



Sun Cluster Data Service for Solaris Containers Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819-2664-10
August 2005, Revision A

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



050511 @ 11223



Contents

Preface	5
Installing and Configuring Sun Cluster HA for Solaris Containers	11
Sun Cluster HA for Solaris Containers Overview	11
Overview of Installing and Configuring Sun Cluster HA for Solaris Containers	13
Planning the Sun Cluster HA for Solaris Containers Installation and Configuration	13
Configuration Restrictions	14
Configuration Requirements	15
Installing and Configuring Zones	17
▼ How to Enable a Zone to Run in a Failover Configuration	17
▼ How to Enable a Zone to Run in a Multiple-Masters Configuration	18
▼ How to Install and Configure a Zone	18
Verifying the Installation and Configuration of a Zone	19
▼ How to Verify the Installation and Configuration of a Zone	19
Installing the Sun Cluster HA for Solaris Containers Packages	20
▼ How to Install the Sun Cluster HA for Solaris Containers Packages	20
Registering and Configuring Sun Cluster HA for Solaris Containers	21
Specifying Configuration Parameters for the Zone Boot Resource	22
Writing a Zone Script	25
Specifying Configuration Parameters for the Zone Script Resource	26
Writing an SMF Service Probe	28
Specifying Configuration Parameters for the Zone SMF Resource	29
▼ How to Create and Enable Resources for the Zone Boot Component	30
▼ How to Create and Enable Resources for the Zone Script Component	31
▼ How to Create and Enable Resources for the Zone SMF Component	31

Verifying the Sun Cluster HA for Solaris Containers and Configuration	32
▼ How to Verify the Sun Cluster HA for Solaris Containers Installation and Configuration	32
Patching the Global Zone and Local Zones	33
▼ How to Apply a Non Rebooting Patch to the Global Zone and Local Zones	33
▼ How to Apply a Rebooting Patch to the Global Zone and Local Zones	33
Tuning the Sun Cluster HA for Solaris Containers Fault Monitors	35
Operation of the Sun Cluster HA for Solaris Containers Parameter File	35
Operation of the Fault Monitor for the Zone Boot Component	35
Operation of the Fault Monitor for the Zone Script Component	36
Operation of the Fault Monitor for the Zone SMF Component	36
Debugging Sun Cluster HA for Solaris Containers	37
▼ How to Activate Debugging for Sun Cluster HA for Solaris Containers	37
A Files for Configuring Sun Cluster HA for Solaris Zones Resources	39
Listing of sczbt_config	39
Listing of sczsh_config	41
Listing of sczsmf_config	42

Index	45
--------------	-----------

Preface

Sun Cluster Data Service for Solaris Containers Guide explains how to install and configure Sun™ Cluster HA for Solaris Zones on both SPARC® and x86 based systems.

Note – In this document, the term “x86” refers to the Intel 32-bit family of microprocessor chips and compatible microprocessor chips made by AMD.

This document is intended for system administrators with extensive knowledge of Sun 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 document assume knowledge of the Solaris™ Operating System and expertise with the volume manager software that is used with Sun Cluster.

Note – Sun Cluster software runs on two platforms, SPARC and x86. The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

UNIX Commands

This document contains information about commands that are specific to installing and configuring Sun 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 Solaris Operating System
- Solaris Operating System man pages
- Other software documentation that you received with your system

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface or Symbol	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>	Command-line placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . Perform a <i>patch analysis</i> . Do <i>not</i> save the file. [Note that some emphasized items appear bold online.]

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Related Documentation

Information about related Sun Cluster topics is available in the documentation that is listed in the following table. All Sun Cluster documentation is available at <http://docs.sun.com>.

Topic	Documentation
Data service administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides
Concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Overview	<i>Sun Cluster Overview for Solaris OS</i>
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
System administration	<i>Sun Cluster System Administration Guide for Solaris OS</i>
Hardware administration	<i>Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Data service development	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
Error messages	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
Command and function reference	<i>Sun Cluster Reference Manual for Solaris OS</i>

For a complete list of Sun Cluster documentation, see the release notes for your release of Sun Cluster at <http://docs.sun.com>.

Related Third-Party Web Site References

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

Note – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun 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. Sun 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, Support, and Training

Sun Function	URL	Description
Documentation	http://www.sun.com/documentation/	Download PDF and HTML documents, and order printed documents
Support and Training	http://www.sun.com/supporttraining/	Obtain technical support, download patches, and learn about Sun courses

Product Training

Sun Microsystems offers training in many Sun technologies through a variety of instructor-led courses and self-paced courses. For information about the training courses that Sun offers and to enroll in a class, visit Sun Microsystems Training at <http://training.sun.com/>.

Getting Help

If you have problems installing or using Sun 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 and serial numbers of your systems
- The release number of the Solaris Operating System (for example, Solaris 8)
- The release number of Sun Cluster (for example, Sun Cluster 3.0)

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
SPARC: <code>prtdiag -v</code>	Displays system diagnostic information
<code>scinstall -pv</code>	Displays Sun Cluster release and package version information

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

Installing and Configuring Sun Cluster HA for Solaris Containers

This chapter explains how to install and configure Sun Cluster HA for Solaris Containers.

This chapter contains the following sections.

- “Overview of Installing and Configuring Sun Cluster HA for Solaris Containers” on page 13
- “Sun Cluster HA for Solaris Containers Overview” on page 11
- “Planning the Sun Cluster HA for Solaris Containers Installation and Configuration” on page 13
- “Installing and Configuring Zones” on page 17
- “Verifying the Installation and Configuration of a Zone” on page 19
- “Installing the Sun Cluster HA for Solaris Containers Packages ” on page 20
- “Registering and Configuring Sun Cluster HA for Solaris Containers” on page 21
- “Verifying the Sun Cluster HA for Solaris Containers and Configuration” on page 32
- “Tuning the Sun Cluster HA for Solaris Containers Fault Monitors” on page 35
- “Debugging Sun Cluster HA for Solaris Containers” on page 37

Sun Cluster HA for Solaris Containers Overview

A Solaris Container is a complete runtime environment for applications. Solaris 10 Resource Manager and Solaris Zones software partitioning technology are both parts of the container. These components address different qualities the container can deliver and work together to create a complete container. The zones portion of the

container provides a virtual mapping from the application to the platform resources. Zones allow application components to be isolated from one application even though the zones share a single instance of the Solaris Operating System. Resource management features permit you to allocate the quantity of resources that a workload receives.

The Solaris Zones facility in the Solaris Operating System provides an isolated and secure environment in which to run applications on your system. When you create a zone, you produce an application execution environment in which processes are isolated from the rest of the system.

This isolation prevents processes that are running in one zone from monitoring or affecting processes that are running in other zones. Even a process that is running with superuser credentials cannot view or affect activity in other zones. A zone also provides an abstract layer that separates applications from the physical attributes of the machine on which they are deployed. Examples of these attributes include physical device paths.

