing Online a Resource Group

This example shows how to bring online the resource group `resource-group-1` and verify its status. All resources in this resource and their fault monitors are also enabled.

```
# clresourcegroup online -emM resource-group-1
# clresourcegroup status
```

Next Steps If you brought resource groups online *without* enabling their resources and fault monitors, enable the fault monitors of any resources that you require to be enabled. For more information, see [“How to Enable a Resource Fault Monitor” on page 79](#).

See Also The `clresourcegroup(1CL)` man page.

Switching Resource Groups to Preferred Primaries

The `clresourcegroup remaster` command can switch the resource group offline from current primaries in order to switch the resource group online on preferred nodes or zones. Similar to the `clresourcegroup online` command, the RGM computes the most preferred set of nodes or zones based on the resource group's `NodeList` property, but also takes into account the following factors:

- The `RG_affinities` property settings
- The `Load_factors` property settings, compared against node load limits and current loads
- The failure history of each resource group

▼ How to Switch Resource Groups to Preferred Primaries

Perform this task from any cluster node.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.
- 2 Type the command to switch the resource groups to their preferred primaries.
 - If you have intentionally disabled a resource or a fault monitor that must remain disabled, type the following command:

```
# clresourcegroup remaster rg-list
```

rg-list Specifies a comma-separated list of the names of the resource groups to switch to the preferred primaries. The resource groups must exist. The list may contain one resource group name or more than one resource group name.
 - If you require the resources and their fault monitors to be enabled when the resource groups are switched to the preferred primary, type the following command:

```
# clresourcegroup remaster -emM rg-list
```

rg-list Specifies a comma-separated list of the names of the resource groups to switch to the preferred primaries. The resource groups must exist. The list can contain one

resource group name or more than one resource group name.

- 3 **Verify that each resource group that you specified in [Step 2](#) has been switched to the preferred primary.**

The output from this command indicates the new nodes for the resource groups that were switched.

```
# clresourcegroup status
```

Enabling a Resource

You can enable a resource that you neglected to enable when you brought online a resource group.

▼ How to Enable a Resource

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have created and have the name of the resource that you intend to enable.

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.**

- 2 **Enable the resource.**

```
# clresource enable [-n node-zone-list] resource
```

-n node-zone-list Specifies a comma-separated, ordered list of nodes on which to enable the resource. If you specify a global-cluster non-voting node, the format of each entry in the list is *node:zone*. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource is enabled on all nodes in its resource group's node list.

Note – If you specify more than one node with the *-n* option, you can specify only one resource.

resource Specifies the name of the resource that you want to enable.

3 Verify that the resource has been enabled.

clresource status

The output from this command indicates the state of the resource that you have enabled.

See Also The `clresource(1CL)` man page.

Quiescing Resource Groups

To stop a resource group from continuously switching from one node to another when a START or STOP method fails, bring it to a quiescent state. To bring a resource group to a quiescent state, you issue the `clresourcegroup quiesce` command.

When you quiesce a resource group, resource methods that are executing are allowed to run until they are completed. If a serious problem occurs, you might need to quiesce a resource group immediately. To do so, you specify the `-k` command option, which kills the following methods:

- `Prenet_start`
- `Start`
- `Monitor_start`
- `Monitor_stop`
- `Stop`
- `Postnet_stop`

Note – The `Init`, `Finis Boot`, and `Update` methods are not killed when you specify this command option.

However, if you immediately quiesce a resource group by killing methods, you might leave one of its resources in an error state such as `Start_failed` or `Stop_failed`. You must clear these error states yourself.

▼ How to Quiesce a Resource Group

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**

2 Quiesce the resource group.

```
# clresourcegroup quiesce resource-group
```

▼ How to Quiesce a Resource Group Immediately

1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

2 Immediately quiesce the resource group.

```
# clresourcegroup quiesce -k resource-group
```

The `Prenet_start`, `Start`, `Monitor_start`, `Monitor_stop`, `Stop`, and `Postnet_stop` methods that are associated with the resource group are killed immediately. The resource group is brought to a quiescent state.

The `clresourcegroup quiesce -k` command blocks until the specified resource group has reached a quiescent state.

Suspending and Resuming the Automatic Recovery Actions of Resource Groups

You can temporarily suspend the automatic recovery actions of a resource group. You might need to suspend the automatic recovery of a resource group to investigate and fix a problem in the cluster. Or, you might need to perform maintenance on resource group services.

To suspend the automatic recovery actions of a resource group, you issue the `clresourcegroup suspend` command. To resume automatic recovery actions, you issue the `clresourcegroup resume` command.

When you suspend the automatic recovery actions of a resource group, you also bring the resource group to a quiescent state.

A suspended resource group is *not* automatically restarted or failed over until you explicitly issue the command that resumes automatic recovery. Whether online or offline, suspended data services remain in their current state. You can still manually switch the resource group to a different state on specified nodes. You can also still enable or disable individual resources in the resource group.

A dependency or affinity is suspended and not enforced when you suspend the automatic recovery actions of a resource group that does one of the following:

- Contains a resource that has a restart dependency on another resource

- Declares a strong positive or negative affinity for another resource group

When you suspend one of these categories of resource groups, Oracle Solaris Cluster displays a warning that the dependency or affinity is suspended as well.

Note – Setting the `RG_system` property does not affect your ability to suspend or resume the automatic recovery actions of a resource group. However, if you suspend a resource group for which the `RG_system` property is set to `TRUE`, a warning message is produced. The `RG_system` property specifies that a resource group contains critical system services. If set to `TRUE`, the `RG_system` property prevents users from inadvertently stopping, deleting, or modifying a resource group or its resources.

Immediately Suspending Automatic Recovery by Killing Methods

When you suspend the automatic recovery actions of a resource group, resource methods that are executing are allowed to run until they are completed. If a serious problem occurs, you might need to suspend the automatic recovery actions of a resource group immediately. To do so, you specify the `-k` command option, which kills the following methods:

- `Prenet_start`
- `Start`
- `Monitor_start`
- `Monitor_stop`
- `Stop`
- `Postnet_stop`

Note – The `Init`, `Finis Boot`, and `Update` methods are not killed when you include this command option.

However, if you immediately suspend automatic recovery actions by killing methods, you might leave one of its resources in an error state such as `Start_failed` or `Stop_failed`. You must clear these error states yourself.

▼ How to Suspend the Automatic Recovery Actions of a Resource Group

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Suspend the automatic recovery actions of the resource group.

```
# clresourcegroup suspend resource-group
```

The resource group that you specify is not automatically started, restarted, or failed over until you resume automatic recovery actions. See [“How to Resume the Automatic Recovery Actions of a Resource Group” on page 77](#).

▼ How to Suspend the Automatic Recovery Actions of a Resource Group Immediately

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Immediately suspend the automatic recovery actions of the resource group.

```
# clresourcegroup suspend -k resource-group
```

The `Prenet_start`, `Start`, `Monitor_start`, `Monitor_stop`, `Stop`, and `Postnet_stop` methods that are associated with the resource group are killed immediately. Automatic recovery actions of the resource group is suspended. The resource group is *not* automatically started, restarted, or failed over until you resume automatic recovery actions. See [“How to Resume the Automatic Recovery Actions of a Resource Group” on page 77](#).

The `clresourcegroup suspend -k` command blocks until the specified resource group has reached a quiescent state.

▼ How to Resume the Automatic Recovery Actions of a Resource Group

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Resume the automatic recovery actions of the resource group.

```
# clresourcegroup resume resource-group
```

The resource group that you specify is automatically started, restarted, or failed over.

Disabling and Enabling Resource Monitors

The procedures in this section explain how to disable or enable resource fault monitors, not the resources themselves. A resource can continue to operate normally while its fault monitor is disabled. However, if the fault monitor is disabled and a data service fault occurs, automatic fault recovery is not initiated.

See the `clresource(1CL)` man page for additional information.

Note – Perform these procedures from any cluster node.

▼ How to Disable a Resource Fault Monitor

- 1 On any cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Disable the resource fault monitor.

```
# clresource unmonitor [-n node-zone-list] resource
```

`-n node-zone-list` Specifies a comma-separated, ordered list of nodes on which to unmonitor the resource. If you specify a global-cluster non-voting node, the format of each entry in the list is `node:zone`. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource is unmonitored on all nodes in its resource group's node list.

Note – If you specify more than one node with the `-n` option, you can specify only one resource.

resource Specifies the name of the resource or resources.

- 3 Run the `clresource` command on each cluster node and check for monitored fields (RS Monitored) to verify that the resource fault monitor has been disabled.

```
# clresource show -v
```

Example 2–12 Disabling a Resource Fault Monitor

```
# clresource unmonitor resource-1
# clresource show -v
...
RS Monitored: no...
```

▼ How to Enable a Resource Fault Monitor

- 1 On any cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Enable the resource fault monitor.

```
# clresource monitor [-n node-zone-list] resource
```

`-n node-zone-list` Specifies a comma-separated, ordered list of nodes on which to monitor the resource. If you specify a global-cluster non-voting node, the format of each entry in the list is `node:zone`. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global cluster, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource is monitored on all nodes in its resource group's node list.

Note – If you specify more than one node with the `-n` option, you can specify only one resource.

resource Specifies the name of the resource or resources.

- 3 Run the `clresource` command on each cluster node and check for monitored fields (RS Monitored) to verify that the resource fault monitor has been enabled.

```
# clresource show -v
```

Example 2–13 Enabling a Resource Fault Monitor

```
# clresource monitor resource-1
# clresource show -v
...
RS Monitored: yes...
```

Removing Resource Types

You do not need to remove resource types that are not in use. However, if you want to remove a resource type, follow this procedure.

Note – Perform this procedure from any cluster node.

▼ How to Remove a Resource Type

Removing a resource type involves disabling and removing all resources of that type in the cluster before unregistering the resource type.

Before You Begin To identify all instances of the resource type that you are removing, type the following command:

```
# clresourcetype show -v
```

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**

- 2 **Disable each resource of the resource type that you are removing.**

```
# clresource disable resource
```

resource Specifies the name of the resource to disable.

- 3 **Remove each resource of the resource type that you are removing.**

```
# clresource delete resource
```

resource Specifies the name of the resource to remove.

- 4 **Unregister the resource type.**

```
# clresourcetype unregister resource-type
```

resource-type Specifies the name of the resource type to unregister.

- 5 **Verify that the resource type has been removed.**

```
# clresourcetype show
```


Example 2–14 Removing a Resource Type

This example shows how to disable and remove all of the resources of a resource type (resource-type-1) and then unregister the resource type. In this example, resource-1 is a resource of the resource type resource-type-1.

```
# clresource disable resource-1
# clresource delete resource-1
# clresourcetype unregister resource-type-1
```

See Also The following man pages:

- [clresource\(1CL\)](#)
- [clresourcetype\(1CL\)](#)

Removing Resource Groups

To remove a resource group, you must first remove all of the resources from the resource group.

Note – Perform this procedure from any cluster node.

▼ How to Remove a Resource Group

Before You Begin To identify all resources in the resource group that you are removing, type the following command:

```
# clresource show -v
```

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Run the following command to switch the resource group offline.

```
# clresourcegroup offline resource-group
```

resource-group Specifies the name of the resource group to take offline.

- 3 Disable all of the resources in the resource group that you are removing.

```
# clresource disable resource
```

resource Specifies the name of the resource to disable.

4 Remove all of the resources from the resource group.

For each resource, type the following command.

```
# clresource delete resource
```

resource Specifies the name of the resource to be removed.

5 Remove the resource group.

```
# clresourcegroup delete resource-group
```

resource-group Specifies the name of the resource group to be removed.

6 Verify that the resource group has been removed.

```
# clresourcegroup show
```

Example 2–15 Removing a Resource Group

This example shows how to remove a resource group (*resource-group-1*) after you have removed its resource (*resource-1*).

```
# clresourcegroup offline resource-group-1
# clresource disable resource-1
# clresource delete resource-1
# clresourcegroup delete resource-group-1
```

See Also The following man pages:

- [clresource\(1CL\)](#)
- [clresourcegroup\(1CL\)](#)

Removing Resources

Disable the resource before you remove it from a resource group.

Note – Perform this procedure from any cluster node.

▼ How to Remove a Resource

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**

2 Disable the resource that you are removing.

```
# clresource disable resource
```

resource Specifies the name of the resource to disable.

3 Remove the resource.

```
# clresource delete resource
```

resource Specifies the name of the resource to remove.

4 Verify that the resource has been removed.

```
# clresource show
```

Example 2–16 Removing a Resource

This example shows how to disable and remove a resource (resource-1).

```
# clresource disable resource-1  
# clresource delete resource-1
```

See Also [clresource\(1CL\)](#)

Switching the Current Primary of a Resource Group

Use the following procedure to switch over a resource group from its current primary to another node that is to become the new primary.

▼ How to Switch the Current Primary of a Resource Group

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that the following conditions are met:

- You have the following information:
 - The name of the resource group that you are switching over
 - The names of the nodes where the resource group is to be brought online or to remain online

- The nodes where the resource group is to be brought online or to remain online are in the cluster.
- These nodes have been set up to be potential masters of the resource group that you are switching.

To see a list of potential primaries for the resource group, type the following command:

```
# clresourcegroup show -v
```

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**
- 2 **Switch the resource group to a new set of primaries.**

```
# clresourcegroup switch [-n node-zone-list] resource-group
```

-n node-zone-list Specifies a comma-separated, ordered list of global-cluster non-voting nodes that can master this resource group. The resource group is switched offline on all of the other nodes. The format of each entry in the list is *node:zone*. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource group is switched on all nodes in the resource group's node list.

resource-group Specifies the name of the resource group to switch.

Note – If any resource group that you are switching declares a strong affinity for other resource groups, the attempt to switch might fail or be delegated. For more information, see [“Distributing Online Resource Groups Among Cluster Nodes” on page 151](#).

- 3 **Verify that the resource group has been switched to the new primary.**

The output from this command indicates the state of the resource group that has been switched over.

```
# clresourcegroup status
```

Example 2–17 Switching a Resource Group to a New Primary

This example shows how to switch the resource group `resource-group-1` from its current primary `phys-schost-1` to the potential primary `phys-schost-2`.

1. To verify that the resource group is online on `phys-schost-1`, the following command is run:

```
phys-schost-1# clresourcegroup status
```

```
=== Cluster Resource Groups ===
```

Group Name -----	Node Name -----	Suspended -----	Status -----
resource-group1	phys-schost-1	No	Online
	phys-schost-2	No	Offline

2. To perform the switch, the following command is run:

```
phys-schost-1# clresourcegroup switch -n phys-schost-2 resource-group-1
```

3. To verify that the group is switched to be online on phys-schost-2, the following command is run:

```
phys-schost-1# clresourcegroup status
```

```
=== Cluster Resource Groups ===
```

Group Name -----	Node Name -----	Suspended -----	Status -----
resource-group1	phys-schost-1	No	Offline
	phys-schost-2	No	Online

See Also The [clresourcegroup\(1CL\)](#) page.

Disabling Resources and Moving Their Resource Group Into the UNMANAGED State

At times, you must bring a resource group into the UNMANAGED state before you perform an administrative procedure on it. Before you move a resource group into the UNMANAGED state, you must disable all of the resources that are part of the resource group and bring the resource group offline.

See the [clresourcegroup\(1CL\)](#) man page for additional information.

Note – Perform this procedure from any cluster node.

▼ How to Disable a Resource and Move Its Resource Group Into the UNMANAGED State

Note – When a shared address resource is disabled, the resource might still be able to respond to [ping\(1M\)](#) commands from some hosts. To ensure that a disabled shared address resource cannot respond to ping commands, you must bring the resource's resource group to the UNMANAGED state.

Before You Begin Ensure that you have the following information.

- The name of each resource to be disabled
- The name of the resource group to move into the UNMANAGED state

To determine the resource and resource group names that you need for this procedure, type:

```
# clresourcegroup show -v
```

- 1 On any cluster member, become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.
- 2 Disable all resources in the resource group.

```
# clresource disable [-n node-zone-list] -g resource-group +
```

`-n node-zone-list` Specifies a comma-separated, ordered list of nodes on which to disable the resource. If you specify a global-cluster non-voting node, the format of each entry in the list is `node:zone`. In this format, `node` specifies the node name and `zone` specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only `node`.

This list is optional. If you omit this list, the resource is disabled on all nodes in its resource group's node list.

Note – If you specify more than one node with the `-n` option, you can specify only one resource.

- 3 Switch the resource group offline.

```
# clresourcegroup offline resource-group
```

`resource-group` Specifies the name of the resource group to take offline.

4 Move the resource group into the UNMANAGED state.

```
# clresourcegroup unmanage resource-group
```

resource-group Specifies the name of the resource group to move into the UNMANAGED state.

5 Verify that the resources are disabled and that the resource group is in the UNMANAGED state.

```
# clresourcegroup show resource-group
```

Example 2–18 Disabling a Resource and Moving Its Resource Group Into the UNMANAGED State

This example shows how to disable the resource (resource-1) and then move the resource group (resource-group-1) into the UNMANAGED state.

```
# clresource disable resource-1
# clresourcegroup offline resource-group-1
# clresourcegroup unmanage resource-group-1
# clresourcegroup show resource-group-1
```

```
=== Resource Groups and Resources ===
```

Resource Group:	resource-group-1
RG_description:	<NULL>
RG_mode:	Failover
RG_state:	Unmanaged
Failback:	False
Nodelist:	phys-schost-1 phys-schost-2

```
--- Resources for Group resource-group-1 ---
```

Resource:	resource-1
Type:	SUNW.LogicalHostname:2
Type_version:	2
Group:	resource-group-1
R_description:	
Resource_project_name:	default
Enabled{phys-schost-1}:	False
Enabled{phys-schost-2}:	False
Monitored{phys-schost-1}:	True
Monitored{phys-schost-2}:	True

See Also The following man pages:

- [clresource\(1CL\)](#)
- [clresourcegroup\(1CL\)](#)

Displaying Resource Type, Resource Group, and Resource Configuration Information

Before you perform administrative procedures on resources, resource groups, or resource types, view the current configuration settings for these objects.

Note – You can view configuration settings for resources, resource groups, and resource types from any cluster node.

You can also use the `clresourcetype`, `clresourcegroup`, and `clresource` commands to check status information about specific resource types, resource groups, and resources. For example, the following command specifies that you want to view specific information about the resource `apache-1` only.

```
# clresource show apache-1
```

For more information, see the following man pages:

- [clresourcetype\(1CL\)](#)
- [clresourcegroup\(1CL\)](#)
- [clresource\(1CL\)](#)

Changing Resource Type, Resource Group, and Resource Properties

Oracle Solaris Cluster defines standard properties for configuring resource types, resource groups, and resources. These standard properties are described in the following sections:

- [“Resource Type Properties” on page 27](#)
- [“Resource Properties” on page 28](#)
- [“Resource Group Properties” on page 29](#)

Resources also have extension properties, which are predefined for the data service that represents the resource. For a description of the extension properties of a data service, see the documentation for the data service.

To determine whether you can change a property, see the Tunable entry for the property in the description of the property.

The following procedures describe how to change properties for configuring resource types, resource groups, and resources.

▼ How to Change Resource Type Properties

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the resource type to change.
- The name of the resource type property to change. For resource types, you can change only certain properties. To determine whether you can change a property, see the Tunable entry for the property in the [rt_properties\(5\)](#) man page.

Note – You cannot change the `Installed_nodes` property explicitly. To change this property, specify the `-n installed-node-list` option of the `clresourcetype` command.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Run the `clresourcetype` command to determine the name of the resource type that you need for this procedure.

```
# clresourcetype show -v
```

- 3 Change the resource type property.

For resource types, you can change only certain properties. To determine whether you can change a property, see the Tunable entry for the property in the [rt_properties\(5\)](#) man page.

```
# clresourcetype set -n installed-node-list \
[-p property=new-value] resource-type
```

`-n installed-node-list` Specifies the names of nodes on which this resource type is installed.

`-p property=new-value` Specifies the name of the standard property to change and the new value of the property.

You cannot change the `Installed_nodes` property explicitly. To change this property, specify the `-n installed-node-list` option of the `clresourcetype` command.

- 4 Verify that the resource type property has been changed.

```
# clresourcetype show resource-type
```

Example 2–19 Changing a Resource Type Property

This example shows how to change the `SUNW.apache` property to define that this resource type is installed on the global-cluster voting nodes of (`phys-schost-1` and `phys-schost-2`).

```
# clresourcetype set -n phys-schost-1,phys-schost-2 SUNW.apache
# clresourcetype show SUNW.apache
```

Resource Type:	SUNW.apache:4
RT_description:	Apache Web Server on Sun Cluster
RT_version:	4
API_version:	2
RT_basedir:	/opt/SUNWscapc/bin
Single_instance:	False
Proxy:	False
Init_nodes:	All potential masters
Installed_nodes:	All
Failover:	False
Pkglist:	SUNWscapc
RT_system:	False

▼ How to Change Resource Group Properties

This procedure explains how to change resource group properties. For a description of resource group properties, see the [rg_properties\(5\)](#) man page.

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the resource group to change
- The name of the resource group property to change and its new value

1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

2 Change the resource group property.

```
# clresourcegroup set -p property=new-value resource-group
-p property           Specifies the name of the property to change
resource-group       Specifies the name of the resource group
```

3 Verify that the resource group property has been changed.

```
# clresourcegroup show resource-group
```

Example 2–20 Changing a Resource Group Property

This example shows how to change the Failback property for the resource group (resource-group-1).

```
# clresourcegroup set-p Failback=True resource-group-1
# clresourcegroup show resource-group-1
```

▼ How to Change Resource Properties

This procedure explains how to change extension properties and standard properties of a resource.

- For a description of standard resource properties, see the [r_properties\(5\)](#) man page.
- For a description of the extension properties of a resource, see the documentation for the resource's resource type.

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the resource with the property to change
- The name of the property to change

1 On a cluster member, become superuser or assume a role that provides solaris.cluster.modify RBAC authorization.

2 View the current resource property settings.

```
# clresource show -v resource
```

3 Change the resource property.

```
# clresource set -p standard-property=new-value | -p "extension-property \
[{node-specifier}]"=new-value resource
```

```
-p standard-property=new-value
```

Specifies the name of the standard property to change.

```
-p "extension-property[{node-specifier}]"=new-value
```

Specifies the name of the extension property to change.

node-specifier is an *optional* qualifier to the -p and -x options. This qualifier indicates that the extension property or properties on *only* the specified node or nodes are to be set when the resource is created. The specified extension properties on other nodes in the cluster are not set. If you do not include *node-specifier*, the specified extension properties on all nodes in

the cluster are set. You can specify a node name or a node identifier for *node-specifier*. Examples of the syntax of *node-specifier* include the following:

```
-p "myprop{phys-schost-1}"
```

The braces ({}) indicate that you are setting the specified extension property on only node `phys-schost-1`. For most shells, the double quotation marks (") are required.

You can also use the following syntax to set an extension property in two different global-cluster non-voting nodes on two different global-cluster voting nodes:

```
-x "myprop{phys-schost-1:zoneA,phys-schost-2:zoneB}"
```

Note – The extension property that you specify with *node-specifier* must be declared in the RTR file as a per-node property.

resource

Specifies the name of the resource.

4 Verify that the resource property has been changed.

```
# clresource show -v resource
```

Example 2–21 Changing a Standard Resource Property

This example shows how to change the system-defined `Start_timeout` property for the resource (`resource-1`).

```
# clresource set -p start_timeout=30 resource-1
# clresource show -v resource-1
```

Example 2–22 Changing an Extension Resource Property

This example shows how to change an extension property (`Log_level`) for the resource (`resource-1`).

```
# clresource set -p Log_level=3 resource-1
# clresource show -v resource-1
```

▼ How to Modify a Logical Hostname Resource or a Shared Address Resource

By default, logical hostname resources and shared address resources use name services for name resolution. You might configure a cluster to use a name service that is running on the same cluster. During the failover of a logical hostname resource or a shared address resource, a

name service that is running on the cluster might also be failing over. If the logical hostname resource or the shared address resource uses the name service that is failing over, the resource fails to fail over.

Note – Configuring a cluster to use a name server that is running on the same cluster might impair the availability of other services on the cluster.

To prevent such a failure to fail over, modify the logical hostname resource or the shared address resource to bypass name services. To modify the resource to bypass name services, set the `CheckNameService` extension property of the resource to `false`. You can modify the `CheckNameService` property at any time.

Note – If your version of the resource type is earlier than 2, you must upgrade the resource type before you attempt to modify the resource. For more information, see [“Upgrading a Preregistered Resource Type” on page 100](#).

- 1 **On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**

- 2 **Change the resource property.**

```
# clresource set -p CheckNameService=false resource
```

`-p CheckNameService=false` Sets the `CheckNameService` extension property of the resource to `false`.

resource Specifies the name of the logical hostname resource or shared address resource that you are modifying.

Clearing the STOP_FAILED Error Flag on Resources

When the `Failover_mode` resource property is set to `NONE` or `SOFT`, a failure of the resource's `STOP` method causes the following effects:

- The individual resource goes into the `STOP_FAILED` state.
- The resource group that contains the resource goes into the `ERROR_STOP_FAILED` state.

