

Oracle® Solaris Cluster Data Service for Solaris Containers Guide

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.

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.

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 END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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.

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.

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 END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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.

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. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Preface	5
1 Installing and Configuring HA for Solaris Containers	9
HA for Solaris Containers Overview	9
Overview of Installing and Configuring HA for Solaris Containers	11
Planning the HA for Solaris Containers Installation and Configuration	11
Configuration Restrictions	12
Configuration Requirements	14
Installing and Configuring Zones	15
▼ How to Enable a Zone to Run in Failover Configuration	15
▼ How to Enable a Zone to Run in a Multiple-Masters Configuration	16
▼ How to Install a Zone and Perform the Initial Internal Zone Configuration	16
Verifying the Installation and Configuration of a Zone	18
▼ How to Verify the Installation and Configuration of a Zone	18
Installing the HA for Solaris Containers Packages	19
▼ How to Install the HA for Solaris Containers Packages	19
Registering and Configuring HA for Solaris Containers	21
Specifying Configuration Parameters for the Zone Boot Resource	21
Writing Scripts for the Zone Script Resource	26
Specifying Configuration Parameters for the Zone Script Resource	27
Writing a Service Probe for the Zone SMF Resource	29
Specifying Configuration Parameters for the Zone SMF Resource	30
▼ How to Create and Enable Resources for the Zone Boot Component	32
▼ How to Create and Enable Resources for the Zone Script Component	32
▼ How to Create and Enable Resources for the Zone SMF Component	33
Verifying the HA for Solaris Containers and Configuration	33
▼ How to Verify the HA for Solaris Containers Installation and Configuration	34
Patching the Global Zone and Non-Global Zones	34

▼ How to Patch to the Global Zone and Non-Global Zones	34
Tuning the HA for Solaris Containers Fault Monitors	35
Operation of the HA for Solaris Containers Parameter File	36
Operation of the Fault Monitor for the Zone Boot Component	36
Operation of the Fault Monitor for the Zone Script Component	36
Operation of the Fault Monitor for the Zone SMF Component	37
Tuning the HA for Solaris Containers Stop_timeout property	37
Choosing the Stop_timeout value for the Zone Boot Component	37
Choosing the Stop_timeout value for the Zone Script Component	37
Choosing the Stop_timeout value for the Zone SMF Component	38
Denying Cluster Services for a Non-Global Zone	38
Debugging HA for Solaris Containers	38
▼ How to Activate Debugging for HA for Solaris Containers	39
 A Files for Configuring HA for Solaris Containers Resources	 41
Listing of sczbt_config	41
Listing of sczsh_config	46
Listing of sczsmf_config	47
 Index	 49

Preface

Oracle Solaris Cluster Data Service for Solaris Containers 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	Description	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 UNIX system prompts and superuser prompts for shells that are included in the Oracle Solaris OS. In command examples, the shell prompt indicates whether the command should be executed by a regular user or a user with privileges.

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	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>

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#sys_sw.

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

For a complete list of Oracle Solaris Cluster documentation, see the release notes for your version of Oracle Solaris Cluster software.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

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 operating environment (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 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 -v</code>	Displays Oracle Solaris Cluster release and package version information for each node

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

Installing and Configuring HA for Solaris Containers

This chapter explains how to install and configure Oracle Solaris Cluster HA for Solaris Containers (HA for Solaris Containers).

This chapter contains the following sections.

- “HA for Solaris Containers Overview” on page 9
- “Overview of Installing and Configuring HA for Solaris Containers” on page 11
- “Planning the HA for Solaris Containers Installation and Configuration” on page 11
- “Installing and Configuring Zones” on page 15
- “Verifying the Installation and Configuration of a Zone” on page 18
- “Installing the HA for Solaris Containers Packages” on page 19
- “Registering and Configuring HA for Solaris Containers” on page 21
- “Verifying the HA for Solaris Containers and Configuration” on page 33
- “Patching the Global Zone and Non-Global Zones” on page 34
- “Tuning the HA for Solaris Containers Fault Monitors” on page 35
- “Tuning the HA for Solaris Containers `Stop_timeout` property” on page 37
- “Denying Cluster Services for a Non-Global Zone” on page 38
- “Debugging HA for Solaris Containers” on page 38

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 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 administrator of the global zone.

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

- The orderly booting, shutdown and fault monitoring of a zone through the `sczbt` component.
- The orderly startup, shutdown and fault monitoring of an application within the zone, using scripts or commands through the `sczsh` component.
- The orderly startup, shutdown and fault monitoring of a Solaris Service Management Facility (SMF) service within the zone through the `sczsmf` component.

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

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

You can also choose to set up non-global zones that do not participate in the cluster. A root user logged into one of these zones is not able to discover or disrupt operation of the cluster. See [“Denying Cluster Services for a Non-Global Zone” on page 38](#) for more information. For instructions, see [“How to Deny Cluster Services For a Non-Global Zone” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*](#) and [“How to Allow Cluster Services For a Non-Global Zone” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*](#)

For conceptual information about failover data services, multiple-masters data services, and scalable data services, see [Oracle Solaris Cluster Concepts Guide](#).