Every Solaris system contains a global zone. The global zone has a dual function. The global zone is both the default zone for the system and the zone that is used for system-wide administrative control. Non-global zones are referred to as zones and are created by the global administrator.

Sun Cluster HA for Solaris Containers enables Sun Cluster to manage Solaris Zones by providing components to perform the following operations:

- The orderly booting and shutdown of a zone
- The orderly startup, shutdown, and fault monitoring of an application within the zone via scripts or commands
- The orderly startup, shutdown, and fault monitoring of a Solaris Service Management Facility (SMF) service within the zone

You can configure Sun Cluster HA for Solaris Containers as a failover service or a multiple-masters service. You *cannot* configure Sun Cluster HA for Solaris Containers as a scalable service.

When a Solaris Zone is managed by the Sun Cluster HA for Solaris Containers data service, the Solaris Zone becomes a failover Solaris Zone, or multiple-masters Solaris Zone, across the Sun Cluster nodes. The failover is managed by the Sun Cluster HA for Solaris Containers data service, which runs only within the global zone.

For conceptual information about failover data services, multiple-masters data services, and scalable data services, see *Sun Cluster Concepts Guide for Solaris OS*.

Overview of Installing and Configuring Sun Cluster HA for Solaris Containers

The following table summarizes the tasks for installing and configuring Sun Cluster HA for Solaris Containers and provides cross-references to detailed instructions for performing these tasks. Perform the tasks in the order that they are listed in the table.

TABLE 1 Tasks for Installing and Configuring Sun Cluster HA for Solaris Containers

Task	Instructions
Plan the installation	“Planning the Sun Cluster HA for Solaris Containers Installation and Configuration” on page 13
Install and configure the Solaris Zones	“Installing and Configuring Zones” on page 17
Verify installation and configuration	“How to Verify the Installation and Configuration of a Zone” on page 19
Install Sun Cluster HA for Solaris Containers Packages	“Installing the Sun Cluster HA for Solaris Containers Packages ” on page 20
Register and configure Sun Cluster HA for Solaris Containers components	“Registering and Configuring Sun Cluster HA for Solaris Containers” on page 21
Verify Sun Cluster HA for Solaris Containers Installation and Configuration	“Verifying the Sun Cluster HA for Solaris Containers and Configuration” on page 32
Tune the Sun Cluster HA for Solaris Containers fault monitors	“Tuning the Sun Cluster HA for Solaris Containers Fault Monitors” on page 35
Debug Sun Cluster HA for Solaris Containers	“Debugging Sun Cluster HA for Solaris Containers” on page 37

Planning the Sun Cluster HA for Solaris Containers Installation and Configuration

This section contains the information you need to plan your Sun Cluster HA for Solaris Containers installation and configuration.

Configuration Restrictions

The configuration restrictions in the subsections that follow apply only to Sun Cluster HA for Solaris Containers.



Caution – Your data service configuration might not be supported if you do not observe these restrictions.

Restrictions for Zone Network Addresses

The configuration of a zone's network addresses depends on the level of high availability you require. You can choose between no HA, HA through the use of IPMP, or HA through the use of IPMP and `SUNW.LogicalHostName`.

Your choice of a zone's network addresses configuration affects some configuration parameters for the zone boot resource. For more information, see [“Registering and Configuring Sun Cluster HA for Solaris Containers”](#) on page 21

- If HA for the zone's addresses is not required then configure the zone's addresses by using the `zonecfg` utility.
- If HA through IPMP protection is required then configure the zone's addresses by using the `zonecfg` utility and then place the zone's addresses on an adapter within an IPMP group.
- If HA through IPMP protection and protection against the failure of all physical interfaces is required, choose one option from the following list:
 - If you require the `SUNW.LogicalHostName` resource type to manage one or a subset of the zone's addresses, configure a `SUNW.LogicalHostName` resource for those zone's addresses and not by using the `zonecfg` utility. Use the `zonecfg` utility only to configure the zones's addresses that are not required to be under the control of the `SUNW.LogicalHostName` resource type.
 - If you require the `SUNW.LogicalHostName` resource type to manage all the zone's addresses, configure a `SUNW.LogicalHostName` resource with a list of the zone's addresses and do not configure them by using the `zonecfg` utility.
 - Otherwise, configure the zone's addresses by using the `zonecfg` utility and configure a separate redundant IP address for use by a `SUNW.LogicalHostName` resource, which must not be configured using the `zonecfg` utility.

Restrictions for a Failover Zone

The zone path of a zone in a failover configuration must reside on a highly available local file system. The zone must be configured on each cluster node where the zone can reside.

The zone is active on only one node at a time, and the zone's address is plumbed on only one node at a time. Application clients can then reach the zone through the zone's address, wherever that zone resides within the cluster.

Restrictions for a Multiple-Masters Zone

The zone path of a zone in a multiple-masters configuration must reside on the local disks of each node. The zone must be configured with the same name on each node that can master the zone.

Each zone that is configured to run within a multiple-masters configuration must also have a zone-specific address. Load balancing for applications in these configurations is typically provided by an external load balancer. You must configure this load balancer for the address of each zone. Application clients can then reach the zone through the load balancer's address.

Restrictions for the Zone Path of a Zone

The zone path of a zone that Sun Cluster HA for Solaris Containers manages cannot reside on a global file system.

- If the zone is in a failover configuration the zone path must reside on a highly available local file system.
- If the zone is in a multiple-masters configuration, the zone path must reside on the local disks of each node.

Restrictions on Major Device Numbers in

`/etc/name_to_major`

For shared devices, Sun Cluster requires that the major and minor device numbers are identical on all nodes in the cluster. If the device is required for a zone, ensure that the major device number is the same in `/etc/name_to_major` on all nodes in the cluster that will host the zone.

Configuration Requirements

The configuration requirements in this section apply only to Sun Cluster HA for Solaris Containers.



Caution – If your data service configuration does not conform to these requirements, the data service configuration might not be supported.

Dependencies Between Sun Cluster HA for Solaris Containers Components

The dependencies between the Sun Cluster HA for Solaris Containers components are described in the following table:

TABLE 2 Dependencies Between Sun Cluster HA for Solaris Containers Components

Component	Dependency
Zone boot resource	<p>SUNW.HAStoragePlus - In a failover configuration, the zone's zone path must be on a highly available file system managed by a SUNW.HAStoragePlus resource .</p> <p>SUNW.LogicalHostName - This dependency is required only if the zone's address is managed by a SUNW.LogicalHostName resource .</p>
Zone script resource	Zone boot resource
Zone SMF resource	Zone boot resource

These dependencies are set when you register and configure Sun Cluster HA for Solaris Containers. For more information, see [“Registering and Configuring Sun Cluster HA for Solaris Containers”](#) on page 21.