In this situation, you cannot perform the following operations:

- Bringing online the resource group on any node
- Adding resources to the resource group
- Removing resources from the resource group
- Changing the properties of the resource group
- Changing the properties of resources in the resource group

▼ How to Clear the STOP_FAILED Error Flag on Resources

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the following information.

- The name of the node where the resource is STOP_FAILED
- The name of the resource and resource group that are in STOP_FAILED state

1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

2 Identify which resources have gone into the STOP_FAILED state and on which nodes.

```
# clresource status
```

3 Manually stop the resources and their monitors on the nodes on which they are in STOP_FAILED state.

This step might require that you kill processes or run commands that are specific to resource types or other commands.

4 Clear the STOP_FAILED error flag on the resources.

```
# clresource clear -f STOP_FAILED -n nodelist resource
```

`-f STOP_FAILED` Specifies the flag name.

`-n nodelist` Specifies a comma-separated list of the names of the nodes where the resource is in the STOP_FAILED state. The list may contain one node name or more than one node name.

`resource` Specifies the name of the resource.

5 Check the resource group state on the nodes where you cleared the STOP_FAILED flag in [Step 4](#).

```
# clresourcegroup status
```

The resource group state should now be OFFLINE or ONLINE.

The resource group remains in the ERROR_STOP_FAILED state in the following combination of circumstances:

- The resource group was being switched offline when the STOP method failure occurred.
- The resource that failed to stop had a dependency on other resources in the resource group.

6 If the resource group remains in the ERROR_STOP_FAILED state, correct the error as follows.

a. Switch the resource group offline on the appropriate nodes.

```
# clresourcegroup offline resource-group
```

resource-group Specifies the name of the resource group to switch offline.

b. Switch the resource group to the ONLINE state.

See Also The following man pages:

- [clresource\(1CL\)](#)
- [clresourcegroup\(1CL\)](#)

Clearing the Start_failed Resource State

The Start_failed resource state indicates that a Start or Prenet_start method failed or timed out on a resource, but its resource group came online anyway. The resource group comes online even though the resource has been placed in a faulted state and might not be providing service. This state can occur if the resource's Failover_mode property is set to None or to another value that prevents the failover of the resource group.

Unlike the Stop_failed resource state, the Start_failed resource state does *not* prevent you or the Oracle Solaris Cluster software from performing actions on the resource group. You need only to execute a command that restarts the resource.

Use any one of the following procedures to clear this condition.

▼ How to Clear a Start_failed Resource State by Switching Over a Resource Group

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that the following conditions are met:

- You have the following information:
 - The name of the resource group that you are switching over
 - The name of the node on which to switch over the resource group
- The nodes where the resource group is to be brought online or to remain online are in the cluster.

1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

2 Switch the resource group to the new node.

```
# clresourcegroup switch [-n node-zone-list] resource-group
-n node-zone-list    Specifies a comma-separated, ordered list of nodes that can master this
                    resource group. This resource group is switched offline on all of the other
                    nodes. The format of each entry in the list is node:zone. In this format,
                    node specifies the node name and zone specifies the name of a
                    global-cluster non-voting node. To specify the global cluster-voting node,
                    or to specify a node without global-cluster non-voting nodes, specify only
                    node.

                    This list is optional. If you omit this list, the resource group is switched on
                    all nodes in the resource group's node list.

resource-group       Specifies the name of the resource group to switch.
```

Note – If any resource group that you are switching declares a strong affinity for other resource groups, the attempt to switch might fail or be delegated. For more information, see [“Distributing Online Resource Groups Among Cluster Nodes” on page 151](#).

3 Verify that the resource group has been switched to the new node and that the `Start_failed` resource state is cleared.

```
# clresourcegroup status
```

The output from this command indicates the state of the resource and the resource group that has been switched over.

Example 2–23 Clearing a `Start_failed` Resource State by Switching Over a Resource Group

This example shows how to clear a `Start_failed` resource state that has occurred on the `rscon` resource in the `resource-group-1` resource group. The command clears this condition by switching the resource group to the global cluster voting node `phys-schost-2`.

1. To verify that the resource is in the `Start_failed` resource state on `phys-schost-1`, the following command is run:

```
# clresource status

=== Cluster Resources ===

Resource Name      Node Name      Status      Message
-----
rscon              phys-schost-1  Faulted     Faulted
                  phys-schost-2  Offline     Offline
```


hastor	phys-schost-1	Online	Online
	phys-schost-2	Offline	Offline

- To perform the switch, the following command is run:

```
# clresourcegroup switch -n phys-schost-2 resource-group-1
```

- To verify that the resource group is switched to be online on phys-schost-2 and that the Start_failed resource status is cleared, the following command is run:

```
# clresource status
```

```
=== Cluster Resources ===
```

Resource Name	Node Name	Status	Message
-----	-----	-----	-----
rscon	phys-schost-1	Offline	Offline
	phys-schost-2	Online	Online
hastor	phys-schost-1	Online	Online
	phys-schost-2	Offline	Offline

See Also The `clresourcegroup(1CL)` man page.

▼ How to Clear a Start_failed Resource State by Restarting a Resource Group

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that the following conditions are met:

- You have the following information:
 - The name of the resource group that you are restarting
 - The name of the node on which to restart the resource group
- The nodes where the resource group is to be brought online or to remain online are cluster nodes.

- On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- Restart the resource group.

```
# clresourcegroup restart -n node resource-group
```

- n node Specifies the name of the node on which the resource group is to be restarted. This resource group is switched offline on all of the other nodes.
- resource-group Specifies the name of the resource group to restart.

3 Verify that the resource group has been restarted on the new node and that the Start_failed resource state is cleared.

clresourcegroup status

The output from this command indicates the state of the resource and the resource group that has been restarted.

Example 2–24 Clearing a Start_failed Resource State by Restarting a Resource Group

This example shows how to clear a Start_failed resource state that has occurred on the rscon resource in the resource-group-1 resource group. The command clears this condition by restarting the resource group on the global-cluster voting node phys-schost-1.

1. To verify that the resource is in the Start_failed resource state on phys-schost-1, the following command is run:

clresource status

=== Cluster Resources ===

Resource Name	Node Name	Status	Message
-----	-----	-----	-----
rscon	phys-schost-1	Faulted	Faulted
	phys-schost-2	Offline	Offline
hastor	phys-schost-1	Online	Online
	phys-schost-2	Offline	Offline

2. To restart the resource, the following command is run:

clresourcegroup restart -n phys-schost-1 -g resource-group-1

3. To verify that the resource group is restarted on phys-schost-1 and that the Start_failed resource status is cleared, the following command is run:

clresource status

=== Cluster Resources ===

Resource Name	Node Name	Status	Message
-----	-----	-----	-----
rscon	phys-schost-1	Offline	Offline
rscon	phys-schost-2	Online	Online
hastor	phys-schost-1	Online	Online
hastor	phys-schost-2	Offline	Offline

See Also The [clresourcegroup\(1CL\)](#) man page.

▼ How to Clear a Start_failed Resource State by Disabling and Enabling a Resource

Note – Perform this procedure from any cluster node.

Before You Begin Ensure that you have the name of the resource that you are disabling and enabling.

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Disable and then enable the resource.

```
# clresource disable resource
# clresource enable resource
```

resource Specifies the name of the resource.

- 3 Verify that the resource has been disabled and enabled and that the Start_failed resource state is cleared.

```
# clresource status
```

The output from this command indicates the state of the resource that has been disabled and re-enabled.

Example 2–25 Clearing a Start_failed Resource State by Disabling and Enabling a Resource

This example shows how to clear a Start_failed resource state that has occurred on the `rscon` resource by disabling and enabling the resource.

1. To verify that the resource is in the Start_failed resource state, the following command is run:

```
# clresource status
```

```
=== Cluster Resources ===
```

Resource Name	Node Name	Status	Message
-----	-----	-----	-----
rscon	phys-schost-1	Faulted	Faulted
	phys-schost-2	Offline	Offline
hastor	phys-schost-1	Online	Online
	phys-schost-2	Offline	Offline

2. To disable and re-enable the resource, the following commands are run:

```
# clresource disable rscon
# clresource enable rscon
```

3. To verify that the resource is re-enabled and that the `Start_` failed resource status is cleared, the following command is run:

```
# clresource status

=== Cluster Resources ===

Resource Name      Node Name      Status      Message
-----
rscon              phys-schost-1  Online      Online
                   phys-schost-2  Offline     Offline

hastor             phys-schost-1  Online      Online
                   phys-schost-2  Offline     Offline
```

See Also The `clresource(1CL)` man page.

Upgrading a Preregistered Resource Type

In Sun Cluster 3.1 9/04, the following preregistered resource types were enhanced:

- `SUNW.LogicalHostname`, which represents a logical hostname
- `SUNW.SharedAddress`, which represents a shared address

The purpose of these enhancements is to enable you to modify logical hostname resources and shared address resources to bypass name services for name resolution.

Upgrade these resource types if all conditions in the following list apply:

- You are upgrading from an earlier version of Oracle Solaris Cluster.
- You need to use the new features of the resource types.

For general instructions that explain how to upgrade a resource type, see “[Upgrading a Resource Type](#)” on page 42. The information that you need to complete the upgrade of the preregistered resource types is provided in the subsections that follow.

Information for Registering the New Resource Type Version

To determine the version of the resource type that is registered, use one command from the following list:

- `clresourcetype list`
- `clresourcetype list -v`

EXAMPLE 2-26 Registering a New Version of the SUNW.LogicalHostname Resource Type

This example shows the command for registering version 4 of the SUNW.LogicalHostname resource type during an upgrade.

```
# clresourcetype register SUNW.LogicalHostname:4
```

Information for Migrating Existing Instances of the Resource Type

The information that you need to migrate an instance of a preregistered resource type is as follows:

- You can perform the migration at any time.
- If you need to use the new features of the preregistered resource type, the required value of the `Type_version` property is 2.
- If you are modifying the resource to bypass name services, set the `CheckNameService` extension property of the resource to `false`.

EXAMPLE 2-27 Migrating a Logical Hostname Resource

This example shows the command for migrating the logical hostname resource `lhostrs`. As a result of the migration, the resource is modified to bypass name services for name resolution.

```
# clresource set -p CheckNameService=false -p Type_version=2 lhostrs
```

Reregistering Preregistered Resource Types After Inadvertent Deletion

The resource types `SUNW.LogicalHostname` and `SUNW.SharedAddress` are preregistered. All of the logical hostname and shared address resources use these resource types. You never need to register these two resource types, but you might inadvertently delete them. If you have deleted resource types inadvertently, use the following procedure to reregister them.

Note – If you are upgrading a preregistered resource type, follow the instructions in [“Upgrading a Preregistered Resource Type” on page 100](#) to register the new resource type version.

Note – Perform this procedure from any cluster node.

▼ How to Reregister Preregistered Resource Types After Inadvertent Deletion

- Reregister the resource type.

```
# clresourcetype register SUNW.resource-type
```

resource-type Specifies the resource type to add (reregister). The resource type can be either `SUNW.LogicalHostname` or `SUNW.SharedAddress`.

Example 2–28 Reregistering a Preregistered Resource Type After Inadvertent Deletion

This example shows how to reregister the `SUNW.LogicalHostname` resource type.

```
# clresourcetype register SUNW.LogicalHostname
```

See Also The `clresourcetype(1CL)` man page.

Adding or Removing a Node to or From a Resource Group

The procedures in this section enable you to perform the following tasks.

- Configuring a cluster node to be an additional master of a resource group
- Removing a node from a resource group

The procedures are slightly different, depending on whether you plan to add or remove the node to or from a failover or scalable resource group.

Failover resource groups contain network resources that both failover and scalable services use. Each IP subnetwork connected to the cluster has its own network resource that is specified and included in a failover resource group. The network resource is either a logical hostname or a shared address resource. Each network resource includes a list of IPMP groups that it uses. For failover resource groups, you must update the complete list of IPMP groups for each network resource that the resource group includes (the `netiflist` resource property).

The procedure for scalable resource groups involves the following steps:

1. Repeating the procedure for failover groups that contain the network resources that the scalable resource uses
2. Changing the scalable group to be mastered on the new set of hosts

For more information, see the `clresourcegroup(1CL)` man page.

Note – Run either procedure from any cluster node.

Adding a Node to a Resource Group

The procedure to follow to add a node to a resource group depends on whether the resource group is a scalable resource group or a failover resource group. For detailed instructions, see the following sections:

- “How to Add a Node to a Scalable Resource Group” on page 103
- “How to Add a Node to a Failover Resource Group” on page 104

You must supply the following information to complete the procedure.

- The names and node IDs of all of the cluster nodes and names of zones
- The names of the resource groups to which you are adding the node
- The name of the IPMP group that is to host the network resources that are used by the resource group on all of the nodes

Also, be sure to verify that the new node is already a cluster member.

▼ How to Add a Node to a Scalable Resource Group

- 1 **For each network resource that a scalable resource in the resource group uses, make the resource group where the network resource is located run on the new node.**
See [Step 1](#) through [Step 5](#) in the following procedure for details.
- 2 **Add the new node to the list of nodes that can master the scalable resource group (the `nodeList` resource group property).**

This step overwrites the previous value of `nodeList`, and therefore you must include all of the nodes that can master the resource group here.

```
# clresourcegroup set [-n node-zone-list] resource-group
```

`-n node-zone-list` Specifies a comma-separated, ordered list of nodes that can master this resource group. This resource group is switched offline on all of the other nodes. The format of each entry in the list is *node:zone*. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the `NodeList` property is set to all nodes in the cluster.

resource-group Specifies the name of the resource group to which the node is being added.

- 3 (Optional) Update the scalable resource's `Load_balancing_weights` property to assign a weight to the node that you are adding to the resource group.

Otherwise, the weight defaults to 1. See the `clresourcegroup(1CL)` man page for more information.

▼ How to Add a Node to a Failover Resource Group

- 1 Display the current node list and the current list of IPMP groups that are configured for each resource in the resource group.

```
# clresourcegroup show -v resource-group | grep -i nodelist
# clresourcegroup show -v resource-group | grep -i netiflist
```

Note – The output of the command line for `nodelist` and `netiflist` identifies the nodes by node name. To identify node IDs, run the command `clnode show -v | grep -i node-id`.

- 2 Update `netiflist` for the network resources that the node addition affects.

This step overwrites the previous value of `netiflist`, and therefore you must include all the IPMP groups here.

```
# clresource set -p netiflist=netiflist network-resource
```

`-p netiflist=netiflist` Specifies a comma-separated list that identifies the IPMP groups that are on each node. Each element in *netiflist* must be in the form of `netif@node`. `netif` can be given as an IPMP group name, such as `sc_ipmp0`. The node can be identified by the node name or node ID, such as `sc_ipmp0@1` or `sc_ipmp@phys-schost-1`.

network-resource Specifies the name of the network resource (logical hostname or shared address) that is being hosted on the *netiflist* entries.

- 3 If the `HASStoragePlus AffinityOn` extension property equals `True`, add the node to the appropriate disk set or device group.

- If you are using Solaris Volume Manager, use the `metaset` command.

```
# metaset -s disk-set-name -a -h node-name
```

`-s disk-set-name` Specifies the name of the disk set on which the `metaset` command is to work

`-a` Adds a drive or host to the specified disk set

`-h node-name` Specifies the node to be added to the disk set

- **SPARC: If you are using Veritas Volume Manager, use the `clsetup` utility.**
 - a. **On any active cluster member, start the `clsetup` utility.**

```
# clsetup
```

The Main Menu is displayed.
 - b. **On the Main Menu, type the number that corresponds to the option for device groups and volumes.**
 - c. **On the Device Groups menu, type the number that corresponds to the option for adding a node to a VxVM device group.**
 - d. **Respond to the prompts to add the node to the VxVM device group.**
- 4 **Update the node list to include all of the nodes that can now master this resource group.**

This step overwrites the previous value of `nodelist`, and therefore you must include all of the nodes that can master the resource group here.

```
# clresourcegroup set [-n node-zone-list] resource-group
```

`-n node-zone-list` Specifies a comma-separated, ordered list of global-cluster non-voting nodes that can master this resource group. This resource group is switched offline on all the other nodes. The format of each entry in the list is `node:zone`. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the `Nodelist` property is set to all nodes in the cluster.

`resource-group` Specifies the name of the resource group to which the node is being added.
- 5 **Verify the updated information.**

```
# clresourcegroup show -v resource-group | grep -i nodelist
# clresourcegroup show -v resource-group | grep -i netiflist
```

Example 2–29 Adding a Node to a Resource Group

This example shows how to add a global-cluster voting node (`phys-schost-2`) to a resource group (`resource-group-1`) that contains a logical hostname resource (`schost-2`).

```
# clresourcegroup show -v resource-group-1 | grep -i nodelist
( Nodelist:   phys-schost-1 phys-schost-3
# clresourcegroup show -v resource-group-1 | grep -i netiflist
( Res property name: NetIfList
```

```
Res property class: extension
List of IPMP
interfaces on each node
Res property type: stringarray
Res property value: sc_ipmp0@1 sc_ipmp0@3
```

(Only nodes 1 and 3 have been assigned IPMP groups. You must add an IPMP group for node 2.)

```
# clresource set -p netiflist=sc_ipmp0@1,sc_ipmp0@2,sc_ipmp0@3 schost-2
# metaset -s red -a -h phys-schost-2
# clresourcegroup set -n phys-schost-1,phys-schost-2,phys-schost-3 resource-group-1
# clresourcegroup show -v resource-group-1 | grep -i nodelist
Nodelist:      phys-schost-1 phys-schost-2
               phys-schost-3
# clresourcegroup show -v resource-group-1 | grep -i netiflist
Res property value: sc_ipmp0@1 sc_ipmp0@2
                   sc_ipmp0@3
```

Removing a Node From a Resource Group

The procedure to follow to remove a node from a resource group depends on whether the resource group is a scalable resource group or a failover resource group. For detailed instructions, see the following sections:

- [“How to Remove a Node From a Scalable Resource Group” on page 107](#)
- [“How to Remove a Node From a Failover Resource Group” on page 108](#)
- [“How to Remove a Node From a Failover Resource Group That Contains Shared Address Resources” on page 109](#)

To complete the procedure, you must supply the following information.

- Node names and node IDs of all of the cluster nodes

```
# clnode show -v | grep -i "Node ID"
```
- The name of the resource group or the names of the resource groups from which you plan to remove the node

```
# clresourcegroup show | grep "Nodelist"
```
- Names of the IPMP groups that are to host the network resources that are used by the resource groups on all of the nodes

```
# clresourcegroup show -v | grep "NetIfList.*value"
```

Additionally, be sure to verify that the resource group **is not mastered** on the node that you are removing. If the resource group **is mastered** on the node that you are removing, run the `clresourcegroup` command to switch the resource group offline from that node. The following `clresourcegroup` command brings the resource group offline from a given node, provided that *new-masters* does not contain that node.

```
# clresourcegroup switch -n new-masters resource-group
```

- n new-masters* Specifies the nodes that is now to master the resource group.
- resource-group* Specifies the name of the resource group that you are switching . This resource group is mastered on the node that you are removing.

For more information, see the [clresourcegroup\(1CL\)](#) man page.



Caution – If you plan to remove a node from all the resource groups, and you use a scalable services configuration, first remove the node from the scalable resource groups. Then remove the node from the failover groups.

▼ How to Remove a Node From a Scalable Resource Group

A scalable service is configured as two resource groups, as follows.

- One resource group is a scalable group that contains the scalable service resource.
- One resource group is a failover group that contains the shared address resources that the scalable service resource uses.

Additionally, the `RG_dependencies` property of the scalable resource group is set to configure the scalable group with a dependency on the failover resource group. For information about this property, see the [rg_properties\(5\)](#) man page.

For details about scalable service configuration, see [Oracle Solaris Cluster Concepts Guide](#).

Removing a node from the scalable resource group causes the scalable service to no longer be brought online on that node. To remove a node from the scalable resource group, perform the following steps.

1 Remove the node from the list of nodes that can master the scalable resource group (the `node-list` resource group property).

```
# clresourcegroup set [-n node-zone-list] scalable-resource-group
```

-n node-zone-list Specifies a comma-separated, ordered list of nodes that can master this resource group. This resource group is switched offline on all the other nodes. The format of each entry in the list is `node:zone`. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the `NodeList` property is set to all nodes in the cluster.

scalable-resource-group Specifies the name of the resource group from which the node is being removed.

- 2 (Optional) Remove the node from the failover resource group that contains the shared address resource.

For details, see [“How to Remove a Node From a Failover Resource Group That Contains Shared Address Resources”](#) on page 109.

- 3 (Optional) Update the `Load_balancing_weights` property of the scalable resource to remove the weight of the node that you are removing from the resource group.

See Also The `clresourcegroup(1CL)` man page.

▼ How to Remove a Node From a Failover Resource Group

Perform the following steps to remove a node from a failover resource group.



Caution – If you plan to remove a node from all of the resource groups, and you use a scalable services configuration, first remove the node from the scalable resource groups. Then use this procedure to remove the node from the failover groups.

Note – If the failover resource group contains shared address resources that scalable services use, see [“How to Remove a Node From a Failover Resource Group That Contains Shared Address Resources”](#) on page 109.

- 1 **Update the node list to include all of the nodes that can now master this resource group.**

This step removes the node and overwrites the previous value of the node list. Be sure to include all of the nodes that can master the resource group here.

```
# clresourcegroup set [-n node-zone-list] failover-resource-group
```

-n node-zone-list

Specifies a comma-separated, ordered list of nodes that can master this resource group. This resource group is switched offline on all the other nodes. The format of each entry in the list is `node:zone`. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the `NodeList` property is set to all nodes in the cluster.

failover-resource-group

Specifies the name of the resource group from which the node is being removed.

- 2 Display the current list of IPMP groups that are configured for each resource in the resource group.

```
# clresourcegroup show -v failover-resource-group | grep -i netiflist
```

- 3 Update `netiflist` for network resources that the removal of the node affects.

This step overwrites the previous value of `netiflist`. Be sure to include all of the IPMP groups here.

```
# clresource set -p netiflist=netiflist network-resource
```

Note – The output of the preceding command line identifies the nodes by node name. Run the command line `clnode show -v | grep -i "Node ID"` to find the node ID.

<code>-p netiflist=netiflist</code>	Specifies a comma-separated list that identifies the IPMP groups that are on each node. Each element in <i>netiflist</i> must be in the form of <code>netif@node</code> . <code>netif</code> can be given as an IPMP group name, such as <code>sc_ipmp0</code> . The node can be identified by the node name or node ID, such as <code>sc_ipmp0@1</code> or <code>sc_ipmp@phys-schost-1</code> .
<code>network-resource</code>	Specifies the name of the network resource that is hosted on the <code>netiflist</code> entries.

Note – Oracle Solaris Cluster does not support the use of the adapter name for `netif`.

- 4 Verify the updated information.

```
# clresourcegroup show -vfailover-resource-group | grep -i nodelist
# clresourcegroup show -vfailover-resource-group | grep -i netiflist
```

▼ How to Remove a Node From a Failover Resource Group That Contains Shared Address Resources

In a failover resource group that contains shared address resources that scalable services use, a node can appear in the following locations.

- The node list of the failover resource group
- The `auxnodelist` of the shared address resource

To remove the node from the node list of the failover resource group, follow the procedure [“How to Remove a Node From a Failover Resource Group” on page 108](#).

To modify the `auxnodelist` of the shared address resource, you must remove and recreate the shared address resource.

If you remove the node from the failover group's node list, you can continue to use the shared address resource on that node to provide scalable services. To continue to use the shared address resource, you must add the node to the `auxnodeList` of the shared address resource. To add the node to the `auxnodeList`, perform the following steps.

Note – You can also use the following procedure to **remove** the node from the `auxnodeList` of the shared address resource. To remove the node from the `auxnodeList`, you must delete and recreate the shared address resource.

- 1 **Switch the scalable service resource offline.**
- 2 **Remove the shared address resource from the failover resource group.**
- 3 **Create the shared address resource.**

Add the node ID or node name of the node that you removed from the failover resource group to the `auxnodeList`.

```
# clressharedaddress create -g failover-resource-group \
-X new-auxnodeList shared-address
```

failover-resource-group The name of the failover resource group that used to contain the shared address resource.

new-auxnodeList The new, modified `auxnodeList` with the desired node added or removed.

shared-address The name of the shared address.

Example – Removing a Node From a Resource Group

This example shows how to remove a node (`phys-schost-3`) from a resource group (`resource-group-1`) that contains a logical hostname resource (`schost-1`).

```
# clresourcegroup show -v resource-group-1 | grep -i nodelist
Nodelist:          phys-schost-1 phys-schost-2
                  phys-schost-3
# clresourcegroup set -n phys-schost-1,phys-schost-2 resource-group-1
# clresourcegroup show -v resource-group-1 | grep -i netiflist
( Res property name: NetIfList
Res property class: extension
( List of IPMP
interfaces on each node
( Res property type: stringarray
  Res property value: sc_ipmp0@1 sc_ipmp0@2
                    sc_ipmp0@3

(sc_ipmp0@3 is the IPMP group to be removed.)

# clresource set -p netiflist=sc_ipmp0@1,sc_ipmp0@2 schost-1
```

```
# clresourcegroup show -v resource-group-1 | grep -i nodelist
Nodelist:      phys-schost-1 phys-schost-2
# clresourcegroup show -v resource-group-1 | grep -i netiflist
Res property value: sc_ipmp0@1 sc_ipmp0@2
```

Migrating the Application From a Global-Cluster Voting Node to a Global-Cluster Non-Voting Node

You can migrate the application resources from a global-cluster voting node to a global-cluster non-voting node.

Note – The data services you want to migrate should be scalable and also be supported in global-cluster non-voting nodes.

▼ How to Migrate the Application From a Global-Cluster Voting Node to a Global-Cluster Non-Voting Node

The procedure assumes a three node cluster with a global-cluster non-voting node created on each of the three nodes. The configuration directory that is made highly available using the HAStoragePlus resource should also be accessible from the global-cluster non-voting nodes.

- 1 **Create the failover resource group with the global-cluster voting node that holds the shared address that the scalable resource group is to use.**

```
# clresourcegroup create -n node1,node2,node3 sa-resource-group
```

sa-resource-group Specifies your choice of the name of the failover resource group to add. This name must begin with an ASCII character.

- 2 **Add the shared address resource to the failover resource group.**

```
# clressharedaddress create -g sa-resource-group -h hostnamelist, ... \
[-X auxnodelist] -N netiflist network-resource
```

-g sa-resource-group Specifies the resource group name. In the node list of a shared address resource, do not specify more than one global-cluster non-voting node on the same global-cluster voting node. Specify the same list of *nodename:zonename* pairs as the node list of the scalable resource group.

-h hostnamelist, ... Specifies a comma-separated list of shared address hostnames.

-X auxnodelist Specifies a comma-separated list of node names or IDs or zones that identify the cluster nodes that can host the shared address but never

serve as primary if failover occurs. These nodes are mutually exclusive, with the nodes identified as potential masters in the resource group's node list. If no auxiliary node list is explicitly specified, the list defaults to the list of all cluster node names that are not included in the node list of the resource group that contains the shared address resource.

Note – To ensure that a scalable service runs in all global-cluster non-voting nodes that were created to master the service, the complete list of nodes must be included in the node list of the shared address resource group or the `auxnodeList` of the shared address resource. If all the global-cluster non-voting nodes are listed in the node list, the `auxnodeList` can be omitted.

`-N netiflist`

Specifies an optional, comma-separated list that identifies the IPMP groups that are on each node. Each element in *netiflist* must be in the form of `netif@node`. `netif` can be given as an IPMP group name, such as `sc_ipmp0`. The node can be identified by the node name or node ID, such as `sc_ipmp0@1` or `sc_ipmp@phys-schost-1`.

Note – Oracle Solaris Cluster does not support the use of the adapter name for `netif`.

network-resource

Specifies an optional resource name of your choice.

3 Create the scalable resource group.

```
# clresourcegroup create -S [-p Maximum primaries=m] [-p Desired primaries=n] \
-n node1,node2,node3 resource-group-1
```

`-S`

Specifies that the resource group is to be multi-mastered. If the `-p Maximum primaries` and `-p Desired primaries` options are omitted, both properties are set to the number of nodes in the resource group's node list.

`-p Maximum primaries=m`

Specifies the maximum number of active primaries for this resource group.

`-p Desired primaries=n`

Specifies the number of active primaries on which the resource group should attempt to start.

resource-group-1

Specifies your choice of the name of the scalable resource group to add. This name must begin with an ASCII character.

- 4 Create the **HAStoragePlus** resource **hastorageplus-1**, and define the filesystem mount points.

```
# clresource create -g resource-group-1 -t SUNW.HAStoragePlus \
-p FilesystemMountPoints=/global/resource-group-1 hastorageplus-1
```

The resource is created in the enabled state.

- 5 Register the resource type for the application.

```
# clresourcetype register resource-type
```

resource-type Specifies name of the resource type to add. See the release notes for your release of Oracle Solaris Cluster to determine the predefined name to supply.

- 6 Add the application resource to **resource-group-1**, and set the dependency to **hastorageplus-1**.

```
# clresource create -g resource-group-1 -t SUNW.application \
[-p "extension-property[{node-specifier}]"=value, ...] -p Scalable=True \
-p Resource_dependencies=network-resource -p Port_list=port-number/protocol \
-p Resource_dependencies_offline_restart=hastorageplus-1 resource
```

- 7 Bring the failover resource group online.

```
# clresourcegroup online sa-resource-group
```

- 8 Bring the scalable resource group online on all the nodes.

```
# clresourcegroup online resource-group-1
```

- 9 Install and boot *zone1* on each of the nodes, *node1*, *node2*, *node3*.

- 10 Bring the application resource group offline on two nodes (*node1*, *node2*).

Note – Ensure the shared address is online on *node3*.

```
# clresourcegroup switch -n node3 resource-group-1
```

resource-group-1 Specifies the name of the resource group to switch.

- 11 Update the **nodeList** property of the failover resource group to include the global-cluster non-voting node of the corresponding nodes removed from the node list.

```
# clresourcegroup set -n node1:zone1,node2:zone1,node3 sa-resource-group
```

- 12 Update the **nodeList** property of the application resource group to include the global-cluster non-voting node of the corresponding nodes removed from node list.

```
# clresourcegroup set node1:zone1,node2:zone1,node3 resource-group-1
```

- 13 Bring the failover resource group and application resource group online only on the newly added zones.

Note – The failover resource group will be online only on *node1:zone1* and application resource group will be online only on *node1:zone1* and *node2:zone1*.

```
# clresourcegroup switch -n node1:zone1 sa-resource-group
```

```
# clresourcegroup switch -n node1:zone1,node2:zone1 resource-group-1
```

- 14 Update the `nodeList` property of both the resource groups to include the global-cluster non-voting node of *node3* by removing the global node, *node3* from the list.**

```
# clresourcegroup set node1:zone1,node2:zone1,node3:zone1 sa-resource-group
```

```
# clresourcegroup set node1:zone1,node2:zone1,node3:zone1 resource-group-1
```

- 15 Bring both the resource groups online on the global-cluster non-voting nodes.**

```
# clresourcegroup switch -n node1:zone1 sa-resource-group
```

```
# clresourcegroup switch -n node1:zone1,node2:zone1,node3:zone1 resource-group-1
```

Synchronizing the Startups Between Resource Groups and Device Groups

After a cluster boots or services fail over to another node, global devices and local and cluster file systems might require time to become available. However, a data service can run its `START` method before global devices and local and cluster file systems come online. If the data service depends on global devices or local and cluster file systems that are not yet online, the `START` method times out. In this situation, you must reset the state of the resource groups that the data service uses and restart the data service manually.

To avoid these additional administrative tasks, use the `HASStoragePlus` resource type. Add an instance of `HASStoragePlus` to all resource groups whose data service resources depend on global devices or local and cluster file systems. Instances of these resource types perform the following operations:

- Forcing the `START` method of the other resources in the same resource group to wait until global devices and local and cluster file systems become available.

If an application resource is configured on top of an `HASStoragePlus` resource, the application resource must define the offline restart dependency on the underlying `HASStoragePlus` resource. This ensures that the application resource comes online after the dependent `HASStoragePlus` resource comes online, and goes offline before the `HASStoragePlus` resource goes offline.

The following command creates an offline restart dependency from an application resource to a HASP resource:

```
# clrs set -p Resource_dependencies_offline_restart=hasp_rs applicaton_rs
```

To create an HASToragePlus resource, see [“How to Set Up the HASToragePlus Resource Type for New Resources” on page 118](#).

Managed Entity Monitoring by HASToragePlus

All entities that are managed by the HASToragePlus resource type are monitored. The SUNWHASToragePlus resource type provides a fault monitor to monitor the health of the entities managed by the HASP resource, including global devices, file systems, and ZFS storage pools. The fault monitor runs fault probes on a regular basis. If one of the entities becomes unavailable, the resource is restarted or a failover to another node is performed. If more than one entity is monitored, the fault monitor probes them all at the same time. Ensure that all configuration changes to the managed entities are completed before you enable monitoring.

Note – Version 9 of the HASToragePlus resource fault monitor probes the devices and file systems it manages by reading and writing to the file systems. If a read operation is blocked by any software on the I/O stack and the HASToragePlus resource is required to be online, the user must disable the fault monitor. For example, you must unmonitor the HASToragePlus resource managing the Sun StorageTek Availability Suite Remote Replication volumes because Sun StorageTek Availability Suite blocks reading from any bitmap volume or any data volume in the NEED SYNC state. The HASToragePlus resource managing the Sun StorageTek Availability Suite volumes must be online at all times.

For more information on the properties that enable monitoring for managed entities, see the [SUNW.HASToragePlus\(5\)](#) man page.

For instructions on enabling and disabling monitoring for managed entities, see [“How to Enable a Resource Fault Monitor” on page 79](#).

Depending on the type of managed entity, the fault monitor probes the target by reading or writing to it. If more than one entity is monitored, the fault monitor probes them all at the same time.

TABLE 2-2 What the Fault Monitor Verifies

Monitored Entity	What the Fault Monitor Verifies
Global device	<ul style="list-style-type: none"> ■ The device group is online or degraded. ■ The device is readable.
Raw device group	<ul style="list-style-type: none"> ■ The device group is online or degraded. ■ For each device of the device group, its path (<code>/dev/global/rdisk/device</code>) is available. ■ Partitions of every device are readable.
Solaris Volume Manager device group	<ul style="list-style-type: none"> ■ The device group is online or degraded. ■ The path of the metaset (<code>/dev/md/metaset</code>) is valid. ■ The Solaris Volume Manager reported status from the primary of the device group: <ul style="list-style-type: none"> ■ The unmirrored metadevice is not in any of the following error states: Needs Maintenance, Last Erred, or Unavailable. ■ At least one submirror of a mirror is not in an error state. An error with some, but not all submirrors, is treated as partial error. ■ The unmirrored metadevice is readable from the primary. ■ Some submirrors of a mirror are readable. An error with some, but not all, submirrors is treated as partial error.
VxVM device group	<ul style="list-style-type: none"> ■ The device group is online or degraded. ■ The path of the VxVM disk group (<code>/dev/vx/rdsk/diskgroup</code>) is valid. ■ On the primary of the device group, VxVM reports all volumes in an Enabled state. ■ Every VxVM volume is readable from the primary.
File systems (including UFS, VxFS, QFS and PxFS)	<ul style="list-style-type: none"> ■ The file system is mounted. ■ Every device under the file system is readable. ■ The file system is readable, if the <code>I00ption</code> property is set to <code>ReadOnly</code>. ■ The file system is writable, if the <code>I00ption</code> property is set to <code>ReadWrite</code>. ■ If the file system is mounted read-only but the <code>I00ption</code> property is set to <code>ReadWrite</code>, the fault monitor issues a warning and then tries to read it (rather than write to it). ■ To avoid having the <code>HASStoragePlus</code> resource go offline when a file system hits its quota, set the <code>I00ption</code> to <code>ReadOnly</code>. The <code>ReadOnly</code> option ensures that the fault monitor will not attempt to write to the file system.

TABLE 2-2 What the Fault Monitor Verifies (Continued)

Monitored Entity	What the Fault Monitor Verifies
ZFS storage pool	<ul style="list-style-type: none"> ■ The pool status is OK or Degraded. ■ Each non-legacy file system is mounted. ■ Each non-legacy file system is readable, if the <code>I00pt ion</code> property is set to <code>ReadOnly</code>. ■ Each non-legacy file system is writable, if the <code>I00pt ion</code> property is set to <code>ReadWrite</code>. ■ If a non-legacy file system is mounted read-only but the <code>I00pt ion</code> property is set to <code>ReadWrite</code>, the fault monitor issues a warning and then tries to read it (rather than write to it). ■ To avoid having the <code>HAStoragePlus</code> resource go offline when a file system hits its quota, set the <code>I00pt ion</code> to <code>ReadOnly</code>. The <code>ReadOnly</code> option ensures that the fault monitor will not attempt to write to the file system. <p>Note – When all connections to a top-level ZFS storage device are lost, queries about the ZFS storage pool or associated file system will hang. To prevent the fault monitor from hanging, you must set the <code>fail_mode</code> property of the ZFS storage pool to <code>panic</code>.</p>

For instructions on enabling a resource fault monitor, see [“How to Enable a Resource Fault Monitor” on page 79](#).

Troubleshooting Monitoring for Managed Entities

If monitoring is not enabled on the managed entities, perform the following troubleshooting steps:

1. Ensure that the `hastorageplus_probe` process is running.
2. Look for error messages on the console.
3. Enable debug messages to the `syslog` file.

```
# mkdir -p /var/cluster/rgm/rt/SUNW.HASStoragePlus:9
# echo 9 > /var/cluster/rgm/rt/SUNW.HASStoragePlus:9/loglevel
```

You should also check the `/etc/syslog.conf` file to ensure that messages with the `daemon.debug` facility level are logged to the `/var/adm/messages` file. Add the `daemon.debug` entry to the `/var/adm/messages` action if it is not already present.

Additional Administrative Tasks to Configure HAStoragePlus Resources for a Zone Cluster

When you configure HAStoragePlus resources for a zone cluster, you need to perform the following additional tasks before performing the steps for global cluster:

- While configuring file systems like UFS or standalone QFS in file system mount points, the file systems need to be configured to the zone cluster. For more information about configuring a file system to a zone cluster, see [“How to Add a Local File System to a Zone Cluster” in Oracle Solaris Cluster Software Installation Guide](#).
- While configuring global devices in global device paths, the devices need to be configured to the zone cluster. For more information about configuring global devices to a zone cluster, see [“Adding Storage Devices to a Zone Cluster” in Oracle Solaris Cluster Software Installation Guide](#).
- While configuring the ZFS file systems using Zpools, the ZFS pool needs to be configured to the zone cluster. For more information about configuring a ZFS file system to a zone cluster, see [“How to Add a ZFS Storage Pool to a Zone Cluster” in Oracle Solaris Cluster Software Installation Guide](#).

Note – The cluster file system is not supported for a zone cluster.

▼ How to Set Up the HAStoragePlus Resource Type for New Resources

In the following example, the resource group `resource-group-1` contains the following data services.

- Java System Web Server, which depends on `/global/resource-group-1`
- Oracle, which depends on `/dev/global/dsk/d5s2`
- NFS, which depends on `dsk/d6`

Note – To create an HAStoragePlus resource with Solaris ZFS (Zettabyte File System) as a highly available local file system see [“How to Set Up the HAStoragePlus Resource Type to Make a Local Solaris ZFS Highly Available” on page 131](#) section.

To create an HAStoragePlus resource `hastorageplus-1` for new resources in `resource-group-1`, read [“Synchronizing the Startups Between Resource Groups and Device Groups” on page 114](#) and then perform the following steps.

To create an HAStoragePlus resource, see [“Enabling Highly Available Local File Systems” on page 124](#).

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.

- 2 Create the resource group `resource-group-1`.

```
# clresourcegroup create resource-group-1
```

- 3 Determine whether the resource type is registered.

The following command prints a list of registered resource types.

```
# clresourcetype show | egrep Type
```

- 4 If you need to, register the resource type.

```
# clresourcetype register SUNW.HAStoragePlus
```

- 5 Create the `HAStoragePlus` resource `hastorageplus-1`, and define the filesystem mount points and global device paths.

```
# clresource create -g resource-group-1 -t SUNW.HAStoragePlus \  
-p GlobalDevicePaths=/dev/global/dsk/d5s2,dsk/d6 \  
-p FilesystemMountPoints=/global/resource-group-1 hastorageplus-1
```

`GlobalDevicePaths` can contain the following values.

- Global device group names, such as `nfs-dg`, `dsk/d5`
- Paths to global devices, such as `/dev/global/dsk/d1s2`, `/dev/md/nfsdg/dsk/d10`

`FilesystemMountPoints` can contain the following values.

- Mount points of local or cluster file systems, such as `/local-fs/nfs`, `/global/nfs`

Note – `HAStoragePlus` has a `Zpools` extension property that is used to configure ZFS file system storage pools and a `ZpoolsSearchDir` extension property that is used to specify the location to search for the devices of ZFS file system storage pools. The default value for the `ZpoolsSearchDir` extension property is `/dev/dsk`. The `ZpoolsSearchDir` extension property is similar to the `-d` option of the `zpool (1M)` command.

The resource is created in the enabled state.

- 6 Add the resources (Java System Web Server, Oracle, and NFS) to `resource-group-1`, and set their dependency to `hastorageplus-1`.

For example, for Java System Web Server, run the following command.

```
# clresource create -g resource-group-1 -t SUNW.iws \  
-p Confdir_list=/global/iws/schost-1 -p Scalable=False \  
-p Resource_dependencies=schost-1 -p Port_list=80/tcp \  
-p Resource_dependencies_offline_restart=hastorageplus-1 resource
```

The resource is created in the enabled state.

- 7 **Verify that you have correctly configured the resource dependencies.**

```
# clresource show -v resource | egrep Resource_dependencies_offline_restart
```

- 8 **Set resource-group-1 to the MANAGED state, and bring resource-group-1 online.**

```
# clresourcegroup online -M resource-group-1
```

More Information Affinity Switchovers

The HASStoragePlus resource type contains another extension property, AffinityOn, which is a Boolean that specifies whether HASStoragePlus must perform an affinity switchover for the global devices that are defined in GlobalDevicePaths and FileSystemMountPoints extension properties. For details, see the [SUNW.HASStoragePlus\(5\)](#) man page.

Note – The setting of the AffinityOn flag is ignored for scalable services. Affinity switchovers are not possible with scalable resource groups.

▼ How to Set Up the HASStoragePlus Resource Type for Existing Resources

Before You Begin Read “[Synchronizing the Startups Between Resource Groups and Device Groups](#)” on page 114.

- 1 **Determine whether the resource type is registered.**

The following command prints a list of registered resource types.

```
# clresourcetype show | egrep Type
```

- 2 **If you need to, register the resource type.**

```
# clresourcetype register SUNW.HASStoragePlus
```

- 3 **Create the HASStoragePlus resource hastorageplus-1.**

```
# clresource create -g resource-group \  
-t SUNW.HASStoragePlus -p GlobalDevicePaths= ... \  
-p FileSystemMountPoints=... -p AffinityOn=True hastorageplus-1
```

The resource is created in the enabled state.

- 4 **Set up the dependency for each of the existing resources, as required.**

```
# clresource set -p Resource_Dependencies_offline_restart=hastorageplus-1 resource
```

- 5 **Verify that you have correctly configured the resource dependencies.**

```
# clresource show -v resource | egrep Resource_dependencies_offline_restart
```


Configuring an HAStoragePlus Resource for Cluster File Systems

When an HAStoragePlus resource is configured for cluster file systems and brought online, it ensures that these file systems are available. The cluster file system is supported on the UNIX File System (UFS) and Veritas File System (VxFS). The instructions in this section apply to HAStoragePlus resources with UFS. Use HAStoragePlus with local file systems if the data service is I/O intensive. See [“How to Change the Cluster File System to Local File System in an HAStoragePlus Resource” on page 149](#) for information about how to change the file system of an HAStoragePlus resource.

The cluster file systems can be configured for zone clusters in the HAStoragePlus resources using the loopback mount mechanism. The SUNW.HAStoragePlus resource type makes the cluster file system available to a zone cluster by mounting the file system in the global cluster. The resource type then performs a loopback mount on the zone cluster nodes where the resource group is online.

Note – If you have a failover resource group, the resource group will be online on only one node. If you use a scalable resource group, the `Desired_primaries` property defines the number of nodes the resource group will have online.

The cluster file systems configured in the HAStoragePlus resource type for zone clusters should be authorized for use in zone clusters using the `clzonecluster(1CL)`. For more information, see [“How to Add a Cluster File System to a Zone Cluster” in *Oracle Solaris Cluster Software Installation Guide*](#).

Sample Entries in /etc/vfstab for Cluster File Systems

The following examples show entries in the `/etc/vfstab` file for global devices that are to be used for cluster file systems.

Note – The entries in the `/etc/vfstab` file for cluster file systems should contain the `global` keyword in the mount options.

EXAMPLE 2-30 Entries in `/etc/vfstab` for a Global Device With Solaris Volume Manager

This example shows entries in the `/etc/vfstab` file for a global device that uses Solaris Volume Manager.

```
/dev/md/kappa-1/dsk/d0    /dev/md/kappa-1/rdisk/d0
/global/local-fs/nfs ufs      5  yes      logging,global
```

EXAMPLE 2-31 Entries in `/etc/vfstab` for a Global Device With VxVM

This example shows entries in the `/etc/vfstab` file for a global device that uses VxVM.

```
/dev/vx/dsk/kappa-1/appvol    /dev/vx/rdsk/kappa-1/appvol
/global/local-fs/nfs vxfs      5 yes      log,global
```

▼ How to Set Up the HAStoragePlus Resource for Cluster File Systems

- 1 On any node in the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Create a failover or scalable resource group as desired.

Perform the following step to create a failover group.

```
# clresourcegroup create resource-group
```

Perform the following step to create a scalable group.

```
# clresourcegroup create -S [-p Maximum primaries=m] [-p Desired primaries=n] \
[-n node-zone-list] resource-group
```

- 3 Register the HAStoragePlus resource type.

```
# clresourcetype register SUNW.HAStoragePlus
```

- 4 Create the HAStoragePlus resource and define the filesystem mount points.

```
# clresource create -g resource-group -t SUNW.HAStoragePlus \
-p FileSystemMountPoints="mount-point-list" hasp-resource
```

The resource is created in the enabled state.

- 5 Add the data service resources to *resource-group*, and set their dependency to *hasp-resource*.

```
# clresource set -p Resource_dependencies_offline_restart= \
hasp-resource application-resource
```

- 6 Bring online and in a managed state the resource group that contains the HAStoragePlus resource.

```
# clresourcegroup online -M resource-group
```

Example 2-32 Setting up the HAStoragePlus Resource Type with a Cluster File System in a Global Cluster

This example shows how to configure the HAStoragePlus resource with a cluster file system `/global/ufs` in a global cluster for a failover resource group.

```

phys-schost-1# vi /etc/vfstab
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck      point      type     pass     at boot    options
#
/dev/md/apachedg/dsk/d0 /dev/md/apachedg/rdisk/d0 /global/ufs ufs 2 yes global, logging
# clresourcegroup create hasp-rg
# clresourcetype register SUNW.HAStoragePlus
# clresource create -g hasp-rg -t SUNW.HAStoragePlus -p \
FileSystemMountPoints=/global/ufs hasp-rs
# clresourcegroup online -M hasp-rg

```

Example 2–33 Setting up the HAStoragePlus Resource Type with a Cluster File System in a Zone Cluster

This example shows how to configure the HAStoragePlus resource with a cluster file system /global/ufs in a zone cluster for a scalable resource group. The cluster file system is available for the zone cluster nodes on the mount point /zone/ufs. This example configuration makes the global file system /global/ufs mounted in a global cluster and later loopback mounted on two zone-cluster nodes where the resource group is online.

```

phys-schost-1# vi /etc/vfstab
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck      point      type     pass     at boot    options
#
/dev/md/apachedg/dsk/d0 /dev/md/apachedg/rdisk/d0 /global/ufs ufs 2 yes global, logging
# clzonecluster configure sczone
clzc:sczone> add fs
clzc:sczone:fs> set dir=/zone/ufs
clzc:sczone:fs> set special=/global/ufs
clzc:sczone:fs> set type=lofs
clzc:sczone:fs> end
clzc:sczone:fs> exit
# clresourcegroup create -Z sczone -p desired_primaries=2 -p maximum_primaries=2 hasp-rg
# clresourcetype register -Z sczone SUNW.HAStoragePlus
# clresource create -Z sczone -g hasp-rg -t SUNW.HAStoragePlus -p
FileSystemMountPoints=/zone/ufs hasp-rs
# clresourcegroup online -Z sczone -M hasp-rg

```

▼ How to Delete an HAStoragePlus Resource Type for Cluster File Systems

- Disable and delete the HAStoragePlus resource configured for cluster file systems.

```
# clresource delete -F -g resource-group -t SUNW.HAStoragePlus resource
```

Enabling Highly Available Local File Systems

Using a highly available local file system improves the performance of I/O intensive data services. To make a local file system highly available in an Oracle Solaris Cluster environment, use the `HASStoragePlus` resource type.

You can specify global or local file systems. Global file systems are accessible from all nodes in a cluster. Local file systems are accessible from a single cluster node. Local file systems that are managed by a `SUNW.HASStoragePlus` resource are mounted on a single cluster node. These local file systems require the underlying devices to be Oracle Solaris Cluster global devices.

These file system mount points are defined in the format `paths [, . . .]`. You can specify both the path in a global-cluster non-voting node and the path in a global-cluster voting node, in this format:

Non-GlobalZonePath:GlobalZonePath

The global-cluster voting node path is optional. If you do not specify a global-cluster voting node path, Oracle Solaris Cluster assumes that the path in the global-cluster non-voting node and in the global-cluster voting node are the same. If you specify the path as *Non-GlobalZonePath:GlobalZonePath*, you must specify *GlobalZonePath* in the global-cluster voting node's `/etc/vfstab`.

The default setting for this property is an empty list.