The zone script resource and SMF resource are optional. If used, multiple instances of the zone script resource and SMF resource can be deployed within the same resource group as the zone boot resource. Furthermore, if more elaborate dependencies are required then refer to the `r_properties(5)` and `rg_properties(5)` man pages for further dependencies and affinities settings.

Parameter File Directory for Sun Cluster HA for Solaris Containers

The boot component and script component of Sun Cluster HA for Solaris Containers require a parameter file to pass configuration information to the data service. You must create a directory for these files. The directory location must be available on the node that is to host the zone and must not be in the zone's zone path. The directory must be accessible *only* from the global zone. The parameter file for each component is created automatically when the resource for the component is registered.

Installing and Configuring Zones

Installing and configuring Solaris Zones involves the following tasks:

1. Enabling a zone to run in your chosen data service configuration, as explained in the following sections:
 - “How to Enable a Zone to Run in a Failover Configuration” on page 17
 - “How to Enable a Zone to Run in a Multiple-Masters Configuration” on page 18
2. Installing and configuring a zone, as explained in:
 - “How to Install and Configure a Zone” on page 18

Perform this task for each zone that you are installing and configuring. This section explains only the special requirements for installing Solaris Zones for use with Sun Cluster HA for Solaris Containers. For complete information about installing and configuring Solaris Zones, see *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

▼ How to Enable a Zone to Run in a Failover Configuration

- Steps**
1. Register the `SUNW.HAStoragePlus` resource type.

```
# scrgadm -a -t SUNW.HAStoragePlus
```

2. Create a failover resource group.

```
# scrgadm -a -g solaris-zone-resource-group
```

3. Create a resource for the zone’s disk storage.

```
# scrgadm -a -j solaris-zone-has-resource \
-g solaris-zone-resource-group \
-t SUNW.HAStoragePlus \
-x FilesystemMountPoints=solaris-zone-instance-mount-points
```

4. (Optional) Create a resource for the zone’s logical hostname.

```
# scrgadm -a -L -j solaris-zone-logical-hostname-resource-name \
-g solaris-zone-resource-group \
-l solaris-zone-logical-hostname
```

5. Enable the failover resource group.

```
# scswitch -Z -g solaris-zone-resource-group
```

▼ How to Enable a Zone to Run in a Multiple-Masters Configuration

Steps 1. Create a scalable resource group.

```
# scrgadm -a -g solaris-zone-resource-group \  
-y Maximum primaries=max-number \  
-y Desired primaries=desired-number
```

2. Enable the scalable resource group.

```
# scswitch -z -g solaris-zone-resource-group
```

▼ How to Install and Configure a Zone

Note – For complete information about installing a zone refer to *System Administration Guide: Solaris Containers-Resource Management and Solaris Zones*.

Before You Begin Determine the following requirements for the deployment of the zone with Sun Cluster:

- The number of Solaris Zone instances that are to be deployed.
- The cluster file system will be used by each Solaris Zone instance.

Steps 1. Install the zone.

Note – If the zone that you are installing is to become a failover zone, the zone's zone path must specify a highly available local file system. The file system must be managed by the SUNW.HAStoragePlus resource that you created in [Step 3](#).

2. Log in to the zone's console.

```
# zlogin -C zone
```

You are prompted to configure the zone.

3. Follow the prompts to configure the zone.

After you configure the zone has been created, an entry will exist in the `/etc/zones/index` file.

4. Disconnect from the zone's console.

Use the escape sequence that you defined for the zone. If you did not define an escape sequence, use the default escape sequence as follows:

```
# ~.
```

5. **Determine the new zone's index entry by listing the contents of the `/etc/zones/index`.**

You need the new zone's index entry for [“How to Enable a Zone to Run in a Failover Configuration”](#) on page 17

```
# cat /etc/zones/index
```

6. **Make the zone available to all nodes in the cluster.**

Perform the following steps on each cluster node.

- a. **Log in to each cluster node.**
- b. **To prevent a loss of data, create a backup copy of the `/etc/zones/index` file.**

```
# cd /etc/zones
# cp index index_backup
```

- c. **Using a plain text editor, add the entry for the zone to the `/etc/zones/index` file on the node.**
- d. **Copy the `zone.xml` file to the `/etc/zones/index` directory on the node.**

```
# rcp zone-install-node:/etc/zones/zone.xml .
```

Verifying the Installation and Configuration of a Zone

Before you install the Sun Cluster HA for Solaris Containers packages, verify that the zones that you created are correctly configured to run in a cluster. This verification does not verify that the zones are highly available because the Sun Cluster HA for Solaris Containers data service is not yet installed.

▼ How to Verify the Installation and Configuration of a Zone

Perform this procedure for each zone that you created in [“Installing and Configuring Zones”](#) on page 17

Steps 1. Start the Zone.

```
# zoneadm -z zone boot
```

2. Log in to the zone.

```
# zlogin -z zone
```

3. Confirm that the zone has reached the **svc:/milestone/multi-user-server:default milestone.**

```
# svcs -a | grep milestone
online Apr_10 svc:/milestone/network:default
online Apr_10 svc:/milestone/devices:default
online Apr_10 svc:/milestone/single-user:default
online Apr_10 svc:/milestone/sysconfig:default
online Apr_10 svc:/milestone/name-services:default
online Apr_10 svc:/milestone/multi-user:default
online Apr_10 svc:/milestone/multi-user-server:default
online Apr_10 svc:/system/cluster/cl-svc-cluster-milestone:default
```

4. Stop the zone.

```
# zoneadm -z halt zone
```

Installing the Sun Cluster HA for Solaris Containers Packages

If you did not install the Sun Cluster HA for Solaris Containers packages during your initial Sun Cluster installation, perform this procedure to install the packages. Perform this procedure on each cluster node where you are installing the Sun Cluster HA for Solaris Containers packages. To complete this procedure, you need the Sun Cluster Agents CD-ROM.

If you are installing more than one data service simultaneously, perform the procedure in “Installing the Software” in *Sun Cluster Software Installation Guide for Solaris OS*.

Install these packages *only* in the global zone. To ensure that these packages are not propagated to any local zones that are created after you install the packages, use the `scinstall` utility to install these packages.

▼ How to Install the Sun Cluster HA for Solaris Containers Packages

Perform this procedure on all nodes that can run Sun Cluster HA for Solaris Containers.

- Steps**
1. **Load the Sun Cluster Agents CD-ROM into the CD-ROM drive.**
 2. **Run the `scinstall` utility with no options.**
The `scinstall` utility prompts you for additional information.
 3. **Chose the menu option, Add Support for New Data Service to this Cluster Node**
This step starts the `scinstall` utility in interactive mode.
 4. **Provide the pathname to the Sun Cluster Agents CD-ROM.**
The utility refers to the CD as “data services cd.”
 5. **Chose the menu option, q) done.**
 6. **Type yes for the question, Do you want to see more data services?**
The utility refers to the CD as “data services cd.”
 7. **Specify the data service to install.**
The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.
 8. **Exit the `scinstall` utility.**
 9. **Unload the CD from the CD-ROM drive.**
-

Registering and Configuring Sun Cluster HA for Solaris Containers

Before you perform this procedure, ensure that the Sun Cluster HA for Solaris Containers data service packages are installed.