You can use the `SUNW.HASStoragePlus` resource type to make a file system available to a global-cluster non-voting node. To enable the `SUNW.HASStoragePlus` resource type to do this, you must create a mount point in the global-cluster voting node and in the global-cluster non-voting node. The `SUNW.HASStoragePlus` resource type makes the file system available to the global-cluster non-voting node by mounting the file system in the global cluster. The resource type then performs a loopback mount in the global-cluster non-voting node. Each file system mount point should have an equivalent entry in `/etc/vfstab` on all cluster nodes and in the global-cluster voting node. The `SUNW.HASStoragePlus` resource type does not check `/etc/vfstab` in global-cluster non-voting nodes.

You can use the `SUNW.HASStoragePlus` resource type to make a file system available to zone cluster nodes. The file systems configured in the `SUNW.HASStoragePlus` resource type for zone clusters should be authorized for use in zone clusters using the `clzonecluster(1CL)` command. For more information, see [“Adding File Systems to a Zone Cluster” in Oracle Solaris Cluster Software Installation Guide](#).

Note – Local file systems include the UNIX File System (UFS), Quick File System (QFS), Veritas File System (VxFS), and Solaris ZFS (Zettabyte File System). Systems with Solaris ZFS are only mounted directly into the non-global zone.

The instructions for each Oracle Solaris Cluster data service that is I/O intensive explain how to configure the data service to operate with the `HASStoragePlus` resource type. For more information, see the individual Oracle Solaris Cluster data service guides.

Note – Do *not* use the `HASStoragePlus` resource type to make a root file system highly available.

Oracle Solaris Cluster provides the following tools for setting up the `HASStoragePlus` resource type to make local file systems highly available:

- **Oracle Solaris Cluster Manager.** For more information, see the Oracle Solaris Cluster Manager online help.
- **The `clsetup(1CL)` utility.**
- **Oracle Solaris Cluster maintenance commands.**

Oracle Solaris Cluster Manager and the `clsetup` utility enable you to add resources to the resource group interactively. Configuring these resources interactively reduces the possibility for configuration errors that might result from command syntax errors or omissions. Oracle Solaris Cluster Manager and the `clsetup` utility ensure that all required resources are created and that all required dependencies between resources are set.

Configuration Requirements for Highly Available Local File Systems

Any file system on multihost disks must be accessible from any host that is directly connected to those multihost disks. To meet this requirement, configure the highly available local file system as follows:

- Ensure that the disk partitions of the local file system reside on global devices.
- Set the `AffinityOn` extension property of the `HASStoragePlus` resource that specifies these global devices to `True`.
The `Zpool`s extension property of the `HASStoragePlus` resource ignores the `AffinityOn` extension property.
- Create the `HASStoragePlus` resource in a failover resource group.
- Ensure that the failback settings for the device groups and the resource group that contains the `HASStoragePlus` resource are identical.

Note – The use of a volume manager with the global devices for a highly available local file system is optional.

Format of Device Names for Devices Without a Volume Manager

If you are not using a volume manager, use the appropriate format for the name of the underlying storage device. The format to use depends on the type of storage device as follows:

- For block devices: `/dev/global/dsk/dDsS`
- For raw devices: `/dev/global/rdsk/dDsS`

The replaceable elements in these device names are as follows:

- *D* is an integer that specifies the device ID (DID) instance number.
- *S* is an integer that specifies the slice number.

Sample Entries in `/etc/vfstab` for Highly Available Local File Systems

The following examples show entries in the `/etc/vfstab` file for global devices that are to be used for highly available local file systems.

Note – Solaris ZFS (Zettabyte File System) does not use the `/etc/vfstab` file.

EXAMPLE 2-34 Entries in `/etc/vfstab` for a Global Device Without a Volume Manager

This example shows entries in the `/etc/vfstab` file for a global device on a physical disk without a volume manager.

```
/dev/global/dsk/d1s0      /dev/global/rdsk/d1s0
/global/local-fs/nfs ufs   5   no   logging
```

EXAMPLE 2-35 Entries in `/etc/vfstab` for a Global Device With Solaris Volume Manager

This example shows entries in the `/etc/vfstab` file for a global device that uses Solaris Volume Manager.

```
/dev/md/kappa-1/dsk/d0    /dev/md/kappa-1/rdsk/d0
/global/local-fs/nfs ufs   5   no   logging
```

EXAMPLE 2-36 Entries in `/etc/vfstab` for a Global Device With VxVM

This example shows entries in the `/etc/vfstab` file for a global device that uses VxVM.

```
/dev/vx/dsk/kappa-1/appvol /dev/vx/rdsk/kappa-1/appvol
/global/local-fs/nfs vxfs   5   no   log
```

Note – The same file system entries must be added to the zone cluster configuration when you configure the file system for a zone cluster using the `SUNW.HAStoragePlus` resource type.

▼ How to Set Up the HAStoragePlus Resource Type by Using the `clsetup` Utility

The following instructions explain how to set up the `HAStoragePlus` resource type by using the `clsetup` utility. Perform this procedure from any global-cluster voting node.

This procedure provides the long forms of the Oracle Solaris Cluster maintenance commands. Most commands also have short forms. Except for the forms of the command names, the commands are identical.

Before You Begin Ensure that the following prerequisites are met:

- Ensure that the required volumes, disk groups and file systems are created.

1 Become superuser on any cluster voting node.

2 Start the `clsetup` utility.

```
# clsetup
```

The `clsetup` main menu is displayed.

3 Type the number that corresponds to the option for data services and press Return.

The Data Services menu is displayed.

4 Type the number that corresponds to the option for configuring highly available storage and press Return.

The `clsetup` utility displays the list of prerequisites for performing this task.

5 Verify that the prerequisites are met, and press Return to continue.

The `clsetup` utility displays a list of the cluster nodes that can master the highly available `HAStoragePlus` resource.

6 Select the nodes that can master the highly available `HAStoragePlus` resource.

- To accept the default selection of all listed nodes in an arbitrary order, type `a` and press Return.

- **To select a subset of the listed nodes, type a comma-separated or space-separated list of the numbers that correspond to the nodes. Then press Return.**

Ensure that the nodes are listed in the order in which the nodes are to appear in the HASToragePlus resource group's node list. The first node in the list is the primary node of this resource group.

- **To select all nodes in a particular order, type a comma-separated or space-separated ordered list of the numbers that correspond to the nodes and press Return.**

7 To confirm your selection of nodes, type `d` and press Return.

The `clsetup` utility displays a list of types of shared storage type where data is to be stored.

8 Type the numbers that correspond to type of shared storage that you are using for storing the data and press Return.

The `clsetup` utility displays the file system mount points that are configured in the cluster. If there are no existing mount points, the `clsetup` utility allows you to define a new mount point.

9 Specify the default mount directory, the raw device path, the Global Mount option and the Check File System Periodically option and press Return.

The `clsetup` utility returns you the properties of the mount point that the utility will create.

10 To create the mount point, type `d` and press Return.

The `clsetup` utility displays the available file system mount points.

Note – You can use the `c` option to define another new mount point.

11 Select the file system mount points.

- **To accept the default selection of all listed file system mount points in an arbitrary order, type `a` and press Return.**
- **To select a subset of the listed file system mount points, type a comma-separated or space-separated list of the numbers that correspond to the file system mount points and press Return.**

12 To confirm your selection of nodes, type `d` and press Return.

The `clsetup` utility displays the global disk sets and device groups that are configured in the cluster.

13 Select the global device groups.

- To accept the default selection of all listed device groups in an arbitrary order, type **a** and press Return.
- To select a subset of the listed device groups, type a comma-separated or space-separated list of the numbers that correspond to the device groups and press Return.

14 To confirm your selection of nodes, type **d and press Return.**

The `clsetup` utility displays the names of the Oracle Solaris Cluster objects that the utility will create.

15 If you require a different name for any Oracle Solaris Cluster object, change the name as follows.**a. Type the number that corresponds to the name that you are changing and press Return.**

The `clsetup` utility displays a screen where you can specify the new name.

b. At the New Value prompt, type the new name and press Return.

The `clsetup` utility returns you to the list of the names of the Oracle Solaris Cluster objects that the utility will create.

16 To confirm your selection of Oracle Solaris Cluster object names, type **d and press Return.**

The `clsetup` utility displays information about the Oracle Solaris Cluster configuration that the utility will create.

17 To create the configuration, type **c and Press Return.**

The `clsetup` utility displays a progress message to indicate that the utility is running commands to create the configuration. When configuration is complete, the `clsetup` utility displays the commands that the utility ran to create the configuration.

18 (Optional) Type **q and press Return repeatedly until you quit the `clsetup` utility.**

If you prefer, you can leave the `clsetup` utility running while you perform other required tasks before using the utility again. If you choose to quit `clsetup`, the utility recognizes your existing resource group when you restart the utility.

19 Verify that the `HASStoragePlus` resource has been created.

Use the `clresource(1CL)` utility for this purpose.

```
# clresource show name_of_rg
```

▼ How to Set Up the HASStoragePlus Resource Type to Make File Systems Highly Available Other Than Solaris ZFS

The following procedure explains how to set up the HASStoragePlus resource type to make file systems other than Solaris ZFS highly available.

- 1 On any node in the global cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Create a failover resource group.

```
# clresourcegroup create resource-group
```
- 3 Register the HASStoragePlus resource type.

```
# clresourcetype register SUNW.HASStoragePlus
```
- 4 Create the HASStoragePlus resource and define the file system mount points.

```
# clresource create -g resource-group \  
-t SUNW.HASStoragePlus -p FileSystemMountPoints=mount-point-list hasp-resource
```
- 5 Bring online and in a managed state the resource group that contains the HASStoragePlus resource.

```
# clresourcegroup online -M resource-group
```

Example 2–37 Setting Up the HASStoragePlus Resource Type to Make a UFS File System Highly Available for the Global Cluster

This example assumes that the file system `/web-1` is configured to the HASStoragePlus resource to make the file system highly available for the global cluster.

```
phys-schost-1# vi /etc/vfstab
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck      point      type     pass      at boot    options
#
# /dev/md/apachedg/dsk/d0 /dev/md/apachedg/rdisk/d0 /web-1 ufs 2 no logging
# clresourcegroup create hasp-rg
# clresourcetype register SUNW.HASStoragePlus
# clresource create -g hasp-rg -t SUNW.HASStoragePlus -p FileSystemMountPoints=/global/ufs-1 hasp-rs
# clresourcegroup online -M hasp-rg
```

Example 2–38 Setting Up the HASStoragePlus Resource Type to Make a UFS File System Highly Available for a Zone Cluster

This example assumes that the file system `/web-1` is configured to the HASStoragePlus resource to make the file system highly available for a zone cluster `sczone`. When a local file system is

configured as a highly available file system for a zone cluster using the SUNW.HAStoragePlus resource type, the HAStoragePlus resource reads the file system information in the zone cluster configuration.

```
# clzonecluster configure sczone
clzc:sczone> add fs
clzc:sczone:fs> set dir=/web-1
clzc:sczone:fs> set special=/dev/md/apachedg/dsk/d0
clzc:sczone:fs> set raw=/dev/md/apachedg/rdisk/d0
clzc:sczone:fs> set type=ufs
clzc:sczone:fs> add options [logging]
clzc:sczone:fs> end
clzc:sczone:fs> exit
# clresourcegroup create -Z sczone hasp-rg
# clresourcetype register -Z sczone SUNW.HAStoragePlus
# clresource create -Z sczone -g hasp-rg \
-t SUNW.HAStoragePlus -p FileSystemMountPoints=/web-1 hasp-rs
# clresourcegroup online -Z sczone -M hasp-rg
```

▼ How to Set Up the HAStoragePlus Resource Type to Make a Local Solaris ZFS Highly Available

You perform the following primary tasks to make a local Solaris ZFS (Zettabyte File System) highly available:

- Create a ZFS storage pool.
- Create a ZFS file system in that ZFS storage pool.
- Set up the HAStoragePlus resource that manages the ZFS storage pool.

This section describes how to complete both tasks.

1 Create a ZFS storage pool.



Caution – Do not add a configured quorum device to a ZFS storage pool. When a configured quorum device is added to a storage pool, the disk is relabeled as an EFI disk, the quorum configuration information is lost, and the disk no longer provides a quorum vote to the cluster. Once a disk is in a storage pool, you can configure that disk as a quorum device. Or, you can unconfigure the disk, add it to the storage pool, then reconfigure the disk as a quorum device.

Observe the following requirements when you create a ZFS storage pool in an Oracle Solaris Cluster configuration:

- Ensure that all of the devices from which you create a ZFS storage pool are accessible from all nodes in the cluster. These nodes must be configured in the node list of the resource group to which the HAStoragePlus resource belongs.
- Ensure that the Oracle Solaris device identifier that you specify to the `zpool (1M)` command, for example `/dev/dsk/c0t0d0`, is visible to the `cldevice list -v` command.

Note – The ZFS storage pool can be created using a full disk or a disk slice. It is preferred to create a ZFS storage pool using a full disk by specifying an Oracle Solaris logical device as ZFS file system performs better by enabling the disk write cache. ZFS file system labels the disk with EFI when a full disk is provided.

See “Creating a ZFS Storage Pool” in *Solaris ZFS Administration Guide* for information about how to create a ZFS storage pool.

2 In the ZFS storage pool that you just created, create a ZFS file system.

You can create more than one ZFS file system in the same ZFS storage pool.

Note – HAStoragePlus does not support file systems created on ZFS file system volumes.

Do not place a ZFS file system in the `FilesystemMountPoints` extension property.

See “Creating a ZFS File System Hierarchy” in *Solaris ZFS Administration Guide* for information about how to create a ZFS file system in a ZFS storage pool.

3 On any node in the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

4 Create a failover resource group.

```
# clresourcegroup create resource-group
```

5 Register the HAStoragePlus resource type.

```
# clresourcetype register SUNW.HAStoragePlus
```

6 Create an HAStoragePlus resource for the local ZFS file system.

```
# clresource create -g resource-group -t SUNW.HAStoragePlus \  
-p Zpools=zpool -p ZpoolsSearchDir=/dev/dsk \  
resource
```

The default location to search for devices of ZFS storage pools is `/dev/dsk`. It can be overridden by using the `ZpoolsSearchDir` extension property.

The resource is created in the enabled state.

7 Bring online and in a managed state the resource group that contains the HAStoragePlus resource.

```
# clresourcegroup online -M resource-group
```

Example 2–39 Setting Up the HAStoragePlus Resource Type to Make a Local ZFS Highly Available for the Global Cluster

The following example shows the commands to make a local ZFS file system highly available.

```
phys-schost-1% su
Password:
# cldevice list -v
```

DID Device	Full Device Path
-----	-----
d1	phys-schost-1:/dev/rdisk/c0t0d0
d2	phys-schost-1:/dev/rdisk/c0t1d0
d3	phys-schost-1:/dev/rdisk/c1t8d0
d3	phys-schost-2:/dev/rdisk/c1t8d0
d4	phys-schost-1:/dev/rdisk/c1t9d0
d4	phys-schost-2:/dev/rdisk/c1t9d0
d5	phys-schost-1:/dev/rdisk/c1t10d0
d5	phys-schost-2:/dev/rdisk/c1t10d0
d6	phys-schost-1:/dev/rdisk/c1t11d0
d6	phys-schost-2:/dev/rdisk/c1t11d0
d7	phys-schost-2:/dev/rdisk/c0t0d0
d8	phys-schost-2:/dev/rdisk/c0t1d0

you can create a ZFS storage pool using a disk slice by specifying a Solaris device identifier:

```
# zpool create HAZpool c1t8d0s2
```

or you can create a ZFS storage pool using disk slice by specifying a logical device identifier

```
# zpool create HAZpool /dev/did/dsk/d3s2
```

```
# zfs create HAZpool/export
```

```
# zfs create HAZpool/export/home
```

```
# clresourcegroup create hasp-rg
```

```
# clresourcetype register SUNW.HAStoragePlus
```

```
# clresource create -g hasp-rg -t SUNW.HAStoragePlus -p Zpools=HAZpool hasp-rs
```

```
# clresourcegroup online -M hasp-rg
```

Example 2–40 Setting Up the HAStoragePlus Resource Type to Make a Local ZFS Highly Available for a Zone Cluster

The following example shows the steps to make a local ZFS file system highly available in a zone cluster *sczone*.

```
phys-schost-1# cldevice list -v
# zpool create HAZpool c1t8d0
# zfs create HAZpool/export
# zfs create HAZpool/export/home
# clzonecluster configure sczone
clzc:sczone> add dataset
clzc:sczone:fs> set name=HAZpool
clzc:sczone:fs> end
clzc:sczone:fs> exit
# clresourcegroup create -Z sczone hasp-rg
# clresourcetype register -Z sczone SUNW.HAStoragePlus
# clresource create -Z sczone -g hasp-rg -t SUNW.HAStoragePlus \
-p Zpools=HAZpool hasp-rs
# clresourcegroup online -Z -sczone -M hasp-rg
```

▼ How to Delete an HASStoragePlus Resource That Makes a Local Solaris ZFS Highly Available

- Disable and delete the HASStoragePlus resource that makes a local Solaris ZFS (Zettabyte File System) highly available.

```
# clresource delete -F -g resource-group -t SUNW.HASStoragePlus resource
```

Sharing a Failover File System Across Zone Clusters

You can use the SUNW.HASStoragePlus resource type to share a failover file system directory managed by a global cluster resource to a zone cluster. This method consolidates the storage and shares a failover file system with different applications running on different zone clusters. For information on adding a file system to a zone cluster, see [“Adding File Systems to a Zone Cluster”](#) in *Oracle Solaris Cluster Software Installation Guide*.

This section explains the requirements and procedures for sharing a failover file system directory across zone clusters.

Configuration Requirements for Sharing a Failover File System Directory to a Zone Cluster

The directory of a failover file system managed by a global cluster resource can be shared to a zone cluster. To share a failover file system directory, the configuration must meet the following requirements:

- Create an HASStoragePlus resource in a failover resource group in a global cluster with the file system where the directory to be shared belongs.
- The directory of the failover file system that you want to share must be configured to a zone cluster as an lofs file system.
- Create an HASStoragePlus resource in a failover resource group in a zone cluster with the lofs file system.
- The zone cluster resource must have an offline restart dependency on the global cluster resource.
- The zone cluster resource's resource group must have a strong positive affinity or strong positive affinity with failover delegation on the global cluster resource's resource group.

Note – The applications sharing a failover file system will experience an availability impact due to co-location of the applications. An application failure on a node and its intent to fail over might have a cascading effect on other applications and the applications would be forced to fail over to another node. Mitigate the problem by reducing the number of applications that are sharing the file system. If the file system that is being shared is UFS or VxFS, you can choose to configure the cluster file system to a zone cluster. See [“How to Set Up the HAStoragePlus Resource for Cluster File Systems” on page 122.](#)

▼ How to Set Up the HAStorage Plus Resource Type to Share a Failover File System Directory to a Zone Cluster

The following procedure explains how to set up the HAStoragePlus resource type to share a failover file system (for example, UFS, VxFS, and QFS) or a ZFS pool directory to a zone cluster called *zoneclustername*.

- 1 On any node in the global cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization. Perform the steps from a node in the global cluster, because the dependencies and affinities from a zone cluster to a global cluster can only be set by an authorized global cluster node administrator.
- 2 Create a failover resource group in the global cluster.

```
# clresourcegroup create gc-hasp-resource-group
```
- 3 Register the HAStoragePlus resource type in the global cluster.

```
# clresourcetype register SUNW.HAStoragePlus
```
- 4 Create an HAStoragePlus resource in a failover resource group of the global cluster with the failover file system that contains the directory that you want to share to a zone cluster.

```
# clresource create -g gc-hasp-resource-group -t HAStoragePlus \
-p FilesystemMountPoints=mount-point \
-p Zpools=pool gc-hasp-resource
```
- 5 Bring the global cluster failover resource group online.

```
# clresourcegroup online -M gc-hasp-resource-group
```
- 6 Configure the directory of the failover file system that is being shared to the zone cluster as an `lofs` file system.

```
# clzonecluster configure zoneclustername
clzc:zoneclustername> add fs
clzc:zoneclustername:fs> set dir = shared-dir-mount-point-in-zc
clzc:zoneclustername:fs> set special = shared-directory
```

```
clzc:zoneclustername:fs> set type = lofs
clzc:zoneclustername:fs> end
clzc:zoneclustername> exit
#
```

- 7 Create a failover resource group in the zone cluster that has a strong positive affinity or strong positive affinity with failover delegation on the failover resource group of the global cluster.

```
# clresourcegroup create -Z zoneclustername \
-p RG_affinities=++global:gc-hasp-resource-group \
zc-hasp-resource-group
OR
# clresourcegroup create -Z zoneclustername \
-p RG_affinities=+++global:gc-hasp-resource-group zc-hasp-resource-group
```

- 8 Register the HASStoragePlus resource type in the zone cluster.

```
# clresourcetype register -Z zoneclustername SUNW.HASStoragePlus
```

- 9 Create an HASStoragePlus resource in a failover resource group of the zone cluster. Configure the zone cluster with the lofs file system for a shared directory with a dependency on the global cluster resource that you want to share to the zone cluster.

```
# clresource create -Z zoneclustername -t SUNW.HASStoragePlus -g zc-hasp-resource-group \
-p FilesystemMountPoints=shared-dir-mount-point-in-zc \
-p Resource_dependencies_offline_restart=global:gc-hasp-resource zc-hasp-resource
```

- 10 Bring the zone cluster failover resource group online.

```
# clresourcegroup online -Z zoneclustername -M zc-hasp-resource-group
```

Example 2-41 Setting Up the HASStoragePlus Resource Type to Share a UFS Failover File System Directory to a Zone Cluster

The following example shows how to share the /local/fs/home directory of a UFS failover file system (/local/fs) to a zone cluster called *sczone*.

```
# clresourcegroup create gc-hasp-rg
# clresourcetype register -Z sczone SUNW.HASStoragePlus
# vi /etc/vfstab /dev/md/dg1/dsk/d0 /dev/md/dg1/rdisk/d0 /local/fs ufs 2 no logging
# clresource create -g gc-hasp-rg -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=/local/fs gc-hasp-rs
# clresourcegroup online -M gc-hasp-rg
```

The steps above ensure that the *gc-hasp-rs* resource running in the global cluster manages the failover file system /local/fs.

```
# clzonecluster configure sczone
clzc:sczone> add fs
clzc:sczone:fs> set dir = /share/local/fs/home
clzc:sczone:fs> set special = /local/fs/home
clzc:sczone:fs> set type = lofs
clzc:sczone:fs> end
clzc:sczone> exit
```


The configuration above makes the failover file system's directory */local/fs/home* available in the zone cluster *sczone* at mount point */share/local/fs/home*.

```
# clresourcegroup create -Z sczone \
-p RG_affinities=++global:gc-hasp-rg zc-hasp-rg
# clresourcetype register -Z sczone SUNW.HAStoragePlus
# clresource create -Z sczone -t HAStoragePlus -g zc-hasp-rg \
-p FilesystemMountPoints=/share/local/fs/home \
-p Resource_dependencies_offline_restart=global:gc-hasp-rs zc-hasp-rs
# clresourcegroup online -Z sczone -M zc-hasp-rg
```

The steps above create a zone cluster resource that manages the shared directory as an *lofs* file system. The steps in this example are applicable to VxFS and QFS file systems.

Example 2-42 Setting Up the HAStoragePlus Resource Type to Share a ZFS Pool Directory to a Zone Cluster

The following example shows how to share the ZFS pool "tank" directory */tank/home* to a zone cluster called *sczone*.

```
# clresourcegroup create gc-hasp-rg
# clresourcetype register SUNW.HAStoragePlus
# clresource create -g gc-hasp-rg -t SUNW.HAStoragePlus \
-p Zpools=tank gc-hasp-rs
# clresourcegroup online -M gc-hasp-rg
```

The steps above ensure that the ZFS failover file system is managed by *gc-hasp-rs* running in the global cluster.

```
# clzonecluster configure sczone
clzc:sczone> add fs
clzc:sczone:fs> set dir = /share/tank/home
clzc:sczone:fs> set special = /tank/home
clzc:sczone:fs> set type = lofs
clzc:sczone:fs> end
clzc:sczone> exit
#
```

The configuration above makes the ZFS pool "tank" directory */tank/home* available in the zone cluster *sczone* at mountpoint */share/tank/home*.

```
# clresourcegroup create -Z sczone \
-p RG_affinities=++global:gc-hasp-rg zc-hasp-rg
# clresourcetype register -Z sczone SUNW.HAStoragePlus
# clresource create -Z sczone -t HAStoragePlus -g zc-hasp-rg \
-p FilesystemMountPoints=/share/tank/home \
-p Resource_dependencies_offline_restart=global:gc-hasp-rs zc-hasp-rs
# clresourcegroup online -Z sczone -M zc-hasp-rg
```

The steps above create a zone cluster resource that manages the shared directory as an *lofs* file system.

Upgrading From HAStorage to HAStoragePlus

HAStorage is not supported in the current release of Oracle Solaris Cluster software. Equivalent functionality is supported by HAStoragePlus. For instructions for upgrading from HAStorage to HAStorage, see the subsections that follow.

Note – The resource groups that have the HAStorage resource configured will be in STOP_FAILED state since HAStorage is no longer supported. Clear the ERROR_STOP_FAILED flag for the resource and follow the instructions to upgrade HAStorage to HAStoragePlus.

▼ How to Upgrade From HAStorage to HAStoragePlus When Using Device Groups or CFS

The following example uses a simple HA-NFS resource that is active with HAStorage. The ServicePaths are the disk group nfsdg and the AffinityOn property is True. Furthermore, the HA-NFS resource has Resource_Dependencies set to the HAStorage resource.

- 1 **Bring offline the resource group nfs1-rg.**
`# clresourcegroup offline nfs1-rg`
- 2 **Remove the dependencies the application resources has on HAStorage.**
`# clresource set -p Resource_Dependencies="" nfsserver-rs`
- 3 **Disable the HAStorage resource.**
`# clresource disable nfs1storage-rs`
- 4 **Remove the HAStorage resource from the application resource group.**
`# clresource delete nfs1storage-rs`
- 5 **Unregister the HAStorage resource type.**
`# clresourcetype unregister SUNW.HAStorage`
- 6 **Register the HAStoragePlus resource type.**
`# clresourcetype register SUNW.HAStoragePlus`
- 7 **Create the HAStoragePlus resource.**

Note – Instead of using the ServicePaths property of HAStorage, you must use the FilesystemMountPoints property or GlobalDevicePaths property of HAStoragePlus.

- To specify the mount point of a file system, type the following command.

The FilesystemMountPoints extension property must match the sequence that is specified in /etc/vfstab.

```
# clresource create -g nfs1-rg -t \
SUNW.HAStoragePlus -p FilesystemMountPoints=/global/nfsdata -p \
AffinityOn=True nfs1-hastp-rs
```

- To specify global device paths, type the following command.

```
# clresource create -g nfs1-rg -t \
SUNW.HAStoragePlus -p GlobalDevicePaths=nfsdg -p AffinityOn=True nfs1-hastp-rs
```

The resource is created in the enabled state.

- 8 Disable the application server resource.

```
# clresource disable nfsserver-rs
```

- 9 Bring online the nfs1-rg group on a cluster node.

```
# clresourcegroup online nfs1-rg
```

- 10 Set up the dependencies between the application server and HAStoragePlus.

```
# clresource set -p Resource_dependencies_offline_restart=nfs1-hastp-rs nfsserver-rs
```

- 11 Bring online the nfs1-rg group on a cluster node.

```
# clresourcegroup online -eM nfs1-rg
```

▼ How to Upgrade From HAStorage With CFS to HAStoragePlus With Highly Available Local File System

The following example uses a simple HA-NFS resource that is active with HAStorage. The ServicePaths are the disk group nfsdg and the AffinityOn property is True. Furthermore, the HA-NFS resource has Resource_Dependencies set to HAStorage resource.

- 1 Remove the dependencies the application resource has on HAStorage resource.

```
# clresource set -p Resource_Dependencies="" nfsserver-rs
```

- 2 Disable the HAStorage resource.

```
# clresource disable nfs1storage-rs
```

- 3 Remove the HAStorage resource from the application resource group.

```
# clresource delete nfs1storage-rs
```

4 Unregister the HASStorage resource type.

```
# clresourcetype unregister SUNW.HASStorage
```

5 Modify /etc/vfstab to remove the global flag and change “mount at boot” to “no”.

6 Create the HASStoragePlus resource.

Note – Instead of using the ServicePaths property of HASStorage, you must use the FilesystemMountPoints property or GlobalDevicePaths property of HASStoragePlus.

■ **To specify the mount point of a file system, type the following command.**

The FilesystemMountPoints extension property must match the sequence that is specified in /etc/vfstab.

```
# clresource create -g nfs1-rg -t \
SUNW.HASStoragePlus -p FilesystemMountPoints=/global/nfsdata -p \
AffinityOn=True nfs1-hastp-rs
```

■ **To specify global device paths, type the following command.**

```
# clresource create -g nfs1-rg -t \
SUNW.HASStoragePlus -p GlobalDevicePaths=nfsdg -p AffinityOn=True nfs1-hastp-rs
```

The resource is created in the enabled state.

7 Disable the application server resource.

```
# clresource disable nfsserver-rs
```

8 Bring online the nfs1-rg group on a cluster node.

```
# clresourcegroup online nfs1-rg
```

9 Set up the dependencies between the application server and HASStoragePlus.

```
# clresource set -p Resource_dependencies_offline_restart=nfs1-hastp-rs nfsserver-rs
```

10 Bring online the nfs1-rg group on a cluster node.

```
# clresourcegroup online -eM nfs1-rg
```

Modifying Online the Resource for a Highly Available File System

You might need a highly available file system to remain available while you are modifying the resource that represents the file system. For example, you might need the file system to remain available because storage is being provisioned dynamically. In this situation, modify the resource that represents the highly available file system while the resource is online.

In the Oracle Solaris Cluster environment, a highly available file system is represented by an `HASStoragePlus` resource. Oracle Solaris Cluster enables you to modify an online `HASStoragePlus` resource as follows:

- Adding file systems to the `HASStoragePlus` resource
- Removing file systems from the `HASStoragePlus` resource

Note – Oracle Solaris Cluster software does not enable you to rename a file system while the file system is online.

Note – When you remove the file systems configured in the `HASStoragePlus` resources for a zone cluster, you also need to remove the file system configuration from the zone cluster. For information about removing a file system from a zone cluster, see [“How to Remove a File System From a Zone Cluster”](#) in *Oracle Solaris Cluster System Administration Guide*.

▼ How to Add File Systems Other Than Solaris ZFS to an Online `HASStoragePlus` Resource

When you add a local or cluster file system to an `HASStoragePlus` resource, the `HASStoragePlus` resource automatically mounts the file system.

- 1 **On one node of the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.**
- 2 **In the `/etc/vfstab` file on each node of the cluster, add an entry for the mount point of each file system that you are adding.**

For each entry, set the mount at boot field and the mount options field as follows:

- For local file systems
 - Set the mount at boot field to `no`.
 - Remove the `global` flag.
- For cluster file systems
 - If the file system is a cluster file system, set the mount options field to contain the `global` option.

- 3 **Retrieve the list of mount points for the file systems that the `HASStoragePlus` resource already manages.**

```
# scha_resource_get -O extension -R hasp-resource -G hasp-rg FileSystemMountPoints
```

`-R hasp-resource` Specifies the `HASStoragePlus` resource to which you are adding file systems

-G hasp-rg Specifies the resource group that contains the HASStoragePlus resource

4 Modify the FileSystemMountPoints extension property of the HASStoragePlus resource to contain the following mount points:

- The mount points of the file systems that the HASStoragePlus resource already manages
- The mount points of the file systems that you are adding to the HASStoragePlus resource

```
# clresource set -p FileSystemMountPoints="mount-point-list" hasp-resource
```

```
-p FileSystemMountPoints="mount-point-list"
```

Specifies a comma-separated list of mount points of the file systems that the HASStoragePlus resource already manages and the mount points of the file systems that you are adding. The format of each entry in the list is LocalZonePath:GlobalZonePath. In this format, the global path is optional. If the global path is not specified, the global path is the same as the local path.

hasp-resource

Specifies the HASStoragePlus resource to which you are adding file systems.

5 Confirm that you have a match between the mount point list of the HASStoragePlus resource and the list that you specified in [Step 4](#).

```
# scha_resource_get -O extension -R hasp-resource -G hasp-rg \
  FileSystemMountPoints
```

-R hasp-resource Specifies the HASStoragePlus resource to which you are adding file systems.

-G hasp-rg Specifies the resource group that contains the HASStoragePlus resource.

6 Confirm that the HASStoragePlus resource is online and not faulted.

If the HASStoragePlus resource is online and faulted, validation of the resource succeeded, but an attempt by HASStoragePlus to mount a file system failed.

```
# clresource status hasp-resource
```

Example 2-43 Adding a File System to an Online HASStoragePlus Resource

This example shows how to add a file system to an online HASStoragePlus resource.

- The HASStoragePlus resource is named rshasp and is contained in the resource group rghasp.
- The HASStoragePlus resource named rshasp already manages the file system whose mount point is /global/global-fs/fs.
- The mount point of the file system that is to be added is /global/local-fs/fs.

The example assumes that the `/etc/vfstab` file on each cluster node already contains an entry for the file system that is to be added.

```
# scha_resource_get -O extension -R rshasp -G rghasp FileSystemMountPoints
STRINGARRAY
/global/global-fs/fs
# clresource set \
-p FileSystemMountPoints="/global/global-fs/fs,/global/local-fs/fs"
# scha_resource_get -O extension -R rshasp -G rghasp FileSystemMountPoints rshasp
STRINGARRAY
/global/global-fs/fs
/global/local-fs/fs
# clresource status rshasp
```

=== Cluster Resources ===

Resource Name	Node Name	Status	Message
-----	-----	-----	-----
rshasp	node46	Offline	Offline
	node47	Online	Online

▼ How to Remove File Systems Other Than Solaris ZFS From an Online HASoragePlus Resource

When you remove a file system from an HASoragePlus resource, the HASoragePlus resource treats a local file system differently from a cluster file system.

- The HASoragePlus resource automatically unmounts a local file system.
- The HASoragePlus resource does not unmount the cluster file system.



Caution – Before removing a file system from an online HASoragePlus resource, ensure that no applications are using the file system. When you remove a file system from an online HASoragePlus resource, the file system might be forcibly unmounted. If a file system that an application is using is forcibly unmounted, the application might fail or hang.

- 1 On one node of the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Retrieve the list of mount points for the file systems that the HASoragePlus resource already manages.

```
# scha_resource_get -O extension -R hasp-resource -G hasp-rg FileSystemMountPoints
-R hasp-resource    Specifies the HASoragePlus resource from which you are removing file
                    systems.
-G hasp-rg          Specifies the resource group that contains the HASoragePlus resource.
```

- 3 **Modify the `FileSystemMountPoints` extension property of the `HASStoragePlus` resource to contain *only* the mount points of the file systems that are to remain in the `HASStoragePlus` resource.**

```
# clresource set -p FileSystemMountPoints="mount-point-list" hasp-resource
```

```
-p FileSystemMountPoints="mount-point-list"
```

Specifies a comma-separated list of mount points of the file systems that are to remain in the `HASStoragePlus` resource. This list must *not* include the mount points of the file systems that you are removing.

```
hasp-resource
```

Specifies the `HASStoragePlus` resource from which you are removing file systems.

- 4 **Confirm that you have a match between the mount point list of the `HASStoragePlus` resource and the list that you specified in [Step 3](#).**

```
# scha_resource_get -O extension -R hasp-resource -G hasp-rg \
FileSystemMountPoints
```

```
-R hasp-resource
```

Specifies the `HASStoragePlus` resource from which you are removing file systems.

```
-G hasp-rg
```

Specifies the resource group that contains the `HASStoragePlus` resource.

- 5 **Confirm that the `HASStoragePlus` resource is online and not faulted.**

If the `HASStoragePlus` resource is online and faulted, validation of the resource succeeded, but an attempt by `HASStoragePlus` to unmount a file system failed.

```
# clresource status hasp-resource
```

- 6 **(Optional) From the `/etc/vfstab` file on each node of the cluster, remove the entry for the mount point of each file system that you are removing.**

Example 2–44 Removing a File System From an Online `HASStoragePlus` Resource

This example shows how to remove a file system from an online `HASStoragePlus` resource.

- The `HASStoragePlus` resource is named `rshasp` and is contained in the resource group `rghasp`.
- The `HASStoragePlus` resource named `rshasp` already manages the file systems whose mount points are as follows:
 - `/global/global-fs/fs`
 - `/global/local-fs/fs`
- The mount point of the file system that is to be removed is `/global/local-fs/fs`.

```
# scha_resource_get -O extension -R rshasp -G rghasp FileSystemMountPoints
STRINGARRAY
```



```
/global/global-fs/fs
/global/local-fs/fs
# clresource set -p FileSystemMountPoints="/global/global-fs/fs"
# scha_resource_get -O extension -R rshasp -G rghasp FileSystemMountPoints rshasp
STRINGARRAY
/global/global-fs/fs
# clresource status rshasp
```

=== Cluster Resources ===

Resource Name	Node Name	Status	Message
-----	-----	-----	-----
rshasp	node46	Offline	Offline
	node47	Online	Online

▼ How to Add a Solaris ZFS Storage Pool to an Online HAStoragePlus Resource

When you add a Solaris ZFS (Zettabyte File System) storage pool to an online HAStoragePlus resource, the HAStoragePlus resource does the following:

- Imports the ZFS storage pool.
- Mounts all file systems in the ZFS storage pool.

- 1 On any node in the cluster, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Determine the ZFS storage pools that the HAStoragePlus resource already manages.

```
# clresource show -g hasp-resource-group -p Zpools hasp-resource
```

`-g hasp-resource-group` Specifies the resource group that contains the HAStoragePlus resource.

`hasp-resource` Specifies the HAStoragePlus resource to which you are adding the ZFS storage pool.

- 3 Add the new ZFS storage pool to the existing list of ZFS storage pools that the HAStoragePlus resource already manages.

```
# clresource set -p Zpools="zpools-list" hasp-resource
```

`-p Zpools="zpools-list"` Specifies a comma-separated list of existing ZFS storage pool names that the HAStoragePlus resource already manages and the new ZFS storage pool name that you want to add.

`hasp-resource` Specifies the HAStoragePlus resource to which you are adding the ZFS storage pool.

- 4 Compare the new list of ZFS storage pools that the **HASStoragePlus** resource manages with the list that you generated in [Step 2](#).

```
# clresource show -g hasp-resource-group -p Zpools hasp-resource
```

-g hasp-resource-group Specifies the resource group that contains the HASStoragePlus resource.

hasp-resource Specifies the HASStoragePlus resource to which you added the ZFS storage pool.

- 5 Confirm that the **HASStoragePlus** resource is online and not faulted.

If the HASStoragePlus resource is online but faulted, validation of the resource succeeded.

However, an attempt by the HASStoragePlus resource to import and mount the ZFS file system failed. In this case, you need to repeat the preceding set of steps.

```
# clresourcegroup status hasp-resource
```

▼ How to Remove a Solaris ZFS Storage Pool From an Online HASStoragePlus Resource

When you remove a Solaris ZFS (Zettabyte File System) storage pool from an online HASStoragePlus resource, the HASStoragePlus resource does the following:

- Unmounts the file systems in the ZFS storage pool.
- Exports the ZFS storage pool from the node.

- 1 On any node in the cluster, become superuser or assume a role that provides **solaris.cluster.modify** RBAC authorization.
- 2 Determine the ZFS storage pools that the **HASStoragePlus** resource already manages.

```
# clresource show -g hasp-resource-group -p Zpools hasp-resource
```

-g hasp-resource-group Specifies the resource group that contains the HASStoragePlus resource.

hasp-resource Specifies the HASStoragePlus resource from which you are removing the ZFS storage pool.

- 3 Remove the ZFS storage pool from the list of ZFS storage pools that the **HASStoragePlus** resource currently manages.

```
# clresource set -p Zpools="zpools-list" hasp-resource
```

-p Zpools="zpools-list" Specifies a comma-separated list of ZFS storage pool names that the HASStoragePlus resource currently manages, minus the ZFS storage pool name that you want to remove.

hasp-resource Specifies the HASStoragePlus resource from which you are removing the ZFS storage pool.

- 4 **Compare the new list of ZFS storage pools that the HASStoragePlus resource now manages with the list that you generated in [Step 2](#).**

```
# clresource show -g hasp-resource-group -p Zpools hasp-resource
```

-g *hasp-resource-group* Specifies the resource group that contains the HASStoragePlus resource.

hasp-resource Specifies the HASStoragePlus resource from which you removed the ZFS storage pool.

- 5 **Confirm that the HASStoragePlus resource is online and not faulted.**

If the HASStoragePlus resource is online but faulted, validation of the resource succeeded. However, an attempt by the HASStoragePlus resource to unmount and export the ZFS file system failed. In this case, you need to repeat the preceding set of steps.

```
# clresource status -t SUNW.HASStoragePlus +
```

▼ How to Recover From a Fault After Modifying the FileSystemMountPoints Property of an HASStoragePlus Resource

If a fault occurs during a modification of the FileSystemMountPoints extension property, the status of the HASStoragePlus resource is online and faulted. After the fault is corrected, the status of the HASStoragePlus resource is online.

- 1 **Determine the fault that caused the attempted modification to fail.**

```
# clresource status hasp-resource
```

The status message of the faulty HASStoragePlus resource indicates the fault. Possible faults are as follows:

- The device on which the file system should reside does not exist.
- An attempt by the fsck command to repair a file system failed.
- The mount point of a file system that you attempted to add does not exist.
- A file system that you attempted to add cannot be mounted.
- A file system that you attempted to remove cannot be unmounted.

- 2 **Correct the fault that caused the attempted modification to fail.**

- 3 Repeat the step to modify the `FileSystemMountPoints` extension property of the `HASStoragePlus` resource.
- ```
clresource set -p FileSystemMountPoints="mount-point-list" hasp-resource
```
- p `FileSystemMountPoints="mount-point-list"`  
Specifies a comma-separated list of mount points that you specified in the unsuccessful attempt to modify the highly available file system
- hasp-resource*  
Specifies the `HASStoragePlus` resource that you are modifying
- 4 Confirm that the `HASStoragePlus` resource is online and not faulted.
- ```
# clresource status
```

Example 2-45 Status of a Faulty `HASStoragePlus` Resource

This example shows the status of a faulty `HASStoragePlus` resource. This resource is faulty because an attempt by the `fsck` command to repair a file system failed.

```
# clresource status

=== Cluster Resources ===

Resource Name      Node Name      Status      Status Message
-----
rshasp             node46         Offline     Offline
                  node47         Online      Online Faulted - Failed to fsck: /mnt.
```

▼ **How to Recover From a Fault After Modifying the Zpools Property of an `HASStoragePlus` Resource**

If a fault occurs during a modification of the `Zpools` extension property, the status of the `HASStoragePlus` resource is online and faulted. After the fault is corrected, the status of the `HASStoragePlus` resource is online.

- 1 Determine the fault that caused the attempted modification to fail.
- ```
clresource status hasp-resource
```
- The status message of the faulty `HASStoragePlus` resource indicates the fault. Possible faults are as follows:
- The ZFS pool *zpool* failed to import.
  - The ZFS pool *zpool* failed to export.
- 2 Correct the fault that caused the attempted modification to fail.

3 Repeat the step to modify the Zpools extension property of the HAStoragePlus resource.

```
clresource set -p Zpools="zpools-list" hasp-resource
```

-p Zpools="zpools-list" Specifies a comma-separated list of ZFS storage pool names that the HAStoragePlus currently manages, minus the ZFS storage pool name that you want to remove.

hasp-resource Specifies the HAStoragePlus resource that you are modifying

4 Confirm that the HAStoragePlus resource is online and not faulted.

```
clresource status
```

**Note** – If you import a corrupt ZFS pool, the best option is to choose Continue to display an error message. Other choices are Wait (which hangs until success occurs or the node panics) or Panic (which panics the node). Choosing Continue can help increase the cluster's availability.

Example 2–46 Status of a Faulty HAStoragePlus Resource

This example shows the status of a faulty HAStoragePlus resource. This resource is faulty because the ZFS pool *zpool* failed to import.

```
clresource status hasp-resource
```

=== Cluster Resources ===

| Resource Name | Node Name | Status  | Status Message                     |
|---------------|-----------|---------|------------------------------------|
| hasp-resource | node46    | Online  | Faulted - Failed to import:hazpool |
|               | node47    | Offline | Offline                            |

# Changing the Cluster File System to a Local File System in an HAStoragePlus Resource

You can change the file system of an HAStoragePlus resource from a cluster file system to a local file system.

## ▼ How to Change the Cluster File System to Local File System in an HAStoragePlus Resource

1 Bring the failover resource group offline.

```
clresourcegroup offline resource-group
```

**2 Display the HASStoragePlus resource.**

```
clresource show -g resource-group -t SUNW.HASStoragePlus
```

**3 Retrieve the list of mount points for each resource.**

```
clresource show -p FilesystemMountPoints hastorageplus-resource
```

**4 Unmount the cluster file system.**

```
umount mount-points
```

**5 Modify the /etc/vfstab entry of the mount points on all the nodes configured in the node list of the resource group.**

- Remove the `global` keyword from the mount options.
- Modify the `mount at boot` option from `yes` to `no`.

Repeat the steps for all the cluster file systems of all the HASStoragePlus resources configured in the resource group.

**6 Bring the resource group online.**

```
clresourcegroup online -M resource-group
```

## Upgrading the HASStoragePlus Resource Type

The HASStoragePlus resource type enables you to modify highly available file systems online. Upgrade the HASStoragePlus resource type if all conditions in the following list apply:

- You are upgrading from an earlier version of Oracle Solaris Cluster.
- You need to use the new features of the HASStoragePlus resource type.

For general instructions that explain how to upgrade a resource type, see [“Upgrading a Resource Type” on page 42](#). The information that you need to complete the upgrade of the HASStoragePlus resource type is provided in the subsections that follow.

## Information for Registering the New Resource Type Version

To determine the version of the resource type that is registered, use a command from the following list:

- The `cluster show` command displays the name and version of the cluster's resource types.
- The `clresourcetype list -v` command displays the node list of each resource type.

The RTR file for this resource type is `/usr/cluster/lib/rgm/rtrereg/SUNW.HAStoragePlus`.

## Information for Migrating Existing Instances of the Resource Type

Use the following information to migrate instances of the `HAStoragePlus` resource type:

- You can perform the migration to version 9 when the resource fault monitor is disabled. Enable the fault monitor after the migration is complete.
- You can perform the migration to any version other than 9 at any time.
- Set the `Type_version` property to the version of the features you intend to use.
- Ensure that the application has the `resource_dependencies_offline_restart` property set to the `HAStoragePlus` resource that it depends on. This is required for version 9 of the `HAStoragePlus` resource type, but it is optional for other versions of the resource type. The application must have either the `resource_dependencies` property or the `resource_dependencies_offline_restart` property set to the `HAStoragePlus` resource that it depends on if the `HAStoragePlus` version is older than 9.

## Distributing Online Resource Groups Among Cluster Nodes

For maximum availability or optimum performance, some combinations of services require a specific distribution of online resource groups among cluster nodes. Distributing online resource groups involves creating affinities between resource groups for the following purposes:

- Enforcing the required distribution when the resource groups are first brought online
- Preserving the required distribution after an attempt to fail over or switch over a resource group

This section provides the following examples of how to use resource group affinities to distribute online resource groups among cluster nodes:

- Enforcing collocation of a resource group with another resource group
- Specifying a preferred collocation of a resource group with another resource group
- Balancing the load of a set of resource groups
- Specifying that a critical service has precedence
- Delegating the failover or switchover of a resource group
- Combining affinities between resource groups to specify more complex behavior

# Resource Group Affinities

An affinity between resource groups restricts on which nodes the resource groups may be brought online simultaneously. In each affinity, a source resource group declares an affinity for a target resource group or several target resource groups. To create an affinity between resource groups, set the `RG_affinities` resource group property of the source as follows:

**-p `RG_affinities=affinity-list`**

*affinity-list* Specifies a comma-separated list of affinities between the source resource group and a target resource group or several target resource groups. You may specify a single affinity or more than one affinity in the list.

Specify each affinity in the list as follows:

*operator target-rg*

**Note** – Do not include a space between *operator* and *target-rg*.

*operator* Specifies the type of affinity that you are creating. For more information, see [Table 2–3](#).

*target-rg* Specifies the resource group that is the target of the affinity that you are creating.

**TABLE 2–3** Types of Affinities Between Resource Groups

| Operator | Affinity Type                            | Effect                                                                                                                                                                                                                    |
|----------|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| +        | Weak positive                            | If possible, the source is brought online on a node or on nodes where the target is online or starting. However, the source and the target are allowed to be online on different nodes.                                   |
| ++       | Strong positive                          | The source is brought online only on a node or on nodes where the target is online or starting. The source and the target are <i>not</i> allowed to be online on different nodes.                                         |
| -        | Weak negative                            | If possible, the source is brought online on a node or on nodes where the target is <i>not</i> online or starting. However, the source and the target are allowed to be online on the same node.                          |
| --       | Strong negative                          | The source is brought online only on a node or on nodes where the target is not online. The source and the target are <i>not</i> allowed to be online on the same node.                                                   |
| +++      | Strong positive with failover delegation | Same as strong positive, except that an attempt by the source to fail over is delegated to the target. For more information, see <a href="#">“Delegating the Failover or Switchover of a Resource Group”</a> on page 157. |



Weak affinities take precedence over `NodeList` preference ordering.

The current state of other resource groups might prevent a strong affinity from being satisfied on any node. In this situation, the resource group that is the source of the affinity remains offline. If other resource groups' states change to enable the strong affinities to be satisfied, the resource group that is the source of the affinity comes back online.

---

**Note** – Use caution when declaring a strong affinity on a source resource group for more than one target resource group. If all declared strong affinities cannot be satisfied, the source resource group remains offline.

---

## Enforcing Collocation of a Resource Group With Another Resource Group

A service that is represented by one resource group might depend so strongly on a service in a second resource group that both services must run on the same node. For example, an application that is comprised of multiple interdependent service daemons might require that all daemons run on the same node.