Use the configuration and registration files in the following directories to register the Sun Cluster HA for Solaris Containers resources:

- `/opt/SUNWsczone/sczbt/util`
- `/opt/SUNWsczone/sczsh/util`
- `/opt/SUNWsczone/sczsmf/util`

The files define the dependencies that are required between the Sun Cluster HA for Solaris Containers components. For information about these dependencies, see [“Dependencies Between Sun Cluster HA for Solaris Containers Components”](#) on page 16

Registering and configuring Sun Cluster HA for Solaris Containers involves the tasks that are explained in the following sections:

1. "Specifying Configuration Parameters for the Zone Boot Resource" on page 22
2. "Writing a Zone Script" on page 25
3. "Specifying Configuration Parameters for the Zone Script Resource" on page 26
4. "Writing an SMF Service Probe" on page 28
5. "Specifying Configuration Parameters for the Zone SMF Resource" on page 29
6. "How to Create and Enable Resources for the Zone Boot Component" on page 30
7. "How to Create and Enable Resources for the Zone Script Component" on page 31
8. "How to Create and Enable Resources for the Zone SMF Component" on page 31

Specifying Configuration Parameters for the Zone Boot Resource

Sun Cluster HA for Solaris Containers provides a script that automates the process of configuring the zone boot resource. This script obtains configuration parameters from the `sczbt_config` file in the `/opt/SUNWsczone/sczbt/util` directory. To specify configuration parameters for the zone boot resource, edit the `sczbt_config` file.

Each configuration parameter in the `sczbt_config` file is defined as a keyword-value pair. The `sczbt_config` file already contains the required keywords and equals signs. For more information, see "Listing of `sczbt_config`" on page 39. When you edit the `sczbt_config` file, add the required value to each keyword.

The keyword-value pairs in the `sczbt_config` file are as follows:

```
RS=sczbt-rs
RG=sczbt-rg
PARAMETERDIR=sczbt-parameter-directory
SC_NETWORK=true|false
SC_LH=sczbt-lh-rs
FAILOVER=true|false
HAS_RS=sczbt-has-rs
Zonename=zone-name
Zonebootopt=zone-boot-options
Milestone=zone-boot-milestone
```

The meaning and permitted values of the keywords in the `sczbt_config` file are as follows:

`RS=sczbt-rs`

Specifies the name that you are assigning to the zone boot resource. You must specify a value for this keyword.

`RG=sczbt-rg`

Specifies the name of the resource group the zone boot resource will reside in. You must specify a value for this keyword.

`PARAMETERDIR=sczbt parameter directory`

Specifies the directory name that you are assigning to the parameter directory where some variables and their values will be stored. You must specify a value for this keyword.

`SC_NETWORK=true|false`

Specifies whether the zone boot resource is network aware with a `SUNW.LogicalHostName` resource. You must specify a value for this keyword.

- If HA for the zone's addresses is not required then configure the zone's addresses by using the `zonecfg` utility.

```
SC_NETWORK=false
SC_LH=
```

- If HA through IPMP protection is required then configure the zone's addresses by using the `zonecfg` utility and then place the zone's addresses on an adapter within an IPMP group.

```
SC_NETWORK=false
SC_LH=
```

- If HA through IPMP protection and protection against the failure of all physical interfaces is required, choose one option from the following list:

- If you require the `SUNW.LogicalHostName` resource type to manage one or a subset of the zone's addresses, configure a `SUNW.LogicalHostName` resource for those zone's addresses and not by using the `zonecfg` utility. Use the `zonecfg` utility to configure only the zones's addresses that are not to be under the control of the `SUNW.LogicalHostName` resource type.

```
SC_NETWORK=true
SC_LH=Name of the SUNW.LogicalHostName resource
```

- If you require the `SUNW.LogicalHostName` resource type to manage all the zone's addresses, configure a `SUNW.LogicalHostName` resource with a list of the zone's addresses and do not configure them by using the `zonecfg` utility.

```
SC_NETWORK=true
SC_LH=Name of the SUNW.LogicalHostName resources
```

- Otherwise, configure the zone's addresses by using the `zonecfg` utility and configure a separate redundant IP address for use by a `SUNW.LogicalHostName` resource, which must not be configured using the `zonecfg` utility.

```
SC_NETWORK=false
SC_LH=Name of the SUNW.LogicalHostName resource
```

`SC_LH=sczbt-ll-rs`

Specifies the name of the `SUNW.LogicalHostName` resource for the zone boot resource. Refer to [“Restrictions for Zone Network Addresses” on page 14](#) for a description of when to set this variable. This name must be the `SUNW.Logicalhostname` resource name you assigned when you created the resource in [Step 4](#).

`FAILOVER=true|false`

Specifies whether the zone's zone path is on a highly available file system.

`HAS_RS=sczbt-has-rs`

Specifies the name of the `SUNW.HAStoragePlus` resource for the zone boot resource. This name must be the `SUNW.HAStoragePlus` resource name you assigned when you created the resource in “[How to Enable a Zone to Run in a Failover Configuration](#)” on page 17. You must specify a value for this keyword if `FAILOVER=true` is set.

`Zonename=zone-name`

Specifies the zone name. You must specify a value for this keyword.

`Zonebootopt=zone-boot-options`

Specifies the zone boot option to use. Only `-s` is supported. Leaving this variable blank will cause the zone to boot to the `multi-user-server` milestone.

`Milestone=zone-boot-milestone`

Specifies the milestone the zone must reach to be considered as successfully booted. You must specify a value for this keyword.

EXAMPLE 1 Sample `sczbt_config` File

This example shows an `sczbt_config` file in which configuration parameters are set as follows:

- The name of the zone boot resource is `zone1-rs`.
- The name of the resource group for the zone boot resource is `zone1-rg`.
- The name of the parameter file directory for the zone boot resource is `/global/zones/pfiles`.
- Indicates that the zone’s address is managed by a `SUNW.LogicalHostName` resource and is `true`.
- The name of the `SUNW.LogicalHostName` resource name for the zone boot resource is `zone1-lh`.
- Indicates that the zone boot resource’s zone path is managed by a `SUNW.LogicalHostName` resource and is `true`.
- The name of the `SUNW.HAStoragePlus` resource name for the zone boot resource is `zone1-has`.
- The name of the zone is `zone1`.
- Indicates that the zone boot resource’s boot option is `null`.
- Indicates that the zone boot resource’s milestone is `multi-user-server`.

```
RS=zone1-rs
RG=zone1-rg
PARAMETERDIR=/global/zones/pfiles
SC_NETWORK=true
SC_LH=zone1-lh
FAILOVER=true
HAS_RS=zone1-has
Zonename=zone1
Zonebootopt=
Milestone=multi-user-server
```


Writing a Zone Script

The zone script resource provides the ability to run commands or scripts to start, stop and probe an application within a zone. The zone script resource depends on the zone boot resource. The command or script names are passed to the zone script resource when the resource is registered and must meet with the following requirements.

- The command or script must contain the fully qualified path within the zone.
- The command or script must be executable by root.
- The command or script must return one of the following return codes.

TABLE 3 Return codes

0	Successful completion
>0	An error has occurred
201	(Probe only) — An error has occurred that requires an immediate failover of the resource group
>0 & !=201	(Probe only) — An error has occurred that requires a resource restart

Note – For an immediate failover of the zone script resource, you must configure the resource properties `Failover_mode` and `Failover_enabled` to meet the required behavior. Refer to the `r_properties(5)` man page when setting the `Failover_mode` property and `SUNW.gds(5)` man page when setting the `Failover_enabled` property.

EXAMPLE 2 Zone Probe Script for Apache2

This example shows a simple script to test that the Apache2 service is running, beyond the process tree existing. The script `/var/tmp/probe-apache2` must exist within the zone.

```
# cat /var/tmp/probe-apache2
#!/usr/bin/ksh
if "echo GET; exit" | mconnect -p 80
then
    exit 0
else
    exit 100
fi
```

Specifying Configuration Parameters for the Zone Script Resource

Sun Cluster HA for Solaris Containers provides a script that automates the process of configuring zone script resource. This script obtains configuration parameters from the `sczsh_config` file in the `/opt/SUNWsczone/sczsh/util` directory. To specify configuration parameters for the zone script resource, edit the `sczsh_config` file.

Each configuration parameter in the `sczsh_config` file is defined as a keyword-value pair. The `sczsh_config` file already contains the required keywords and equals signs. For more information, see [“Listing of `sczsh_config`” on page 41](#). When you edit the `sczsh_config` file, add the required value to each keyword.

The keyword-value pairs in the `sczsh_config` file are as follows:

```
RS=sczsh-rs
RG=sczbt-rg
SCZBT_RS=sczbt-rs
PARAMETERDIR=sczsh-parameter-directory
Zonename=sczbt-zone-name
ServiceStartCommand=sczsh-start-command
ServiceStopCommand=sczsh-stop-command
ServiceProbeCommand=sczsh-probe-command
```

The meaning and permitted values of the keywords in the `sczsh_config` file are as follows:

`RS=sczsh-rs`

Specifies the name that you are assigning to the zone script resource. You must specify a value for this keyword.

`RG=sczbt-rg`

Specifies the name of the resource group the zone boot resource resides in. You must specify a value for this keyword.

`SCZBT_RS=sczbt-rs`

Specifies the name of the zone boot resource. You must specify a value for this keyword.

`PARAMETERDIR=sczsh parameter directory`

Specifies the directory name that you are assigning to the parameter directory where the following variables and their values will be stored. You must specify a value for this keyword.

`Zonename=sczbt-zone-name`

Specifies the zone name. You must specify a value for this keyword.

`ServiceStartCommand=sczsh-start-command`

Specifies the zone start command or script to run. You must specify a value for this keyword.

`ServiceStopCommand=sczsh-stop-command`

Specifies the zone stop command or script to run. You must specify a value for this keyword

`ServiceProbeCommand=sczsh-probe-command`

Specifies the zone probe command or script to run. You must specify a value for this keyword

EXAMPLE 3 Sample `sczsh_config` File

In this example the zone script resource uses the Apache2 scripts that are available in Solaris 10. Before this example can be used the Apache2 configuration file `http.conf` needs to be configured. For the purpose of this example, the delivered `http.conf-example` can be used. Copy the file as follows:

```
# zlogin zone1
# cd /etc/apache2
# cp http.conf-example http.conf
# exit
```

This example shows an `sczsh_config` file in which configuration parameters are set as follows:

- The name of the zone script resource is `zone1-script-rs`.
- The name of the resource group for the zone script resource is `zone1-rg`.
- The name of the zone boot resource is `zone1-rs`.
- The name of the parameter file directory for the zone script resource is `/global/zones/pfiles`.
- The name of the zone is `zone1`.
- The name of the zone script resource start command and its parameter is `"/lib/svc/method/http-apache2 start"`.
- The name of the zone script resource stop command and its parameter is `"/lib/svc/method/http-apache2 stop"`.
- The name of the zone script resource probe command is `"/var/tmp/probe-apache2"`. This script is shown in [Example 2](#) and must exist in `zone1`.

```
RS="zone1-script-rs"
RG="zone1-rg"
SCZBT_RS="zone1-rs"
PARAMETERDIR="/global/zones/pfiles"
Zonename="zone1"
ServiceStartCommand="/lib/svc/method/http-apache2 start"
ServiceStopCommand="/lib/svc/method/http-apache2 stop"
ServiceProbeCommand="/var/tmp/probe-apache2"
```

Writing an SMF Service Probe

The zone SMF resource provides the ability to enable, disable and probe a SMF service within a zone. The zone SMF resource depends on the zone boot resource. Probing the SMF service is performed by running a command or script against the SMF service. The SMF service and probe command or script names are passed to the zone SMF resource when the resource is registered. The probe command or script must meet to the following requirements.

- The probe command or script must contain the fully qualified path within the zone.
- The probe command or script must be executable by root.
- The probe command or script must return one of the following return codes.

TABLE 4 Return codes

0	Successful completion
100	An error occurred that requires a resource restart
201	An error has occurred that requires an immediate failover of the resource group

Note – For an immediate failover of the zone SMF resource, you must configure the resource properties `Failover_mode` and `Failover_enabled` to meet the required behavior. Refer to the `r_properties(5)` man page when setting the `Failover_mode` property and `SUNW.gds(5)` man page when setting the `Failover_enabled` property.

EXAMPLE 4 Zone SMF Probe Script for Apache2

This example shows a simple script to test that the SMF Apache2 service is running, beyond the process tree existing. The script `/var/tmp/probe-apache2` must exist within the zone.

```
# cat /var/tmp/probe-apache2
#!/usr/bin/ksh
if "echo GET; exit" | mconnect -p 80
then
    exit 0
else
    exit 100
fi
```

Specifying Configuration Parameters for the Zone SMF Resource

Sun Cluster HA for Solaris Containers provides a script that automates the process of configuring the zone SMF resource. This script obtains configuration parameters from the `sczsmf_config` file in the `/opt/SUNWsczone/sczsmf/util` directory. To specify configuration parameters for the zone SMF resource, edit the `sczsmf_config` file.

Each configuration parameter in the `sczsmf_config` file is defined as a keyword-value pair. The `sczsmf_config` file already contains the required keywords and equals signs. For more information, see [“Listing of `sczsmf_config`” on page 42](#). When you edit the `sczsmf_config` file, add the required value to each keyword.