In this situation, force the resource group of the dependent service to be collocated with the resource group of the other service. To enforce collocation of a resource group with another resource group, declare on the resource group a strong positive affinity for the other resource group.

```
clresourcegroup set|create -p RG_affinities=++target-rg source-rg
```

*source-rg*

Specifies the resource group that is the source of the strong positive affinity. This resource group is the resource group *on* which you are declaring a strong positive affinity for another resource group.

```
-p RG_affinities=++target-rg
```

Specifies the resource group that is the target of the strong positive affinity. This resource group is the resource group *for* which you are declaring a strong positive affinity.

A resource group follows the resource group for which it has a strong positive affinity. If the target resource group is relocated to a different node, the source resource group automatically switches to the same node as the target. However, a resource group that declares a strong positive affinity is prevented from failing over to a node on which the target of the affinity is not already running.

---

**Note** – Only failovers that are initiated by a resource monitor are prevented. If a node on which the source resource group and target resource group are running fails, both resource groups fail over to the same surviving node.

---

For example, a resource group rg1 declares a strong positive affinity for resource group rg2. If rg2 fails over to another node, rg1 also fails over to that node. This failover occurs even if all the resources in rg1 are operational. However, if a resource in rg1 attempts to fail over rg1 to a node where rg2 is not running, this attempt is blocked.

The source of a strong positive affinity might be offline on all nodes when you bring online the target of the strong positive affinity. In this situation, the source of the strong positive affinity is automatically brought online on the same node as the target.

For example, a resource group rg1 declares a strong positive affinity for resource group rg2. Both resource groups are initially offline on all nodes. If an administrator brings online rg2 on a node, rg1 is automatically brought online on the same node.

You can use the `clresourcegroup suspend` command to prevent a resource group from being brought online automatically due to strong affinities or cluster reconfiguration.

If you require a resource group that declares a strong positive affinity to be allowed to fail over, you must delegate the failover. For more information, see [“Delegating the Failover or Switchover of a Resource Group” on page 157](#).

#### EXAMPLE 2-47 Enforcing Collocation of a Resource Group With Another Resource Group

This example shows the command for modifying resource group rg1 to declare a strong positive affinity for resource group rg2. As a result of this affinity relationship, rg1 is brought online only on nodes where rg2 is running. This example assumes that both resource groups exist.

```
clresourcegroup set -p RG_affinities=++rg2 rg1
```

## Specifying a Preferred Collocation of a Resource Group With Another Resource Group

A service that is represented by one resource group might use a service in a second resource group. As a result, these services run most efficiently if they run on the same node. For example, an application that uses a database runs most efficiently if the application and the database run on the same node. However, the services can run on different nodes because the reduction in efficiency is less disruptive than additional failovers of resource groups.

In this situation, specify that both resource groups should be colocated if possible. To specify preferred collocation of a resource group with another resource group, declare on the resource group a weak positive affinity for the other resource group.

```
clresourcegroup set|create -p RG_affinities=+target-rg source-rg
```

*source-rg*

Specifies the resource group that is the source of the weak positive affinity. This resource group is the resource group *on* which you are declaring a weak positive affinity for another resource group.

```
-p RG_affinities=+target-rg
```

Specifies the resource group that is the target of the weak positive affinity. This resource group is the resource group *for* which you are declaring a weak positive affinity.

By declaring a weak positive affinity on one resource group for another resource group, you increase the probability of both resource groups running on the same node. The source of a weak positive affinity is first brought online on a node where the target of the weak positive affinity is already running. However, the source of a weak positive affinity does not fail over if a resource monitor causes the target of the affinity to fail over. Similarly, the source of a weak positive affinity does not fail over if the target of the affinity is switched over. In both situations, the source remains online on the node where the source is already running.

---

**Note** – If a node on which the source resource group and target resource group are running fails, both resource groups are restarted on the same surviving node.

---

#### EXAMPLE 2-48 Specifying a Preferred Collocation of a Resource Group With Another Resource Group

This example shows the command for modifying resource group `rg1` to declare a weak positive affinity for resource group `rg2`. As a result of this affinity relationship, `rg1` and `rg2` are first brought online on the same node. But if a resource in `rg2` causes `rg2` to fail over, `rg1` remains online on the node where the resource groups were first brought online. This example assumes that both resource groups exist.

```
clresourcegroup set -p RG_affinities=+rg2 rg1
```

## Distributing a Set of Resource Groups Evenly Among Cluster Nodes

Each resource group in a set of resource groups might impose the same load on the cluster. In this situation, by distributing the resource groups evenly among cluster nodes, you can balance the load on the cluster.

To distribute a set of resource groups evenly among cluster nodes, declare on each resource group a weak negative affinity for the other resource groups in the set.

```
clresourcegroup set|create -p RG_affinities=neg-affinity-list source-rg
```

*source-rg*

Specifies the resource group that is the source of the weak negative affinity. This resource group is the resource group *on* which you are declaring a weak negative affinity for other resource groups.

```
-p RG_affinities=neg-affinity-list
```

Specifies a comma-separated list of weak negative affinities between the source resource group and the resource groups that are the target of the weak negative affinity. The target resource groups are the resource groups *for* which you are declaring a weak negative affinity.

By declaring a weak negative affinity on one resource group for other resource groups, you ensure that a resource group is always brought online on the most lightly loaded node in the cluster. The fewest other resource groups are running on that node. Therefore, the smallest number of weak negative affinities are violated.

#### EXAMPLE 2-49 Distributing a Set of Resource Groups Evenly Among Cluster Nodes

This example shows the commands for modifying resource groups rg1, rg2, rg3, and rg4 to ensure that these resource groups are evenly distributed among the available nodes in the cluster. This example assumes that resource groups rg1, rg2, rg3, and rg4 exist.

```
clresourcegroup set -p RG_affinities=-rg2,-rg3,-rg4 rg1
clresourcegroup set -p RG_affinities=-rg1,-rg3,-rg4 rg2
clresourcegroup set -p RG_affinities=-rg1,-rg2,-rg4 rg3
clresourcegroup set -p RG_affinities=-rg1,-rg2,-rg3 rg4
```

## Specifying That a Critical Service Has Precedence

A cluster might be configured to run a combination of mission-critical services and noncritical services. For example, a database that supports a critical customer service might run in the same cluster as noncritical research tasks.

To ensure that the noncritical services do not affect the performance of the critical service, specify that the critical service has precedence. By specifying that the critical service has precedence, you prevent noncritical services from running on the same node as the critical service.

When all nodes are operational, the critical service runs on a different node from the noncritical services. However, a failure of the critical service might cause the service to fail over to a node where the noncritical services are running. In this situation, the noncritical services are taken offline immediately to ensure that the computing resources of the node are fully dedicated to the mission-critical service.

To specify that a critical service has precedence, declare on the resource group of each noncritical service a strong negative affinity for the resource group that contains the critical service.

```
clresourcegroup set|create -p RG_affinities=--critical-rg noncritical-rg
noncritical-rg
```

Specifies the resource group that contains a noncritical service. This resource group is the resource group *on* which you are declaring a strong negative affinity for another resource group.

```
-p RG_affinities=--critical-rg
```

Specifies the resource group that contains the critical service. This resource group is the resource group *for* which you are declaring a strong negative affinity.

A resource group moves away from a resource group for which it has a strong negative affinity.

The source of a strong negative affinity might be offline on all nodes when you take offline the target of the strong negative affinity. In this situation, the source of the strong negative affinity is automatically brought online. In general, the resource group is brought online on the most preferred node, based on the order of the nodes in the node list and the declared affinities.

For example, a resource group `rg1` declares a strong negative affinity for resource group `rg2`. Resource group `rg1` is initially offline on all nodes, while resource group `rg2` is online on a node. If an administrator takes offline `rg2`, `rg1` is automatically brought online.

You can use the `clresourcegroup suspend` command to prevent the source of a strong negative affinity from being brought online automatically due to strong affinities or cluster reconfiguration.

#### EXAMPLE 2-50 Specifying That a Critical Service Has Precedence

This example shows the commands for modifying the noncritical resource groups `ncrg1` and `ncrg2` to ensure that the critical resource group `mcdbrg` has precedence over these resource groups. This example assumes that resource groups `mcdbrg`, `ncrg1`, and `ncrg2` exist.

```
clresourcegroup set -p RG_affinities=--mcdbrg ncrg1 ncrg2
```

## Delegating the Failover or Switchover of a Resource Group

The source resource group of a strong positive affinity cannot fail over or be switched over to a node where the target of the affinity is not running. If you require the source resource group of a strong positive affinity to be allowed to fail over or be switched over, you must delegate the failover to the target resource group. When the target of the affinity fails over, the source of the affinity is forced to fail over with the target.

---

**Note** – You might need to switch over the source resource group of a strong positive affinity that is specified by the ++ operator. In this situation, switch over the target of the affinity and the source of the affinity at the same time.

---

To delegate failover or switchover of a resource group to another resource group, declare on the resource group a strong positive affinity with failover delegation for the other resource group.

```
clresourcegroup set|create source-rg -p RG_affinities=+++target-rg
```

*source-rg*

Specifies the resource group that is delegating failover or switchover. This resource group is the resource group *on* which you are declaring a strong positive affinity with failover delegation for another resource group.

```
-p RG_affinities=+++target-rg
```

Specifies the resource group to which *source-rg* delegates failover or switchover. This resource group is the resource group *for* which you are declaring a strong positive affinity with failover delegation.

A resource group may declare a strong positive affinity with failover delegation for at most one resource group. However, a given resource group may be the target of strong positive affinities with failover delegation that are declared by any number of other resource groups.

A strong positive affinity with failover delegation is not fully symmetric. The target can come online while the source remains offline. However, if the target is offline, the source cannot come online.

If the target declares a strong positive affinity with failover delegation for a third resource group, failover or switchover is further delegated to the third resource group. The third resource group performs the failover or switchover, forcing the other resource groups to fail over or be switched over also.

#### EXAMPLE 2-51 Delegating the Failover or Switchover of a Resource Group

This example shows the command for modifying resource group `rg1` to declare a strong positive affinity with failover delegation for resource group `rg2`. As a result of this affinity relationship, `rg1` delegates failover or switchover to `rg2`. This example assumes that both resource groups exist.

```
clresourcegroup set -p RG_affinities=+++rg2 rg1
```

## Combining Affinities Between Resource Groups

You can create more complex behaviors by combining multiple affinities. For example, the state of an application might be recorded by a related replica server. The node selection requirements for this example are as follows:

- The replica server must run on a different node from the application.
- If the application fails over from its current node, the application should fail over to the node where the replica server is running.
- If the application fails over to the node where the replica server is running, the replica server must fail over to a different node. If no other node is available, the replica server must go offline.

You can satisfy these requirements by configuring resource groups for the application and the replica server as follows:

- The resource group that contains the application declares a weak positive affinity for the resource group that contains the replica server.
- The resource group that contains the replica server declares a strong negative affinity for the resource group that contains the application.

### EXAMPLE 2-52 Combining Affinities Between Resource Groups

This example shows the commands for combining affinities between the following resource groups.

- Resource group `app-rg` represents an application whose state is tracked by a replica server.
- Resource group `rep-rg` represents the replica server.

In this example, the resource groups declare affinities as follows:

- Resource group `app-rg` declares a weak positive affinity for resource group `rep-rg`.
- Resource group `rep-rg` declares a strong negative affinity for resource group `app-rg`.

This example assumes that both resource groups exist.

```
clresourcegroup set -p RG_affinities=+rep-rg app-rg
clresourcegroup set -p RG_affinities=--app-rg rep-rg
```

## Zone Cluster Resource Group Affinities

The cluster administrator can specify affinities between a resource group in a zone cluster and another resource group in a zone cluster or a resource group on the global cluster.

You can use the following command to specify the affinity between resource groups in zone clusters.

```
clresourcegroup set -p RG_affinities=affinity-typetarget-zc:target-rg source-zc:source-rg
```

The resource group affinity types in a zone cluster can be one of the following:

- + (weak positive)
- ++ (strong positive)
- +++ (strong positive with failover delegation)
- - (weak negative)
- -- (strong negative)

**EXAMPLE 2-53** Specifying a Strong Positive Affinity Between Resource Groups in Zone Clusters

This example shows the command for specifying a strong positive affinity between resource groups in zone clusters.

The resource group RG1 in a zone cluster ZC1 declares a strong positive affinity for a resource group RG2 in a zone cluster ZC2.

If you need to specify a strong positive affinity between a resource group RG1 in a zone cluster ZC1 and a resource group RG2 in another zone cluster ZC2, use the following command:

```
clresourcegroup set -p RG_affinities=++ZC2:RG2 ZC1:RG1
```

**EXAMPLE 2-54** Specifying a Strong Negative Affinity Between a Resource Group in a Zone Cluster and a Resource Group in the Global Cluster

This example shows the command for specifying a strong negative affinity between resource groups in zone clusters. If you need to specify a strong negative affinity between a resource group RG1 in a zone cluster ZC1 and a resource group RG2 in the global cluster, use the following command:

```
clresourcegroup set -p RG_affinities=--global:RG2 ZC1:RG1
```

## Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources

If you require identical resource configuration data on two clusters, you can replicate the data to the second cluster to save the laborious task of setting it up again. Use `scsnapshot` to propagate the resource configuration information from one cluster to another cluster. To save effort, ensure that your resource-related configuration is stable and you do not need to make any major changes to the resource configuration, before copying the information to a second cluster.

Configuration data for resource groups, resource types, and resources can be retrieved from the Cluster Configuration Repository (CCR) and formatted as a shell script. The script can be used to perform the following tasks:



- Replicate configuration data on a cluster that does not have configured resource groups, resource types, or resources
- Upgrade configuration data on a cluster that has configured resource groups, resource types, and resources

The `scsnapshot` tool retrieves configuration data that is stored in the CCR. Other configuration data are ignored. The `scsnapshot` tool ignores the dynamic state of different resource groups, resource types, and resources.

## ▼ How to Replicate Configuration Data on a Cluster Without Configured Resource Groups, Resource Types, and Resources

This procedure replicates configuration data on a cluster that does not have configured resource groups, resource types, and resources. In this procedure, a copy of the configuration data is taken from one cluster and used to generate the configuration data on another cluster.

- 1 Using the system administrator role, log in to any node in the cluster from which you want to copy the configuration data.**

For example, `node1`.

The system administrator role gives you the following role-based access control (RBAC) rights:

- `solaris.cluster.resource.read`
- `solaris.cluster.resource.modify`

- 2 Retrieve the configuration data from the cluster.**

```
node1 % scsnapshot -s scriptfile
```

The `scsnapshot` tool generates a script called *scriptfile*. For more information about using the `scsnapshot` tool, see the [scsnapshot\(1M\)](#) man page.

- 3 Edit the script to adapt it to the specific features of the cluster where you want to replicate the configuration data.**

For example, you might have to change the IP addresses and host names that are listed in the script.

- 4 Launch the script from any node in the cluster where you want to replicate the configuration data.**

The script compares the characteristics of the local cluster to the cluster where the script was generated. If the characteristics are not the same, the script writes an error and ends. A message

asks whether you want to rerun the script, using the `-f` option. The `-f` option forces the script to run, despite any difference in characteristics. If you use the `-f` option, ensure that you do not create inconsistencies in your cluster.

The script verifies that the Oracle Solaris Cluster resource type exists on the local cluster. If the resource type does not exist on the local cluster, the script writes an error and ends. A message asks whether you want to install the missing resource type before running the script again.

## ▼ How to Upgrade Configuration Data on a Cluster With Configured Resource Groups, Resource Types, and Resources

This procedure upgrades configuration data on a cluster that already has configured resource groups, resource types, and resources. This procedure can also be used to generate a configuration template for resource groups, resource types, and resources.

In this procedure, the configuration data on `cluster1` is upgraded to match the configuration data on `cluster2`.

### 1 Using the system administrator role, log on to any node in `cluster1`.

For example, `node1`.

The system administrator role gives you the following RBAC rights:

- `solaris.cluster.resource.read`
- `solaris.cluster.resource.modify`

### 2 Retrieve the configuration data from the cluster by using the image file option of the `scsnapshot` tool:

```
node1% scsnapshot -s scriptfile1 -o imagefile1
```

When run on `node1`, the `scsnapshot` tool generates a script that is called `scriptfile1`. The script stores configuration data for the resource groups, resource types, and resources in an image file that is called `imagefile1`. For more information about using the `scsnapshot` tool, see the [scsnapshot\(1M\)](#) man page.

### 3 Repeat [Step 1](#) through [Step 2](#) on a node in `cluster2`:

```
node2 % scsnapshot -s scriptfile2 -o imagefile2
```

### 4 On `node1`, generate a script to upgrade the configuration data on `cluster1` with configuration data from `cluster2`:

```
node1 % scsnapshot -s scriptfile3 imagefile1 imagefile2
```

This step uses the image files that you generated in [Step 2](#) and [Step 3](#), and generates a new script that is called `scriptfile3`.

- 5 **Edit the script that you generated in [Step 4](#) to adapt it to the specific features of the `cluster1`, and to remove data specific to `cluster2`.**

- 6 **From `node1`, launch the script to upgrade the configuration data.**

The script compares the characteristics of the local cluster to the cluster where the script was generated. If the characteristics are not the same, the script writes an error and ends. A message asks whether you want to rerun the script, using the `-f` option. The `-f` option forces the script to run, despite any difference in characteristics. If you use the `-f` option, ensure that you do not create inconsistencies in your cluster.

The script verifies that the Oracle Solaris Cluster resource type exists on the local cluster. If the resource type does not exist on the local cluster, the script writes an error and ends. A message asks whether you want to install the missing resource type before running the script again.

## Enabling Oracle Solaris SMF Services to Run With Oracle Solaris Cluster

The Service Management Facility (SMF) enables you to automatically start and restart SMF services, during a node boot or service failure. SMF facilitates some degree of high availability to the SMF services on a single host. This feature is similar to the Oracle Solaris Cluster Resource Group Manager (RGM), which facilitates high availability and scalability for cluster applications. SMF services and RGM features are complementary to each other.

Oracle Solaris Cluster includes three SMF proxy resource types that can be used to enable SMF services to run with Oracle Solaris Cluster in a failover, multi-master, or scalable configuration. The following are the proxy resource types:

- `SUNW.Proxy_SMF_failover`
- `SUNW.Proxy_SMF_multimaster`
- `SUNW.Proxy_SMF_scalable`

The SMF proxy resource types enables you to encapsulate a set of interrelated SMF services into a single resource, *SMF proxy resource* to be managed by Oracle Solaris Cluster. In this feature, SMF manages the availability of SMF services on a single node. Oracle Solaris Cluster provides cluster-wide high availability and scalability of the SMF services.

You can use the SMF proxy resource types to integrate your own SMF controlled services into Oracle Solaris Cluster so that these services have cluster-wide service availability without you rewriting callback methods or service manifest. After you integrate the SMF service into the SMF proxy resource, the SMF service is no longer managed by the default restarter. The restarter that is delegated by Oracle Solaris Cluster manages the SMF service.

SMF proxy resources are identical to other resources, with no restriction on their usage. For example, an SMF proxy resource can be grouped with other resources into a resource group. SMF proxy resources can be created and managed the same way as other resources. An SMF

proxy resource differs from other resources in one way. When you create a resource of any of the SMF proxy resource types, you need to specify the extension property `Proxied_service_instances`. You must include information about the SMF services to be proxied by the SMF resource. The extension property's value is the path to a file that contains all the proxied SMF services. Each line in the file is dedicated to one SMF service and specifies `svc_fmri`, path of the corresponding service manifest file.

For example, if the resource has to manage two services, `restarter_svc_test_1:default` and `restarter_svc_test_2:default`, the file should include the following two lines:

```
<svc:/system/cluster/restarter_svc_test_1:default>,</var/svc/manifest/system/cluster/restarter_svc_test_1.xml>

<svc:/system/cluster/restarter_svc_test_2:default>,</var/svc/manifest/system/cluster/restarter_svc_test_2.xml>
```

The services that are encapsulated under an SMF proxy resource can reside in the global cluster or global-cluster non-voting node. However, all the services under the same proxy resource must be in the same zone.



**Caution** – Do not use SMF `svcadm` for disabling or enabling SMF services that are encapsulated in a proxy resource. Do not change the properties of the SMF services (in the SMF repository) that are encapsulated in a proxy resource.

---

- “Encapsulating an SMF Service Into a Failover Proxy Resource Configuration” on page 164
  - “Encapsulating an SMF Service Into a Multi-Master Proxy Resource Configuration” on page 167
  - “Encapsulating an SMF Service Into a Scalable Proxy Resource Configuration” on page 170
- 

## ▼ Encapsulating an SMF Service Into a Failover Proxy Resource Configuration

For information about failover configuration, see “[Creating a Resource Group](#)” on page 49

---

**Note** – Perform this procedure from any cluster node.

---

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Register the proxy SMF failover resource type.

```
clresourcetype register -f \
/opt/SUNWscsmf/etc/SUNW.Proxy_SMF_failover SUNW.Proxy_SMF_failover
```

**3 Verify that the proxy resource type has been registered.**

```
clresource_type show
```

**4 Create the SMF failover resource group.**

```
clresourcegroup create [-n node-zone-list] resource-group
```

*-n node-zone-list* Specifies a comma-separated, ordered list of nodes that can master this resource group. The format of each entry in the list is *node:zone*. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource group is configured on all the global-cluster voting nodes.

---

**Note** – To achieve highest availability, specify global-cluster non-voting nodes on different nodes in the node list of an SMF failover resource group instead of different global-cluster non-voting nodes on the same node.

---

*resource-group* Specifies your choice of the name of the resource group to add. This name must begin with an ASCII character.

**5 Verify that the SMF resource group has been created.**

```
clresourcegroup status resource-group
```

**6 Add an SMF failover application resource to the resource group.**

```
clresource create -g resource-group -t SUNW.Proxy_SMF_failover \
-p Port_list=portnumber/protocol \
-x Proxied_service_instances=/tmp/dns_svcs.txt
```

*-g resource-group* Specifies the name of the SMF failover resource group that you previously created.

*-p Port\_list=portnumber/protocol* Specifies the port number the instance will use to listen for activity. The protocol can be either *tcp* or *udp*.

*-p Proxied\_service\_instances* Specifies the path to the file you created that specifies the mapping of SMF services and their corresponding manifests for the SMF services to be proxied. In the example above, */tmp/dns\_svcs.txt* is the path to the text file.

The resource is created in the enabled state.

**7 Verify that the SMF failover application resource has been added and validated.**

```
clresource show resource
```

**8 Bring the failover resource group online.**

```
clresourcegroup online -M resource-group
```

---

**Note** – If you use the `clresource status` command to view the state of the SMF proxy resource type, the status is displayed as `online` but not `monitored`. This is not an error message. The SMF proxy resource is enabled and running and this status message is displayed because there is no monitoring support provided for the resources of SMF proxy resource type.

---

**Example 2–55 Registering an SMF Proxy Failover Resource Type**

The following example registers the `SUNW.Proxy_SMF_failover` resource type.

```
clresourcetype register SUNW.Proxy_SMF_failover
clresourcetype show SUNW.Proxy_SMF_failover

Resource Type: SUNW.Proxy_SMF_failover
RT_description: Resource type for proxying failover SMF services
RT_version: 3.2
API_version: 6
RT_basedir: /opt/SUNWscsmf/bin
Single_instance: False
Proxy: False
Init_nodes: All potential masters
Installed_nodes: <All>
Failover: True
Pkglist: SUNWscsmf
RT_system: False
Global_zone: False
```

**Example 2–56 Adding an SMF Proxy Failover Application Resource to a Resource Group**

This example shows the addition of a proxy resource type, `SUNW.Proxy_SMF_failover` to a resource group `resource-group-1`.

```
clresource create -g resource-group-1 -t SUNW.Proxy_SMF_failover \
-x proxied_service_instances=/var/tmp/svslst.txt resource-1
clresource show resource-1

=== Resources ===

Resource: resource-1
Type: SUNW.Proxy_SMF_failover
Type_version: 3.2
Group: resource-group-1
R_description:
Resource_project_name: default
```

|                    |      |
|--------------------|------|
| Enabled{phats1}:   | True |
| Monitored{phats1}: | True |

## ▼ Encapsulating an SMF Service Into a Multi-Master Proxy Resource Configuration

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.

- 2 Register the SMF proxy multi-master resource type.

```
clresourcetype register -f \
/opt/SUNWscsmf/etc/SUNW.Proxy_SMF_multimaster SUNW.Proxy_SMF_multimaster
```

- 3 Create the SMF multi-master resource group.

```
clresourcegroup create -S [-p Maximum primaries=m] [-p Desired primaries=n] \
[-n node-zone-list] resource-group
```

-S Specifies that the resource group is to be multi-mastered. If the -p Maximum primaries and -p Desired primaries options are omitted, both properties are set to the number of nodes in the resource group's node list.

-p Maximum primaries=*m* Specifies the maximum number of active primaries for this resource group.

-p Desired primaries=*n* Specifies the number of active primaries on which the resource group should attempt to start.

-n *node-zone-list* Specifies a comma-separated, ordered list of nodes in which this resource group is to be available. The format of each entry in the list is *node:zone*. In this format, *node* specifies the node name and *zone* specifies the name of a global-cluster non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource group is configured on the global-cluster voting nodes.