The keyword-value pairs in the `sczsmf_config` file are as follows:

```
RS=sczsmf-rs
RG=sczbt-rg
SCZBT_RS=sczbt-rs
ZONE=sczbt-zone-name
SERVICE=smf-service
RECURSIVE=true|false
STATE=true|false
SERVICE_PROBE=sczsmf-service-probe
```

The meaning and permitted values of the keywords in the `sczsmf_config` file are as follows:

`RS=sczsmf-rs`

Specifies the name that you are assigning to the zone SMF resource. This must be defined.

`RG=sczbt-rg`

Specifies the name of the resource group the zone boot resource resides in. This must be defined.

`SCZBT_RS=sczbt-rs`

Specifies the name of the zone boot resource. You must specify a value for this keyword.

`ZONE=sczbt-zone-name`

Specifies the zone name. This must be defined.

`SERVICE=smf-service`

Specifies the SMF service to enable/disable. This must be defined.

`RECURSIVE=true|false`

Specifies `true` to enable the service recursively or `false` to just enable the service and no dependents. This must be defined.

`STATE=true|false`

Specifies `true` to wait until the service state is reached or `false` to not wait until the service state is reached. This must be defined.

```
SERVICE_PROBE=sczsmf-service-probe
```

Specify the script to check the SMF service.

EXAMPLE 5 Sample `sczsmf_config` File

In this example the zone SMF resource uses the Apache2 SMF service that is available in Solaris 10. Before this example can be used the Apache2 configuration file `http.conf` needs to be configured. For the purpose of this example, the delivered `http.conf-example` can be used. Copy the file as follows:

```
# zlogin zone1
# cd /etc/apache2
# cp http.conf-example http.conf
# exit
```

This example shows an `sczsmf_config` file in which configuration parameters are set as follows:

- The name of the zone SMF resource is `zone1-smf-rs`.
- The name of the resource group for the zone SMF resource is `zone1-rg`.
- The name of the zone boot resource is `zone1-rs`.
- The name of the zone name is `zone1`.
- The name of the zone SMF service is `apache2`.
- Indicates that the zone SMF service Recursive option is `true`.
- Indicates that the zone SMF service State option is `true`.
- Indicates that the zone SMF service probe name is `/var/tmp/probe-apache2`.
This script is shown in [Example 4](#) and must exist in `zone1`.

```
RS=zone1-smf-rs
RG=zone1-rg
SCZBT_RS=zone1-rs
ZONE=zone1
SERVICE=apache2
RECURSIVE=true
STATE=true
SERVICE_PROBE=/var/tmp/probe-apache2
```

▼ How to Create and Enable Resources for the Zone Boot Component

Before You Begin Ensure you have edited the `sczbt_config` file to specify configuration parameters for the Sun Cluster HA for Solaris Containers zone boot component. For more information, see [“Specifying Configuration Parameters for the Zone Boot Resource”](#) on page 22.

- Steps** 1. Become superuser on one of the nodes in the cluster that will host the zone.

2. Register the `SUNW.gds` resource type.

```
# scrgadm -a -t SUNW.gds
```

3. Go to the directory that contains the script for creating the Sun Cluster HA for Solaris Containers boot resource.

```
# cd /opt/SUNWsczone/sczbt/util
```

4. Run the script that creates the zone boot resource.

```
# ./sczbt_register
```

5. Bring online the zone boot resource.

```
# scswitch -e -j sczbt-rs
```

▼ How to Create and Enable Resources for the Zone Script Component

Before You Begin

Ensure you have edited the `sczsh_config` file to specify configuration parameters for the Sun Cluster HA for Solaris Containers zone script component. For more information, see [“Specifying Configuration Parameters for the Zone Script Resource” on page 26](#).

Steps

1. Go to the directory that contains the script for creating the Sun Cluster HA for Solaris Containers script resource.

```
# cd /opt/SUNWsczone/sczsh/util
```

2. Run the script that creates the zone script resource.

```
# ./sczsh_register
```

3. Bring online the zone script resource.

```
# scswitch -e -j sczsh-rs
```

▼ How to Create and Enable Resources for the Zone SMF Component

Before You Begin

Ensure you have edited the `sczsmf_config` file to specify configuration parameters for the Sun Cluster HA for Solaris Containers zone SMF component. For more information, see [“Specifying Configuration Parameters for the Zone SMF Resource” on page 29](#).

- Steps**
1. Go to the directory that contains the script for creating the Sun Cluster HA for Solaris Containers SMF resource.

```
# cd /opt/SUNWsczone/sczsmf/util
```

2. Run the script that creates the zone SMF resource.

```
# ./sczsmf_register
```

3. Bring online the zone SMF resource.

```
# scswitch -e -j sczsmf-rs
```

Verifying the Sun Cluster HA for Solaris Containers and Configuration

After you install, register, and configure Sun Cluster HA for Solaris Containers, verify the Sun Cluster HA for Solaris Containers installation and configuration. Verifying the Sun Cluster HA for Solaris Containers installation and configuration determines if the Sun Cluster HA for Solaris Containers data service makes your zones highly available.

▼ How to Verify the Sun Cluster HA for Solaris Containers Installation and Configuration

- Steps**
1. Become superuser on a cluster node that is to host the Solaris Zones component.

2. Ensure all the Solaris Zone resources are online.

For each resource, perform the following steps.

- a. Determine whether the resource is online.

```
# scstat -g
```

- b. If the resource is not online, bring online the resource.

```
# scswitch -e -j solaris-zone-resource
```

3. Switch the zone resource group to another cluster node, such as *node2*

```
# scswitch -z -g solaris-zone-resource-group -h node2
```

4. Confirm that the resource is now online on *node2*.

```
# scstat -g
```

Patching the Global Zone and Local Zones

The procedures that follow are required only if you are applying the patch to the global zone *and* to local zones. If you are applying a patch to only the global zone, follow the instructions in Chapter 8, “Patching Sun Cluster Software and Firmware,” in *Sun Cluster System Administration Guide for Solaris OS*.

Before you begin, consult the patch README file to determine whether the patch is a nonrebooting patch or a rebooting patch.

▼ How to Apply a Non Rebooting Patch to the Global Zone and Local Zones

A nonrebooting patch does not require you to reboot a node after you apply the patch on the node. You can apply the patch to a live system.

- Steps**
1. From one node, disable monitoring of every resource in the resource group that contains the zone resource.

```
# scswitch -n -M -j resource-list
```

2. On each node where the zone is *not* booted, comment out the entry for the zone in the `/etc/zones/index` file.

To comment out an entry, add the # character to the start of the line that contains the entry.

3. Apply the patch on all nodes where the zone is configured.
4. Remove the comment from each entry that you edited in [Step 2](#).
5. Enable monitoring of the resources for which you disabled monitoring in [Step 1](#).

```
# scswitch -e -M -j resource-list
```

▼ How to Apply a Rebooting Patch to the Global Zone and Local Zones

A rebooting patches requires you to reboot a node after you apply the patch to the node.