*resource-group* Specifies your choice of the name of the scalable resource group to add. This name must begin with an ASCII character.

- 4 Verify that the SMF proxy multi-master resource group has been created.

```
clresourcegroup show resource-group
```

**5 Add an SMF proxy multi-master resource to the resource group.**

```
clresource create -g resource-group -t SUNW.Proxy_SMF_multimaster \
-p Port_list=portnumber/protocol \
-x Proxied_service_instances=/tmp/dns_svcs.txt
```

|                                  |                                                                                                                                                                                                                                   |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -g resource-group                | Specifies the name of the SMF multi-master resource group that you previously created.                                                                                                                                            |
| -p Port_list=portnumber/protocol | Specifies the port number the instance will use to listen for activity. The protocol can be either tcp or udp.                                                                                                                    |
| -p Proxied_service_instances     | Specifies the path to the file you created that specifies the mapping of SMF services and their corresponding manifests for the SMF services to be proxied. In the example above, /tmp/dns_svcs.txt is the path to the text file. |

The resource is created in the enabled state.

**6 Verify that the SMF proxy multi-master application resource has been added and validated.**

```
clresource show resource
```

**7 Bring the multi-master resource group online.**

```
clresourcegroup online -M resource-group
```

---

**Note** – If you use the `clresource status` command to view the state of the SMF proxy resource type, the status is displayed as `online` but not `monitored`. This is not an error message. The SMF proxy resource is enabled and running and this status message is displayed because there is no monitoring support provided for the resources of SMF proxy resource type.

---

**Example 2–57 Registering an SMF Proxy Multi-Master Resource Type**

The following example registers the `SUNW.Proxy_SMF_multimaster` resource type.

```
clresourcetype register SUNW.Proxy_SMF_multimaster
clresourcetype show SUNW.Proxy_SMF_multimaster

Resource Type: SUNW.Proxy_SMF_multimaster
RT_description: Resource type for proxying multimastered SMF services
RT_version: 3.2
API_version: 6
RT_basedir: /opt/SUNWscsmf/bin
Single_instance: False
Proxy: False
Init_nodes: All potential masters
Installed_nodes: <All>
Failover: True
Pkglist: SUNWscsmf
```



```
RT_system: False
Global_zone: False
```

### Example 2–58 Creating and Adding an SMF Proxy Multi-Master Application Resource to a Resource Group

This example shows the creation and addition of a multi-master proxy resource type `SUNW.Proxy_SMF_multimaster` to a resource group `resource-group-1`.

```
clresourcegroup create -S \
-p Maximum_primaries=2 \
-p Desired_primaries=2 \
-n phys-schost-1, phys-schost-2 resource-group-1
clresourcegroup show resource-group-1

=== Resource Groups and Resources ===

Resource Group: resource-group-1
RG_description: <NULL>
RG_mode: multimastered
RG_state: Unmanaged
RG_project_name: default
RG_affinities: <NULL>
Auto_start_on_new_cluster: True
Failback: False
Nodelist: phys-schost-1 phys-schost-2
Maximum_primaries: 2
Desired_primaries: 2
Implicit_network_dependencies: True
Global_resources_used: <All>
Pingpong_interval: 3600
Pathprefix: <NULL>
RG_System: False
Suspend_automatic_recovery: False

clresource create -g resource-group-1 -t SUNW.Proxy_SMF_multimaster \
-x proxied_service_instances=/var/tmp/svslist.txt resource-1
clresource show resource-1

=== Resources ===

Resource: resource-1
Type: SUNW.Proxy_SMF_multimaster
Type_version: 3.2
Group: resource-group-1
R_description:
Resource_project_name: default
Enabled{phats1}: True
Monitored{phats1}: True
```

# ▼ Encapsulating an SMF Service Into a Scalable Proxy Resource Configuration

For information about scalable configuration, see [“How to Create a Scalable Resource Group” on page 51](#).

---

**Note** – Perform this procedure from any cluster node.

---

- 1 On a cluster member, become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Register the SMF proxy scalable resource type.  

```
clresourcetype register -f \
/opt/SUNWscsmf/etc/SUNW.Proxy_SMF_scalable SUNW.Proxy_SMF_scalable
```
- 3 Create the SMF failover resource group that holds the shared address that the scalable resource group is to use. See [“How to Create a Failover Resource Group” on page 49](#) to create the failover resource group.
- 4 Add the shared address resource to the failover resource group.  
See [“How to Add a Shared Address Resource to a Resource Group Using the Command-Line Interface” on page 63](#). The example below uses a resource called *network-resource*.
- 5 Create the SMF proxy scalable resource group.  

```
clresourcegroup create -S [-p Maximum primaries=m] [-p Desired primaries=n] \
[-n node-zone-list] resource-group
```

|                        |                                                                                                                                                                                                                                                                                            |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -S                     | Specifies that the resource group is to be multi-mastered. If the <code>-p Maximum primaries</code> and <code>-p Desired primaries</code> options are omitted, both properties are set to the number of nodes in the resource group's node list.                                           |
| -p Maximum primaries=m | Specifies the maximum number of active primaries for this resource group.                                                                                                                                                                                                                  |
| -p Desired primaries=n | Specifies the number of active primaries on which the resource group should attempt to start.                                                                                                                                                                                              |
| -n node-zone-list      | Specifies a comma-separated, ordered list of nodes in which this resource group is to be available. The format of each entry in the list is <i>node:zone</i> . In this format, <i>node</i> specifies the global-cluster voting node and <i>zone</i> specifies the name of a global-cluster |

non-voting node. To specify the global-cluster voting node, or to specify a node without global-cluster non-voting nodes, specify only *node*.

This list is optional. If you omit this list, the resource group is created on all nodes in the cluster.

The node list of the scalable resource can contain the same list or a subset of *nodename:zonename* pairs as the node list of the shared address resource.

*resource-group*

Specifies your choice of the name of the scalable resource group to add. This name must begin with an ASCII character.

## 6 Verify that the scalable resource group has been created.

```
clresourcegroup show resource-group
```

## 7 Add an SMF proxy scalable resource to the scalable resource group you created in [Step 5](#).

```
clresource create-g resource-group -t SUNW.Proxy_SMF_scalable \
-p Resource_dependencies=network-resource[,network-resource...] \
-p Scalable=True \
-p Port_list=portnumber/protocol \
-x Proxied_service_instances=/tmp/dns_svcs.txt
```

```
-p Resource_dependencies=network-resource[,network-resource...]
```

Specifies the name of the scalable network resource you created in [Step 3](#) on which this resource depends.

```
-g resource-group
```

Specifies the name of the SMF proxy scalable resource group that you previously created.

```
-p Scalable=True
```

Specifies that this resource uses the network load balancing features of Oracle Solaris Cluster software. For more information, see [“How to Add a Scalable Application Resource to a Resource Group” on page 68](#).

The resource is created in the enabled state.

## 8 Verify that the SMF proxy scalable application resource has been added and validated.

```
clresource show resource
```

## 9 Bring the SMF proxy scalable resource group online.

```
clresourcegroup online -M resource-group
```

---

**Note** – If you use the `clresource status` command to view the state of the SMF proxy resource type, the status is displayed as `onLine` but not `monitored`. This is not an error message. The SMF proxy resource is enabled and running and this status message is displayed because there is no monitoring support provided for the resources of SMF proxy resource type.

---

### Example 2–59 Registering an SMF Proxy Scalable Resource Type

The following example registers the `SUNW.Proxy_SMF_scalable` resource type.

```
clresourcetype register SUNW.Proxy_SMF_scalable
clresourcetype show SUNW.Proxy_SMF_scalable

Resource Type: SUNW.Proxy_SMF_scalable
RT_description: Resource type for proxying scalable SMF services
RT_version: 3.2
API_version: 6
RT_basedir: /opt/SUNWscsmf/bin
Single_instance: False
Proxy: False
Init_nodes: All potential masters
Installed_nodes: <All>
Failover: True
Pkglist: SUNWscsmf
RT_system: False
Global_zone: False
```

### Example 2–60 Creating and Adding an SMF Proxy Scalable Application Resource to a Resource Group

This example shows the creation and addition of a scalable proxy resource type `SUNW.Proxy_SMF_scalable` to a resource group `resource-group-1`.

```
clresourcegroup create -S \
-p Maximum_primaries=2 \
-p Desired_primaries=2 \
-p RG_dependencies=resource-group-2 \
-n phys-schost-1, phys-schost-2 resource-group-1
clresourcegroup show resource-group-1

=== Resource Groups and Resources ===

Resource Group: resource-group-1
RG_description: <NULL>
RG_mode: Scalable
RG_state: Unmanaged
RG_project_name: default
RG_affinities: <NULL>
Auto_start_on_new_cluster: True
Fallback: False
Nodelist: phys-schost-1 phys-schost-2
Maximum_primaries: 2
```

```

Desired_primaries: 2
RG_dependencies: resource-group2
Implicit_network_dependencies: True
Global_resources_used: <All>
Pingpong_interval: 3600
Pathprefix: <NULL>
RG_System: False
Suspend_automatic_recovery: False

clresource create -g resource-group-1 -t SUNW.Proxy_SMF_scalable \
-p resource_dependencies=net-res -p port_list=1080/tcp \
-x proxied_service_instances=/var/tmp/svslst.txt resource-1
clresource show resource-1

=== Resources ===

Resource: resource-1
Type: SUNW.Proxy_SMF_scalable
Type_version: 3.2
Group: resource-group-1
R_description:
Resource_project_name: default
Enabled{phats1}: True
Monitored{phats1}: True

```

You can choose the `resource_dependencies` and port number to use.

## Tuning Fault Monitors for Oracle Solaris Cluster Data Services

Each data service that is supplied with the Oracle Solaris Cluster product has a built-in fault monitor. The fault monitor performs the following functions:

- Detecting the unexpected termination of processes for the data service server
- Checking the health of the data service

The fault monitor is contained in the resource that represents the application for which the data service was written. You create this resource when you register and configure the data service. For more information, see the documentation for the data service.

System properties and extension properties of this resource control the behavior of the fault monitor. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune a fault monitor *only* if you need to modify this preset behavior.

Tuning a fault monitor involves the following tasks:

- Setting the interval between fault monitor probes
- Setting the timeout for fault monitor probes
- Defining the criteria for persistent faults
- Specifying the failover behavior of a resource

Perform these tasks when you register and configure the data service. For more information, see the documentation for the data service.

---

**Note** – A resource's fault monitor is started when you bring online the resource group that contains the resource. You do not need to start the fault monitor explicitly.

---

## Setting the Interval Between Fault Monitor Probes

To determine whether a resource is operating correctly, the fault monitor probes this resource periodically. The interval between fault monitor probes affects the availability of the resource and the performance of your system as follows:

- The interval between fault monitor probes affects the length of time that is required to detect a fault and respond to the fault. Therefore, if you decrease the interval between fault monitor probes, the time that is required to detect a fault and respond to the fault is also decreased. This decrease enhances the availability of the resource.
- Each fault monitor probe consumes system resources such as processor cycles and memory. Therefore, if you decrease the interval between fault monitor probes, the performance of the system is degraded.

The optimum interval between fault monitor probes also depends on the time that is required to respond to a fault in the resource. This time depends on how the complexity of the resource affects the time that is required for operations such as restarting the resource.

To set the interval between fault monitor probes, set the `Thorough_probe_interval` system property of the resource to the interval in seconds that you require.

## Setting the Timeout for Fault Monitor Probes

The timeout for fault monitor probes specifies the length of time that a fault monitor waits for a response from a resource to a probe. If the fault monitor does not receive a response within this timeout, the fault monitor treats the resource as faulty. The time that a resource requires to respond to a fault monitor probe depends on the operations that the fault monitor performs to probe the resource. For information about operations that a data service's fault monitor performs to probe a resource, see the documentation for the data service.

The time that is required for a resource to respond also depends on factors that are unrelated to the fault monitor or the application, for example:

- System configuration
- Cluster configuration
- System load
- Amount of network traffic

To set the timeout for fault monitor probes, set the `Probe_timeout` extension property of the resource to the timeout in seconds that you require.

## Defining the Criteria for Persistent Faults

To minimize the disruption that transient faults in a resource cause, a fault monitor restarts the resource in response to such faults. For persistent faults, more disruptive action than restarting the resource is required:

- For a failover resource, the fault monitor fails over the resource to another node.
- For a scalable resource, the fault monitor takes the resource offline.

A fault monitor treats a fault as persistent if the number of complete failures of a resource exceeds a specified threshold within a specified retry interval. Defining the criteria for persistent faults enables you to set the threshold and the retry interval to accommodate the performance characteristics of your cluster and your availability requirements.

### Complete Failures and Partial Failures of a Resource

A fault monitor treats some faults as a *complete failure* of a resource. A complete failure typically causes a complete loss of service. The following failures are examples of a complete failure:

- Unexpected termination of the process for a data service server
- Inability of a fault monitor to connect to a data service server

A complete failure causes the fault monitor to increase by 1 the count of complete failures in the retry interval.

A fault monitor treats other faults as a *partial failure* of a resource. A partial failure is less serious than a complete failure, and typically causes a degradation of service, but not a complete loss of service. An example of a partial failure is an incomplete response from a data service server before a fault monitor probe is timed out.

A partial failure causes the fault monitor to increase by a fractional amount the count of complete failures in the retry interval. Partial failures are still accumulated over the retry interval.

The following characteristics of partial failures depend on the data service:

- The types of faults that the fault monitor treats as partial failure
- The fractional amount that each partial failure adds to the count of complete failures

For information about faults that a data service's fault monitor detects, see the documentation for the data service.

## Dependencies of the Threshold and the Retry Interval on Other Properties

The maximum length of time that is required for a single restart of a faulty resource is the sum of the values of the following properties:

- `Thorough_probe_interval` system property
- `Probe_timeout` extension property

To ensure that you allow enough time for the threshold to be reached within the retry interval, use the following expression to calculate values for the retry interval and the threshold:

$$retry\_interval \geq 2 \times threshold \times (thorough\_probe\_interval + probe\_timeout)$$

The factor of 2 accounts for partial probe failures that do not immediately cause the resource to be failed over or taken offline.

## System Properties for Setting the Threshold and the Retry Interval

To set the threshold and the retry interval, set the following system properties of the resource:

- To set the threshold, set the `Retry_count` system property to the maximum allowed number of complete failures.
- To set the retry interval, set the `Retry_interval` system property to the interval in seconds that you require.

## Specifying the Failover Behavior of a Resource

The failover behavior of a resource determines how the RGM responds to the following faults:

- Failure of the resource to start
- Failure of the resource to stop
- Failure of the resource's fault monitor to stop

To specify the failover behavior of a resource, set the `Failover_mode` system property of the resource. For information about the possible values of this property, see the description of the `Failover_mode` system property in the [r\\_properties\(5\)](#) man page.

## Denying Cluster Services For a Selected Non-Global Zone

You can choose to set up non-global zones that do not participate in the cluster. A root user logged into one of these zones will not be able to discover or disrupt operation of the cluster. To implement this feature, create a file in the global zone of each node and add the names of the non-global zones on that node that should not be part of the cluster. If the zone name appears in the file, all cluster commands and daemons are disabled in that zone. Ensure that the zone is not running when you add or remove a zone name from this file.



---

**Note** – The zones listed in the `cluster.zone.deny` file might be managed by the HA Container agents (the data services documented in the *Oracle Solaris Cluster Data Service for Solaris Containers Guide*). This feature allows you to manage the zone as a "black box", without any cluster software running inside the zone.

---

Perform these procedures while logged into the global zone of the node on which you want to deny or allow cluster services to non-global zones.

## ▼ How to Deny Cluster Services For a Non-Global Zone

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 From the global zone, halt the non-global zone which will be denied cluster services.  

```
zoneadm halt zonename
```
- 3 Open the `/etc/cluster/cluster.zone.deny` file in a text editor, and add the name of the non-global zone that will be denied cluster services and save the file.  
Each zone name you type in the file must end with a new line.
- 4 Reboot the non-global zone.
- 5 Verify that the selected non-global zone has been denied.

a. Log into the selected non-global zone.

b. Run a cluster command. The command fails if the zone has been denied. For example:

```
clnode status
```

The command returns a *This node is not in cluster mode* message if the zone has been denied.

## ▼ How to Allow Cluster Services For a Non-Global Zone

- 1 Become superuser or assume a role that provides `solaris.cluster.modify` RBAC authorization.
- 2 Halt the non-global zone which will rejoin the cluster.
- 3 Open the `etc/cluster/cluster.zone.deny` file in a text editor, remove the name of the non-global zone that is rejoining the cluster, and save the file.

- 4 Reboot the non-global zone.**
- 5 Verify that the non-global zone has rejoined the cluster.**
  - a. Log into the selected non-global zone.**
  - b. Run a cluster command. For example, `clnode status`.**

```
clnode status
```

# Data Service Configuration Worksheets and Examples

---

This appendix provides worksheets for planning resource-related components of your cluster configuration and examples of completed worksheets for your reference. For worksheets for other components of your cluster configuration, see [Appendix A, “Oracle Solaris Cluster Installation and Configuration Worksheets,”](#) in *Oracle Solaris Cluster Software Installation Guide*.

If necessary, make additional copies of a worksheet to accommodate all the resource-related components in your cluster configuration. To complete these worksheets, follow the planning guidelines in *Oracle Solaris Cluster Software Installation Guide* and [Chapter 1, “Planning for Oracle Solaris Cluster Data Services.”](#) Then refer to your completed worksheets during cluster installation and configuration.

---

**Note** – The data that is used in the worksheet examples is intended as a guide only. The examples do not represent a complete configuration of a functional cluster.

---

## Configuration Worksheets

This appendix contains the following worksheets.

- “Resource Types Worksheet” on page 180
- “Network Resources Worksheet” on page 182
- “Application Resources—Failover Worksheet” on page 184
- “Application Resources—Scalable Worksheet” on page 186
- “Resource Groups—Failover Worksheet” on page 188
- “Resource Groups—Scalable Worksheet” on page 190

# Resource Types Worksheet

Use this worksheet for resource types other than logical host or shared address.

| Resource Type Name | Nodes on Which the Resource Type Runs |
|--------------------|---------------------------------------|
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |
|                    |                                       |

EXAMPLE A-1 Resource Types Worksheet

| Resource Type Name   | Nodes on Which the Resource Type Runs |
|----------------------|---------------------------------------|
| SUNW.nshttp          | phys-schost-1, phys-schost-2          |
| SUNW.oracle_listener | phys-schost-1, phys-schost-2          |
| SUNW.oracle_server   | phys-schost-1, phys-schost-2          |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |
|                      |                                       |

# Network Resources Worksheet

| Component                  | Name                              |       |
|----------------------------|-----------------------------------|-------|
| Resource name              |                                   |       |
| Resource group name        |                                   |       |
| Resource type (circle one) | Logical hostname   Shared address |       |
| Resource type name         |                                   |       |
| Dependencies               |                                   |       |
| Hostnames used             |                                   |       |
| Extension properties       | Name                              | Value |
|                            |                                   |       |
|                            |                                   |       |
|                            |                                   |       |
|                            |                                   |       |
|                            |                                   |       |
|                            |                                   |       |
|                            |                                   |       |
|                            |                                   |       |

EXAMPLE A-2 Network Resources—Shared Address Worksheet

| Component                  | Name                      |                         |
|----------------------------|---------------------------|-------------------------|
| Resource name              | <b>sh-galileo</b>         |                         |
| Resource group name        | <b>rg-shared</b>          |                         |
| Resource type (circle one) | <b>Shared address</b>     |                         |
| Resource type name         | <b>SUNW.SharedAddress</b> |                         |
| Dependencies               | <b>none</b>               |                         |
| Hostnames used             | <b>sh-galileo</b>         |                         |
| Extension properties       | <b>Name</b>               | <b>Value</b>            |
|                            | <b>netiflist</b>          | <b>ipmp0@1, ipmp0@2</b> |
|                            |                           |                         |

EXAMPLE A-3 Network Resources—Logical Hostname Worksheet

| Component                  | Name                        |                         |
|----------------------------|-----------------------------|-------------------------|
| Resource name              | <b>relo-galileo</b>         |                         |
| Resource group name        | <b>rg-oracle</b>            |                         |
| Resource type (circle one) | <b>Logical hostname</b>     |                         |
| Resource type name         | <b>SUNW.LogicalHostname</b> |                         |
| Dependencies               | <b>none</b>                 |                         |
| Hostnames used             | <b>relo-galileo</b>         |                         |
| Extension properties       | <b>Name</b>                 | <b>Value</b>            |
|                            | <b>netiflist</b>            | <b>ipmp0@1, ipmp0@2</b> |
|                            |                             |                         |

# Application Resources—Failover Worksheet

|                      |      |       |
|----------------------|------|-------|
| Component            | Name |       |
| Resource name        |      |       |
| Resource group name  |      |       |
| Resource type name   |      |       |
| Dependencies         |      |       |
| Extension Properties | Name | Value |
|                      |      |       |
|                      |      |       |
|                      |      |       |
|                      |      |       |
|                      |      |       |
|                      |      |       |
|                      |      |       |
|                      |      |       |



EXAMPLE A-4 Application Resources—Failover Worksheet

|                      |                      |                         |
|----------------------|----------------------|-------------------------|
| Component            | Name                 |                         |
| Resource name        | oracle-listener      |                         |
| Resource group name  | rg-oracle            |                         |
| Resource type name   | SUNW.oracle_listener |                         |
| Dependencies         | hasp_resource        |                         |
| Extension Properties | Name                 | Value                   |
|                      | ORACLE_HOME          | /global/oracle/orahome/ |
|                      | LISTENER_NAME        | lsnr1                   |
|                      |                      |                         |
|                      |                      |                         |
|                      |                      |                         |
|                      |                      |                         |
|                      |                      |                         |
|                      |                      |                         |

# Application Resources—Scalable Worksheet

| Component                          | Name |       |
|------------------------------------|------|-------|
| Resource name                      |      |       |
| Logical-host resource group name   |      |       |
| Shared-address resource group name |      |       |
| Logical-host resource type name    |      |       |
| Shared-address resource type name  |      |       |
| Dependencies                       |      |       |
| Extension Properties               | Name | Value |
|                                    |      |       |
|                                    |      |       |
|                                    |      |       |
|                                    |      |       |
|                                    |      |       |
|                                    |      |       |
|                                    |      |       |
|                                    |      |       |

EXAMPLE A-5 Application Resources—Scalable Worksheet

|                                    |            |       |
|------------------------------------|------------|-------|
| Component                          | Name       |       |
| Resource name                      | sh-galileo |       |
| Logical-host resource group name   |            |       |
| Shared-address resource group name | rg-shared  |       |
| Logical-host resource type name    |            |       |
| Shared-address resource type name  |            |       |
| Dependencies                       |            |       |
| Extension Properties               | Name       | Value |
|                                    |            |       |
|                                    |            |       |
|                                    |            |       |
|                                    |            |       |
|                                    |            |       |
|                                    |            |       |
|                                    |            |       |
|                                    |            |       |

# Resource Groups—Failover Worksheet

| Component                    | Notes                                                                                                                                                                                                                             | Name     |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| Resource group name          | The name must be unique within the cluster                                                                                                                                                                                        |          |
| Function                     | Describe the function of this resource group                                                                                                                                                                                      |          |
| Failback? (circle one)       | Will this resource group switch back to the primary node after the primary node has failed and been restored?                                                                                                                     | Yes   No |
| Node list                    | Indicate the cluster nodes that can host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries. |          |
| Dependent disk device groups | List the disk device groups on which this resource group depends.                                                                                                                                                                 |          |
| Configuration directory      | If the resources in this resource group need to create files for administrative purposes, include the subdirectory they use.                                                                                                      |          |

EXAMPLE A-6 Example: Resource Groups—Failover Worksheet

| Component                    | Notes                                                                                                                                                                                                                             | Name                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Resource group name          | The name must be unique within the cluster                                                                                                                                                                                        | <b>rg-oracle</b>                                   |
| Function                     | Describe the function of this resource group                                                                                                                                                                                      | <b>Contains the Oracle resources</b>               |
| Failback? (circle one)       | Will this resource group switch back to the primary node after the primary node has failed and been restored?                                                                                                                     | <b>No</b>                                          |
| Node list                    | Indicate the cluster nodes that can host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries. | <b>1) phys-schost-1</b><br><b>2) phys-schost-2</b> |
| Dependent disk device groups | List the disk device groups on which this resource group depends.                                                                                                                                                                 | <b>schost1-dg</b>                                  |
| Configuration directory      | If the resources in this resource group need to create files for administrative purposes, include the subdirectory they use.                                                                                                      |                                                    |

# Resource Groups—Scalable Worksheet

| Component                   | Notes                                                                                                                                                                                                                             | Name     |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| Resource group name         | The name must be unique within the cluster.                                                                                                                                                                                       |          |
| Function                    |                                                                                                                                                                                                                                   |          |
| Maximum number of primaries |                                                                                                                                                                                                                                   |          |
| Desired number of primaries |                                                                                                                                                                                                                                   |          |
| Failback? (circle one)      | Will this resource group switch back to the primary node, after the primary node has failed?                                                                                                                                      | Yes   No |
| Node list                   | Indicate the cluster nodes that can host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries. |          |
| Dependencies                | List any resource groups on which this resource depends.                                                                                                                                                                          |          |

EXAMPLE A-7 Example: Resource Groups—Scalable Worksheet

| Component                   | Notes                                                                                                                                                                                                                             | Name                                               |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Resource group name         | The name must be unique within the cluster.                                                                                                                                                                                       | <b>rg-http</b>                                     |
| Function                    |                                                                                                                                                                                                                                   | <b>Contains the web server resources</b>           |
| Maximum number of primaries |                                                                                                                                                                                                                                   | <b>2</b>                                           |
| Desired number of primaries |                                                                                                                                                                                                                                   | <b>2</b>                                           |
| Failback? (circle one)      | Will this resource group switch back to the primary node, after the primary node has failed?                                                                                                                                      | <b>No</b>                                          |
| Node list                   | Indicate the cluster nodes that can host this resource group. The first node in this list should be the primary, with others being the secondaries. The order of the secondaries will indicate preference for becoming primaries. | <b>1) phys-schost-1</b><br><b>2) phys-schost-2</b> |
| Dependencies                | List any resource groups on which this resource depends.                                                                                                                                                                          | <b>rg-shared</b>                                   |





# Index

---

## A

### adding

- file systems to HAStoragePlus resource, 141–143
- nodes to resource groups
  - overview, 103
  - failover, 104–106
  - scalable, 103–104
- resources to a resource group
  - overview, 55–70
  - failover application, 65–67
  - logical hostname, 56–58, 59–61
  - scalable application, 68–70
  - shared address, 61–63, 63–65

affinities, resource groups, 152–153

Affinity\_timeout, resource property, 28

allowing cluster services, non-global zone, 177–178

API\_version, resource type property, 27

application binaries, determining location, 16–17

Array\_maxsize, resource property attribute, 30

Array\_minsize, resource property attribute, 30

attributes, resource property, 30

Auto\_start\_on\_new\_cluster, resource group property, 29

automatic recovery actions, 75–78

## B

balancing, load on cluster nodes, 155–156

Boot, resource type property, 27

Boot\_timeout, resource property, 28

bringing online, resource groups, 71–72

bypassing, name services, 92–93

## C

caution notice, removal of file systems, 143

### changing

- resource group properties, 90–91

- resource properties, 91–92

- resource type properties, 89–90

Cheap\_probe\_interval, resource property, 28

CheckNameService extension property, 93

### clearing

- Start\_failed resource state, 95–97, 97–98, 99–100

- STOP\_FAILED error flag, 93–95

clnode command, 12

clsetup utility, 25

### clsetup utility

- logical hostname

  - adding to a resource group, 56–58

- shared address

  - adding to a resource group, 61–63

cluster properties, 27

- Concentrate\_load, 27

### collocation

- enforcing for online resource groups, 153–154

- preferring for online resource groups, 154–155

combining, resource group affinities, 159

### Command-Line Interface

- logical hostname

  - adding to a resource group, 59–61

**Command-Line Interface (*Continued*)**

- shared address
  - adding to a resource group, 63–65
- commands, node information, 12
- complete failures, 175
- Concentrate\_load, resource property, 27
- configuration
  - guidelines, 16–18
  - overview, 21–23
  - planning cluster file system, 17
- configuring and administering, Oracle Solaris Cluster data services, 39–40
- Conn\_threshold, resource property, 28
- considerations, 21
- creating
  - failover application resources, 65–67
  - logical hostname resources, 56–58, 59–61
  - resource groups
    - failover, 49–51
    - scalable, 51–53
  - scalable application resources, 68–70
  - shared address resources, 61–63
    - using CLI, 63–65

**D**

- data services
  - considerations, 21
  - planning, 15–33
  - special requirements, 16
- Default, resource property attribute, 30
- defining, persistent faults, 175–176
- delegating, resource group failover or switchover, 157–158
- denying cluster services, non-global zone, 176–178
- Description, resource property attribute, 30
- description values, rules, 33
- Desired\_primaries, resource group property, 29
- device groups
  - relationship with resource groups, 19
  - synchronizing startups with resource groups, 114–120
- directive, #supgrade, 32
- disabled resources, unexpected behavior, 86

- disabling
  - resource fault monitors, 78–79
  - resources, 86–87, 99–100
  - SMF instances, 18
- displaying, resource type, resource group, and resource configuration, 88
- distributing, online resource groups, 151–160
- downgrading, resource types, 47–48

**E**

- editing
  - HASStoragePlus resources, 151
  - logical hostname resources, 101
  - shared address resources, 101
- enabling
  - Oracle Solaris SMF services, 163–173
  - resource fault monitors, 79
  - resources, 73–74, 99–100
- enumeration literal names, rules, 31
- Enumlist, resource property attribute, 30
- error flags, STOP\_FAILED, 93–95
- error messages
  - failure to modify file system, 147, 148
- /etc/vfstab file
  - adding entries to, 141
  - removing entries from, 144
- extension properties
  - Probe\_timeout
    - effect on restart time, 176
    - tuning, 175
  - resource property attribute, 30
  - resource type, 28

**F**

- Failback, resource group property, 29
- failover, delegating for resource groups, 157–158
- Failover, resource type property, 27
- failover application resources, adding to resource group, 65–67
- failover file system across zone clusters, sharing, 134–137

**Failover\_mode**, resource property, 28  
**Failover\_mode** system property, 176  
 failovers, preserving distribution of online resource groups, 151–160  
 fault monitors  
   disabling, 78–79  
   enabling, 79  
   faults detected by, 176  
   probe interval, 174  
   probe timeout, 174–175  
   resource, 78–79  
   response to faults, 176  
   tuning, 173–176  
 faults  
   modification of file systems, 147–148, 148–149  
   persistent, 175–176  
   responses to, 176  
 file systems  
   adding to **HASStoragePlus** resource, 141–143  
   caution notice, 143  
   failure to modify, 147–148, 148–149  
   highly available  
     enabling, 124  
     modifying, 140–149  
   removing from **HASStoragePlus** resource, 143–145  
   unmounting, 143  
 files  
   /etc/vfstab  
     adding entries to, 141  
     removing entries from, 144  
   RTR, 151  
**Fini**, resource type property, 27  
**Fini\_timeout**, resource property, 28  
 format, resource type names, 32  
 freeing node resources, affinities, 156–157

## G

**Generic\_affinity**, resource property, 28  
**Global\_resources\_used**, resource group property, 29  
**Global\_zone**, resource type property, 27  
**Global\_zone\_override**, resource property, 28

## H

**HASStoragePlus**, 134–137  
   managed entity monitoring, 115  
   upgrading from **HASStorage**, 138–140  
**HASStoragePlus** resource  
   cluster file systems  
     change from local file system, 149–150  
     configuring, 121–123  
**HASStoragePlus** resource type  
   caution notice, 143  
   criteria for using, 20–21  
   failure to modify instances of, 147–148, 148–149  
   modifying instances of, 140–149  
   overview, 19–21  
   resource type versions, 150–151  
   upgrading, 150–151  
 help, 12–13  
 highly available file systems  
   adding file systems to, 141–143  
   caution notice, 143  
   enabling, 124  
   failure to modify, 147–148, 148–149  
   modifying, 140–149  
   removing file systems from, 143–145

## I

**Implicit\_network\_dependencies**, resource group property, 29  
**Init**, resource type property, 27  
**Init\_nodes**, resource type property, 27  
**Init\_timeout**, resource property, 28  
 installation, overview, 21–23  
**Installed\_nodes**, resource type property, 27  
 Internet Protocol (IP) addresses, restrictions, 21  
 intervals, fault monitor probes, 174  
**IOption** property, 116  
**IPsec\_session\_failover**, resource property, 28  
**Is\_logical\_hostname**, resource type property, 27  
**Is\_shared\_address**, resource type property, 27

**L**

## legal names

Resource Group Manager (RGM), 31–33

## load balancing, 155–156

Load\_balancing\_policy, resource property, 28

Load\_balancing\_weights, resource property, 28

Load\_factors, resource group property, 29

local file systems, list of supported, 124

## logical hostname resources

adding to resource group

using CLI, 59–61

using clsetup utility, 56–58

modifying, 92–93

## loopback mount

using HASStoragePlus, 121–123, 124

**M**

Max, resource property attribute, 30

Maximum\_primaries, resource group property, 29

maximum values, restarts, 175

MaxLength, resource property attribute, 30

messages file, 13

## migrating

HASStoragePlus resources, 151

logical hostname resources, 101

shared address resources, 101

migrating to a new resource type version, 43–47

Min, resource property attribute, 30

Minlength, resource property attribute, 30

mission-critical services, 156–157

modifying resources, 92–93

Monitor\_check, resource type property, 27

Monitor\_check\_timeout, resource property, 28

Monitor\_start, resource type property, 27

Monitor\_start\_timeout, resource property, 28

Monitor\_stop, resource type property, 27

Monitor\_stop\_timeout, resource property, 28

Monitored\_switch, resource property, 28

**N**

name services, bypassing, 92–93

Network\_resources\_used (deprecated), resource property, 28

networks, restrictions, 21

NodeList, resource group property, 29

NodeList resource group property, and affinities, 153

## nodes

adding to resource groups

overview, 103

failover, 104–106

scalable, 103–104

distributing resource groups on, 151–160

load balancing, 155–156

off-loading noncritical services, 156–157

removing from resource groups

overview, 106

failover, 108–109

failover with shared addresses, 109–110

scalable, 107–108

## non-global zone

allowing cluster services, 177–178

denying cluster services, 177

noncritical services, off-loading, 156–157

nsswitch.conf, verifying file contents, 17

Num\_resource\_restarts, resource property, 28

Num\_rg\_restarts, resource property, 28

**O**

off-loading noncritical resource groups,

affinities, 156–157

On\_off\_switch, resource property, 28

Oracle Service Management Facility (SMF)

enabling, 163–173

encapsulating into a failover proxy

resource, 164–167

encapsulating into a multi-master resource

type, 167–169

encapsulating into a scalable proxy resource, 170

Oracle Service Management Facility (SMF) services,

enabling, 18

Oracle Solaris Cluster Administrative commands, 25

Oracle Solaris Cluster Manager GUI, 24–25

Outgoing\_Connection, resource property, 28

**P**

partial failures, 175  
 Pathprefix, resource group property, 29  
 Per\_node, resource property attributes, 30  
 performance  
   effect of probe interval on, 174  
   optimizing for mission-critical services, 156–157  
 persistent faults, defining, 175–176  
 ping command, responses from disabled resources, 86  
 Pingpong\_interval, resource group property, 29  
 Pkglist, resource type property, 27  
 planning  
   cluster file system, 17  
   data services, 15–33  
 Port\_list, resource property, 28  
 Postnet\_stop, resource type property, 27  
 Postnet\_stop\_timeout, resource property, 28  
 Preemption\_mode, resource group property, 29  
 Prenet\_start, resource type property, 27  
 Prenet\_start\_timeout, resource property, 28  
 preregistered resource types  
   reregistering after inadvertent deletion, 101–102  
   upgrading, 100–101  
 Priority, resource group property, 29  
 Probe\_timeout extension property  
   effect on restart time, 176  
   tuning, 175  
 properties  
   *See also* extension properties  
   cluster, 27  
   resource, 28–29  
   resource group, 29–30  
   Type\_version, 101, 151  
 Property, resource property attribute, 30  
 property attributes, resource, 30  
 property names, rules, 31  
 property values, rules, 33  
 Proxy, resource type property, 27  
 prtconf -v command, 12  
 prtdiag -v command, 12  
 psrinfo -v command, 12

**Q**

quiescing resource groups, 74–75

**R**

R\_description, resource property, 28  
 recovering  
   from failure to modify file system, 147–148, 148–149  
 recovery actions  
   resuming automatic, 75–78  
   suspending automatic, 75–78  
 registering  
   HASToragePlus resource type  
     during upgrade, 150–151  
   preregistered resource types, 101–102  
   resource types, 40–41  
   SUNW.LogicalHostname resource type  
     after inadvertent deletion, 101–102  
     during upgrade, 100–101  
   SUNW.SharedAddress resource type  
     after inadvertent deletion, 101–102  
     during upgrade, 100–101  
 removing  
   file systems from HASToragePlus resource, 143–145  
   nodes from resource groups  
     overview, 106  
     failover, 108–109  
     failover with shared addresses, 109–110  
     scalable, 107–108  
   resource group, 81–82  
   resource types, 80–81  
   resources, 82–83  
 replicating, configuration data, 161  
 requirements, data services, 16  
 reregistering preregistered resource types, 101–102  
 Resource\_dependencies, resource property, 28  
 Resource\_dependencies\_offline\_restart, resource property, 29  
 Resource\_dependencies\_restart, resource property, 29  
 Resource\_dependencies\_weak, resource property, 29  
 resource fault monitors, 78–79  
 resource group affinities, zone cluster, 159–160

**Resource Group Manager (RGM)**

legal names, 31–33

values, 33

resource group names, rules, 31

resource group properties, 29–30

Auto\_start\_on\_new\_cluster, 29

Desired\_primaries, 29

Failback, 29

Global\_resources\_used, 29

Implicit\_network\_dependencies, 29

Load\_factors, 29

Maximum\_primaries, 29

Nodelist, 29

Pathprefix, 29

Pingpong\_interval, 29

Preemption\_mode, 29

Priority, 29

Resource\_list, 29

RG\_affinities, 29

RG\_dependencies, 30

RG\_description, 30

RG\_is\_frozen, 30

RG\_mode, 30

RG\_name, 30

RG\_project\_name, 30

RG\_slm\_cpu, 30

RG\_slm\_cpu\_min, 30

RG\_slm\_pset\_type, 30

RG\_slm\_type, 30

RG\_state, 30

RG\_system, 30

Suspend\_automatic\_recovery, 30

**resource groups**

adding nodes to, 103

adding resources, 55–70

failover application, 65–67

logical hostname, 56–58, 59–61

scalable application, 68–70

shared address, 61–63, 63–65

affinities, 152–153

bringing online, 71–72

changing properties, 90–91

creating

failover, 49–51

**resource groups, creating (*Continued*)**

scalable, 51–53

delegating failover or switchover, 157–158

displaying configuration information, 88

distributing among nodes, 151–160

enforced collocation, 153–154

enforced separation, 156–157

even distribution, 155–156

failover

adding nodes, 104–106

removing nodes from, 108–109

failover with shared addresses

removing nodes from, 109–110

moving to UNMANAGED state, 86–87

preferred collocation, 154–155

preferred separation, 155–156

quiescing, 74–75

relationship with device groups, 19

removing, 81–82

removing nodes, 106

restarting, 97–98

resuming automatic recovery actions, 75–78

retrieving, replicating or upgrading configuration

data about, 160

scalable

adding nodes, 103–104

removing nodes from, 107–108

suspending automatic recovery actions, 75–78

switching over, 95–97

switching the current primary, 83–85

switching to preferred primaries, 72–73

synchronizing startups with device groups, 114–120

**Resource\_list**

resource group property, 29

resource type property, 27

Resource\_name, resource property, 29

resource names, rules, 31

Resource\_project\_name, resource property, 29

resource properties, 28–29

Affinity\_timeout, 28

Boot\_timeout, 28

Cheap\_probe\_interval, 28

Conn\_threshold, 28

extension, 28

resource properties (*Continued*)

- Failover\_mode, 28
- Fini\_timeout, 28
- Generic\_affinity, 28
- Global\_zone\_override, 28
- Init\_timeout, 28
- IPsec\_session\_failover, 28
- Load\_balancing\_policy, 28
- Load\_balancing\_weights, 28
- Monitor\_check\_timeout, 28
- Monitor\_start\_timeout, 28
- Monitor\_stop\_timeout, 28
- Monitored\_switch, 28
- Network\_resources\_used (deprecated), 28
- Num\_resource\_restarts, 28
- Num\_rg\_restarts, 28
- On\_off\_switch, 28
- Outgoing\_Connection, 28
- Port\_list, 28
- Postnet\_stop\_timeout, 28
- Prenet\_start\_timeout, 28
- R\_description, 28
- Resource\_dependencies, 28
- Resource\_dependencies\_offline\_restart, 29
- Resource\_dependencies\_restart, 29
- Resource\_dependencies\_weak, 29
- Resource\_name, 29
- Resource\_project\_name, 29
- Resource\_state, 29
- Retry\_count, 29
- Retry\_interval, 29
- Round\_robin, 29
- Scalable, 29
- Start\_timeout, 29
- Status, 29
- Status\_msg, 29
- Stop\_timeout, 29
- Thorough\_probe\_interval, 29
- Type, 29
- Type\_version, 29
- UDP\_affinity, 29
- Update\_timeout, 29
- Validate\_timeout, 29
- Weak\_affinity, 29

## resource property attributes, 30

- Array\_maxsize, 30
- Array\_minsize, 30
- Default, 30
- Description, 30
- Enumlist, 30
- Extension, 30
- Max, 30
- Maxlength, 30
- Min, 30
- Minlength, 30
- Per\_node, 30
- Property, 30
- Tunable, 30
- type, 30

Resource\_state, resource property, 29

Resource\_type, resource type property, 27

resource type, ScalMountPoint, 53

resource type names, rules, 32

## resource type properties

- API\_version, 27
- Boot, 27
- Failover, 27
- Fini, 27
- Global\_zone, 27
- Init, 27
- Init\_nodes, 27
- Installed\_nodes, 27
- Is\_logical\_hostname, 27
- Is\_shared\_address, 27
- Monitor\_check, 27
- Monitor\_start, 27
- Monitor\_stop, 27
- Pkglist, 27
- Postnet\_stop, 27
- Prenet\_start, 27
- Proxy, 27
- Resource\_list, 27
- Resource\_type, 27
- RT\_basedir, 27
- RT\_description, 28
- RT\_system, 28
- RT\_version, 28
- Single\_instance, 28

**resource type properties** (*Continued*)

- Start, 28
- Stop, 28
- Update, 28
- Validate, 28
- Vendor\_ID, 28

**resource type registration (RTR) file**, 151**resource types**

- changing properties, 89–90
- displaying configuration information, 88
- downgrading, 47–48
- HASStoragePlus
  - existing resources, 120
  - migrating instances of, 151
  - new resources, 118–120
- LogicalHostname
  - migrating instances of, 101
- migrating to a new resource type version, 43–47
- preregistered
  - reregistering after inadvertent deletion, 101–102
  - upgrading, 100–101
- registering, 40–41
- removing, 80–81
- retrieving, replicating or upgrading configuration data about, 160
- SharedAddress
  - migrating instances of, 101
- unregistering, 80–81
- upgrading, 42–43

**resources**

- adding to a resource group, 55–70
- changing properties, 91–92
- clearing STOP\_FAILED error flag, 93–95
- disabling, 86–87, 99–100
- disabling fault monitors, 78–79
- displaying configuration information, 88
- enabling, 73–74, 99–100
- enabling fault monitors, 79
- failover application
  - adding to a resource group, 65–67
- logical hostname
  - adding to a resource group, 56–58, 59–61
  - modifying, 92–93
- removing, 82–83

**resources** (*Continued*)

- removing resource types, 80–81
- retrieving, replicating or upgrading configuration data about, 160
- scalable application
  - adding to a resource group, 68–70
- shared address
  - adding to a resource group, 61–63, 63–65
  - isolating from hosts when disabled, 86
  - modifying, 92–93
- responses, to faults, 176
- restarting, resource groups, 97–98
- restarts, maximum allowed, 175
- restrictions, 21
- Retry\_count, resource property, 29
- Retry\_count system property, 176
- retry interval, 175
- Retry\_interval, resource property, 29
- Retry\_interval system property, 176
- RG\_affinities, resource group property, 29
- RG\_affinities resource group property, 152–153
- RG\_dependencies, resource group property, 30
- RG\_description, resource group property, 30
- RG\_is\_frozen, resource group property, 30
- RG\_mode, resource group property, 30
- RG\_name, resource group property, 30
- RG\_project\_name, resource group property, 30
- RG\_slm\_cpu, resource group property, 30
- RG\_slm\_cpu\_min, resource group property, 30
- RG\_slm\_pset\_type, resource group property, 30
- RG\_slm\_type, resource group property, 30
- RG\_state, resource group property, 30
- RG\_system, resource group property, 30
- Round\_robin, resource property, 29
- RT\_basedir, resource type property, 27
- RT\_description, resource type property, 28
- RT\_system, resource type property, 28
- RT\_version, resource type property, 28
- rules
  - description values, 33
  - enumeration literal names, 31
  - property names, 31
  - property values, 33
  - resource group names, 31



rules (*Continued*)

- resource names, 31

**S**

- Scalable, resource property, 29

- scalable application resources, adding to resource group, 68–70

- ScalMountPoint resource type, creating, 53

- retrieving, configuration data about resource groups, resource types, and resources, 161

- scsnapshot utility, 160

- setting up

- HASStoragePlus resource type, 124
  - existing resources, 120
  - new resources, 118–120

- shared address resources

- adding to resource group, 63–65

- using clsetup utility, 61–63

- isolating from hosts when disabled, 86

- modifying, 92–93

- sharing a failover file system, zone clusters, 134–137

- show - rev subcommand, 12

- showrev -p command, 12

- Single\_instance, resource type property, 28

- SMF

- See also* Oracle Service Management Facility (SMF)

- enabling, 163–173

- encapsulating into a failover proxy resource, 164–167

- encapsulating into a multi-master proxy resource, 167–169

- encapsulating into a scalable proxy resource, 170

- Oracle Service Management Facility (SMF), 18

- special requirements, identifying, 16

- Start, resource type property, 28

- Start\_failed resource state

- clearing, 95–97, 97–98, 99–100

- Start\_timeout, resource property, 29

- Status, resource property, 29

- Status\_msg, resource property, 29

- Stop, resource type property, 28

- STOP\_FAILED error flag, 93–95

- Stop\_timeout, resource property, 29

- strong negative affinity

- definition, 152

- usage example, 156–157

- strong positive affinity

- definition, 152

- usage example, 153–154

- strong positive affinity with failover delegation

- definition, 152

- usage example, 157–158

- Sun Management Center GUI, 25

- Sun StorageTek Availability Suite, managed entity monitoring, 115

- SUNW.LogicalHostname resource type

- reregistering after inadvertent deletion, 101–102

- resource type versions, 100

- upgrading, 100–101

- SUNW.SharedAddress resource type

- reregistering after inadvertent deletion, 101–102

- resource type versions, 100

- upgrading, 100–101

- Suspend\_automatic\_recovery, resource group

- property, 30

- switching, resource groups, 95–97

- switching the current primary, resource groups, 83–85

- switching to preferred primaries, resource

- groups, 72–73

- switchover, delegating for resource groups, 157–158

- synchronizing startups, resource groups and device groups, 114–120

- syntax

- description values, 33

- enumeration literal names, 31

- property names, 31

- property values, 33

- resource group names, 31

- resource names, 31

- resource type names, 32

- system properties

- See also* extension properties

- See also* properties

- effect on fault monitors, 173

- Failover\_mode, 176

- Retry\_count, 176

- Retry\_interval, 176

system properties (*Continued*)

- Thorough\_probe\_interval
  - effect on restart time, 176
  - tuning, 174

**T**

- technical support, 12–13
- Thorough\_probe\_interval, resource property, 29
- Thorough\_probe\_interval system property
  - effect on restart time, 176
  - tuning, 174
- timeouts
  - fault monitor
    - guidelines for setting, 174–175
- tools
  - clsetup utility, 25
  - Oracle Solaris Cluster Administrative commands, 25
  - Oracle Solaris Cluster Manager GUI, 24–25
  - Sun Management Center GUI, 25
- troubleshooting
  - modification of file systems, 147–148, 148–149
- Tunable, resource property attribute, 30
- tuning, fault monitors, 173–176
- Type, resource property, 29
- type, resource property attributes, 30
- Type\_version, resource property, 29
- Type\_version property, 101, 151

**U**

- UDP\_affinity, resource property, 29
- unmounting, file systems, 143
- unregistering, resource types, 80–81
- Update, resource type property, 28
- Update\_timeout, resource property, 29
- upgrade directive, 32
- upgrading
  - configuration data, 162
  - HASStoragePlus resource type, 150–151
  - preregistered resource types, 100–101
  - resource types, 42–43

**V**

- Validate, resource type property, 28
- Validate\_timeout, resource property, 29
- values, Resource Group Manager (RGM), 33
- /var/adm/messages file, 13
- Vendor\_ID, resource type property, 28
- verifying
  - addition of file systems to HASStoragePlus resource, 142
  - nsswitch.conf file contents, 17
  - removal of file systems from HASStoragePlus resource, 144
- versions
  - HASStoragePlus resource type, 150–151
  - SUNW.LogicalHostname resource type, 100
  - SUNW.SharedAddress resource type, 100
- vfstab file
  - adding entries to, 141
  - adding removing entries from, 144
- volume managers, highly available file systems, 125

**W**

- Weak\_affinity, resource property, 29
- weak negative affinity
  - definition, 152
  - usage example, 155–156
- weak positive affinity
  - definition, 152
  - usage example, 154–155

**Z**

- zone cluster, resource group affinities, 159–160
- Zpools property, recovering from a fault, 148–149