- Steps**
1. Disable the resources that depend on the zones to which you are applying the patch.

```
# scswitch -n -j zdepend-rs-list
```

2. Disable monitoring of the zone resource.

```
# scswitch -n -M -j zone-rs
```

3. Bring the resource groups that contain zone resources online on a node.

```
# scswitch -z -g zone-rg -h node
```

4. On each node where the zone is *not* booted, comment out the entry for the zone in the `/etc/zones/index` file.

To comment out an entry, add the # character to the start of the line that contains the entry.

5. For each node where the zone is *not* booted, perform the following sequence of operations:

- a. Apply the patch.

- b. Reboot the node.

6. Apply the patch on the node where the zone is booted.

7. Remove the comment from each entry that you edited in [Step 4](#).

8. Enable monitoring of the resource for which you disabled monitoring in [Step 2](#).

```
# scswitch -e -M -j zone-rs
```

9. Reboot the node where the zone is booted.

10. Enable the resources that you disabled in [Step 1](#).

```
# scswitch -e -j zdepend-rs-list
```

Next Steps To verify that the patch is correctly applied, switch each resource group that contains zone resources to each node in the resource group's node list. To switch a resource group to another node, type the command:

```
scswitch -z -g zone-rg -h node
```

Tuning the Sun Cluster HA for Solaris Containers Fault Monitors

The Sun Cluster HA for Solaris Containers fault monitors verify that the following components are running correctly:

- Zone boot resource
- Zone script resource
- Zone SMF resource

Each Sun Cluster HA for Solaris Containers fault monitor is contained in the resource that represents Solaris Zones component. You create these resources when you register and configure Sun Cluster HA for Solaris Containers. For more information, see [“Registering and Configuring Sun Cluster HA for Solaris Containers”](#) on page 21.

System properties and extension properties of these resources 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 Sun Cluster installations. Therefore, you should tune the Sun Cluster HA for Solaris Containers fault monitor *only* if you need to modify this preset behavior.

Tuning the Sun Cluster HA for Solaris Containers fault monitors involves the following tasks:

- Setting the interval between fault monitor probes
- Setting the time-out for fault monitor probes
- Defining the criteria for persistent faults
- Specifying the failover behavior of a resource

For more information, see *“Tuning Fault Monitors for Sun Cluster Data Services”* in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Operation of the Sun Cluster HA for Solaris Containers Parameter File

The Sun Cluster HA for Solaris Containers zone boot and script resources uses a parameter file to pass parameters to the start, stop and probe commands. Changes to these parameters take effect at every restart or enabling, disabling of the resource.

Operation of the Fault Monitor for the Zone Boot Component

The fault monitor for the zone boot component ensures that the all requirements for the zone boot component to run are met:

- The Sun Cluster HA for Solaris Containers `zsched` process is running.
If this process is not running, the fault monitor restarts the zone. If this fault persists, the fault monitor fails over the resource group that contains resource for the zone boot component.
- Every host that is managed by a `SUNW.LogicalHostname` resource is operational.
If the host is not operational, the fault monitor fails over the resource group that contains resource for the zone boot component.
- The specified milestone is either online or degraded
If the milestone is not online or degraded, the fault monitor restarts the zone. If this fault persists, the fault monitor fails over the resource group that contains resource for the zone boot component.

To verify the state of the milestone, the fault monitor connects to the zone. If the fault monitor cannot connect to the zone, the fault monitor retries every five seconds for approximately 60% of the probe time-out. If the attempt to connect still fails, then the fault monitor restarts the zone.

Operation of the Fault Monitor for the Zone Script Component

The fault monitor for the zone script component runs the script that you specify for the component. The value that this script returns to the fault monitor determines the action that the fault monitor performs. For more information, see [Table 3](#).

Operation of the Fault Monitor for the Zone SMF Component

The fault monitor for the zone SMF component verifies that the SMF service is not disabled. If the service is disabled, the fault monitor restarts the SMF service. If this fault persists, the fault monitor fails over the resource group that contains resource for the zone SMF component.

If the service is not disabled, the fault monitor runs the SMF service probe that you specify for the component. The value that this probe returns to the fault monitor determines the action that the fault monitor performs. For more information, see [Table 4](#).

Debugging Sun Cluster HA for Solaris Containers

The `config` file in the `/opt/SUNWsczone/xxx/etc` directory enables you to activate debugging for Solaris Zone resources. Where `xxx` represents `sczbt` for the boot component, `sczsh` for the script component and `sczsmf` for the SMF component.

Each component of Sun Cluster HA for Solaris Containers has a `config` that enables you to activate debugging for Solaris Zone resources. The location of this file for each component is as follows:

- For the zone boot component, this file is contained in the `/opt/SUNWsczone/sczbt/etc` directory.
- For the zone script component, this file is contained in the `/opt/SUNWsczone/sczsh/etc` directory.
- For the zone SMF component, this file is contained in the `/opt/SUNWsczone/sczsmf/etc` directory.

▼ How to Activate Debugging for Sun Cluster HA for Solaris Containers

- Steps**
1. **Determine whether debugging for Sun Cluster HA for Solaris Containers is active.**

If debugging is inactive, `daemon.notice` is set in the file `/etc/syslog.conf`.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                  operator
#
```

2. **If debugging is inactive, edit the `/etc/syslog.conf` file to change `daemon.notice` to `daemon.debug`.**

3. **Confirm that debugging for Sun Cluster HA for Solaris Containers is active.**

If debugging is active, `daemon.debug` is set in the file `/etc/syslog.conf`.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                  operator
#
```

4. **Restart the `syslogd` daemon.**

```
# pkill -1 syslogd
```

5. Edit the `/opt/SUNWsczone/sczbt/etc/config` file to change `DEBUG=` to `DEBUG=ALL` or `DEBUG=sczbt-rs`.

```
# cat /opt/SUNWsczone/sczbt/etc/config
#
# Copyright 2005 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# Usage:
#     DEBUG=<RESOURCE_NAME> or ALL
#
DEBUG=ALL
#
```

Note – To deactivate debugging, reverse the preceding steps.

Files for Configuring Sun Cluster HA for Solaris Zones Resources

The `/opt/SUNWsczone/util` directory contains files that automate the process of configuring Sun Cluster HA for Solaris Zones resources. Listings of these files are provided in the following sections:

- “Listing of `sczbt_config`” on page 39
- “Listing of `sczsh_config`” on page 41
- “Listing of `sczsmf_config`” on page 42

Listing of `sczbt_config`

```
#
# Copyright 2005 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# This file will be sourced in by sczbt_register and the parameters
# listed below will be used.
#
# These parameters can be customized in (key=value) form
#
#           RS - Name of the resource
#           RG - Name of the resource group containing RS
# PARAMETERDIR - Name of the parameter file directory
#           SC_NETWORK - Identifies if SUNW.LogicalHostname will be used
#                       true = zone will use SUNW.LogicalHostname
#                       false = zone will use it's own configuration
#
#           The configuration of a zone's network addresses depends on
#           whether you require IPMP protection or protection against
#           the failure of all physical interfaces.
#
#           If you require only IPMP protection, configure the zone's
#           addresses by using the zonecfg utility and then place the
```

```

#           zone's address in an IPMP group.
#
#           To configure this option set
#           SC_NETWORK=false
#           SC_LH=
#
# If IPMP protection is not required, just configure the
# zone's addresses by using the zonecfg utility.
#
#           To configure this option set
#           SC_NETWORK=false
#           SC_LH=
#
# If you require protection against the failure of all physical
# interfaces, choose one option from the following list.
#
# - If you want the SUNW.LogicalHostName resource type to manage
#   the zone's addresses, configure a SUNW.LogicalHostName
#   resource with at least one of the zone's addresses.
#
#           To configure this option set
#           SC_NETWORK=true
#           SC_LH=<Name of the SC Logical Hostname resource>
#
# - Otherwise, configure the zone's addresses by using the
#   zonecfg utility and configure a redundant IP address
#   for use by a SUNW.LogicalHostName resource.
#
#           To configure this option set
#           SC_NETWORK=false
#           SC_LH=<Name of the SC Logical Hostname resource>
#
# Whichever option is chosen, multiple zone addresses can be
# used either in the zone's configuration or using several
# SUNW.LogicalHostname resources.
#
#   e.g. SC_NETWORK=true
#        SC_LH=zone1-lh1,zone1-lh2
#
#       SC_LH - Name of the SC Logical Hostname resource
#       FAILOVER - Identifies if the zone's zone path is on a
#                  highly available local file system
#
#   e.g. FAILOVER=true - highly available local file system
#        FAILOVER=false - local file system
#
#       HAS_RS - Name of the HASStoragePlus SC resource
#
RS=
RG=
PARAMETERDIR=
SC_NETWORK=
SC_LH=
FAILOVER=

```



```

HAS_RS=

#
# The following variable will be placed in the parameter file
#
# Parameters for sczbt (Zone Boot)
#
# Zonename      Name of the zone
# Zonebootopt   Zone boot options
# Milestone     SMF Milestone which needs to be online
#               before the zone is considered booted
#
# Note: Zonebootopt="-s" requires Milestone=single-user
#

Zonename=
Zonebootopt=
Milestone=multi-user-server

```

Listing of sczsh_config

```

#
# Copyright 2005 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# This file will be sourced by sczsh_register and the parameters
# listed below will be used.
#
# These parameters can be customized in (key=value) form
#
#           RS - Name of the resource
#           RG - Name of the resource group containing RS
#           SCZBT_RS - Name of the SC Zone boot resource
#           PARAMETERDIR - Name of the parameter file directory
#           Zonename - Name of the zone
# ServiceStartCommand - Command including all options to start
#                       the service in the configured zone
# ServiceStopCommand - Command including all options to stop
#                       the service in the configured zone
# ServiceProbeCommand - Command including all options to probe
#                       the service in the configured zone
#
RS=""
RG=""
SCZBT_RS=""
PARAMETERDIR=""
#
# The following parameters will be put in the agents parameterfile:
#

```

```
Zonename=""
ServiceStartCommand=""
ServiceStopCommand=""
ServiceProbeCommand=""
```

Listing of sczsmf_config

```
#
# Copyright 2005 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# This file will be sourced in by sczsmf_register and the parameters
# listed below will be used.
#
# These parameters can be customized in (key=value) form
#
#         RS - Name of the resource
#         RG - Name of the resource group containing RS
#     SCZBT_RS - Name of the SC Zone boot resource
#         ZONE - Name of the Zone
#
# For SERVICE, RECURSIVE and STATE, refer to the svcadm(1M)
#         man page
#
#     SERVICE - {FMRI | pattern}
#         FMRI - Fault management resource identifier
#         pattern - Pattern matching a service
#
#     RECURSIVE - {false | true}    Default: true
#         False - Just enable the service and no dependents
#         True - Enable the service and recursively enable
#             its dependents
#
#         RECURSIVE=true equates to svcadm enable "-r"
#
#     STATE - {false | true}    Default: true
#         False - Do not wait until service state is reached
#         True - Wait until service state is reached
#
#         STATE=true equates to svcadm enable/disable "-s"
#
#     SERVICE_PROBE - Script to check the SMF service
#
#         The optional parameter, SERVICE_PROBE, provides the
#         ability to check that the SMF service is working.
#         This must be a script within the zone and must
#         adhere to these return codes,
#
#         0 - The SMF service is working
```

```
#          100 - The SMF service should be restarted
#          201 - The SMF service should initiate a failover of
#                Resource Group
#
#          Note: That return code 201, requires that this resource
#                has an appropriate extension property value for
#                FAILOVER_MODE and FAILOVER_ENABLED=TRUE
#
#          For FAILOVER_MODE refer to the r_properties(5) man page.
#

RS=
RG=
SCZBT_RS=
ZONE=
SERVICE=
RECURSIVE=true
STATE=true
SERVICE_PROBE=
```


Index

C

commands
 node information, 9
 scinstall, 21
 scrgadm, 17, 18, 31
 scswitch, 17, 18, 31, 32
 svcs, 20
 zlogin, 18, 20
 zoneadm, 20
configuration files, 21
 zone boot resource, 22
 zone script resource, 22
 zone SMF resource, 22

D

debugging, Sun Cluster HA for Solaris
 Containers, 37-38
dependencies, 16

E

enable resource group, 17, 18
extension properties, effect on fault
 monitors, 35

F

fault monitors, tuning, 35-36
files
 config, 37
 /opt/SUNWsczone/etc/config, 37
 sczbt_config
 listing, 39-41
 sczsh_config
 listing, 41-42
 sczsmf_config, 42-43

O

/opt/SUNWsczone/sczbt/etc/config
 file, 37
overview, 11-12

P

parameter file directory, 16
patching
 non rebooting patch, 33
 rebooting patch, 33-34
prtconf -v command, 9
prtdiag -v command, 9
psrinfo -v command, 9

R

register

- sczbt_register, 31
- sczsh_register, 31
- sczsmf_register, 32

requirements, 15-16

resource types, fault monitors, 35-36

resources

- Solaris Zones application
debugging, 37-38

restrictions, 14-15

return codes

- zone script resource, 25
- zone SMF resource, 28

S

scinstall -pv command, 9

sczbt_config file, listing, 39-41

sczsh_config file, listing, 41-42

sczsmf_config file, 42-43

showrev -p command, 9

Solaris Zones application, fault monitor, 35-36

Sun Cluster HA for Solaris Containers

- debugging, 37-38

- fault monitors, 35-36

system properties, effect on fault monitors, 35

T

training, 8

tuning, fault monitors, 35-36