Overview of Installing and Configuring HA for Solaris Containers

The following table summarizes the tasks for installing and configuring 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 1-1 Tasks for Installing and Configuring HA for Solaris Containers

Task	Instructions
Plan the installation	“Planning the HA for Solaris Containers Installation and Configuration” on page 11
Install and configure the Solaris Zones	“Installing and Configuring Zones” on page 15
Verify installation and configuration	“How to Verify the Installation and Configuration of a Zone” on page 18
Install HA for Solaris Containers Packages	“Installing the HA for Solaris Containers Packages” on page 19
Register and configure HA for Solaris Containers components	“Registering and Configuring HA for Solaris Containers” on page 21
Verify HA for Solaris Containers installation and configuration	“Verifying the HA for Solaris Containers and Configuration” on page 33
Applying Patches to the global and non-global zones	“Patching the Global Zone and Non-Global Zones” on page 34
Tune the HA for Solaris Containers fault monitors	“Tuning the HA for Solaris Containers Fault Monitors” on page 35
Tune the HA for Solaris Containers <code>Stop_timeout</code> property	“Tuning the HA for Solaris Containers <code>Stop_timeout</code> property” on page 37
Debug HA for Solaris Containers	“Debugging HA for Solaris Containers” on page 38

Planning the HA for Solaris Containers Installation and Configuration

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

Configuration Restrictions

The configuration restrictions in the subsections that follow apply only to 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 (HA) you require. You can choose between no HA, HA through the use of only 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 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 only HA through IPMP protection is required, then configure the zone's addresses by using the `zonecfg` utility and 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 by triggering a failover 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 zone's addresses that are not required to be under the control 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.
 - 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.

If `ip-type=exclusive` option is set with `zonecfg` in the zone configuration for the configured `sczbt` resource, the `SC_NETWORK` variable in the `sczbt_config` file must be set to `false` to successfully register the `sczbt` resource. If `ip-type=exclusive` option is set for the non-global zone, do not configure a resource dependency on the `SUNW.LogicalHostname` resource from the `sczbt` resource.

Restrictions for an HA Container

The zone path of a zone in an HA container 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.

Ensure that the zone's autoboot property is set to `false`. Setting a zone's autoboot property to `false` prevents the zone from being booted when the global zone is booted. The HA for Solaris Containers data service can manage a zone only if the zone is booted under the control of the data service.

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.

Ensure that the zone's autoboot property is set to `false`. Setting a zone's autoboot property to `false` prevents the zone from being booted when the global zone is booted. The HA for Solaris Containers data service can manage a zone only if the zone is booted under the control of the data service.

Restrictions for the Zone Path of a Zone

The zone path of a zone that 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, Solaris 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 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 HA for Solaris Containers Components

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

TABLE 1-2 Dependencies Between HA for Solaris Containers Components

Component	Dependency
Zone boot resource (sczbt)	SUNW.HAS <code>storagePlus</code> - In a failover configuration, the zone's zone path must be on a highly available file system managed by a SUNW.HAS <code>storagePlus</code> resource SUNW.Logical <code>HostName</code> - This dependency is required only if the zone's address is managed by a SUNW.Logical <code>HostName</code> resource
Zone script resource (sczsh)	Zone boot resource
Zone SMF resource (sczsmf)	Zone boot resource

These dependencies are set when you register and configure HA for Solaris Containers. For more information, see [“Registering and Configuring 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 HA for Solaris Containers

The boot component and script component of 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.

Note – If a multiple-masters zone configuration is being deployed, you must ensure that the parameter directory exists on all nodes that will host the zone. You can either place separate copies of the parameter directory on local storage on each node or place the parameter directory on a cluster file system. The advantage of using a cluster file system ensures that only one copy of the parameter directory exists, whereas if you have separate copies of the parameter directory on local storage you must ensure the contents are kept up to date after registering an `sczbt` or `sczsh` resource and before you enable the resource.

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 Failover Configuration” on page 15](#)
 - [“How to Enable a Zone to Run in a Multiple-Masters Configuration” on page 16](#)
2. Installing and configuring a zone, as explained in:
 - [“How to Install a Zone and Perform the Initial Internal Zone Configuration” on page 16](#)

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 HA for Solaris Containers. For complete information about installing and configuring Solaris Zones, see *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*.

▼ How to Enable a Zone to Run in Failover Configuration

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

```
# clresourcetype register SUNW.HAStoragePlus
```

- 2 Create a failover resource group.

```
# clresourcegroup create solaris-zone-resource-group
```

- 3 Create a resource for the zone's disk storage.

This `HAStoragePlus` resource is for the `zonepath`. The file system must be a failover file system.

```
# clresource create \
-g solaris-zone-resource-group \
-t SUNW.HAStoragePlus \
-p FilesystemMountPoints=solaris-zone-instance-mount-points \
solaris-zone-has-resource-name
```

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

```
# clreslogicalhostname create \  
-g solaris-zone-resource-group \  
-h solaris-zone-logical-hostname \  
solaris-zone-logical-hostname-resource-name
```

5 Enable the failover resource group.

```
# clresourcegroup online -M solaris-zone-resource-group
```

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

1 Create a scalable resource group.

```
# clresourcegroup create \  
-p Maximum_primaries=max-number \  
-p Desired_primaries=desired-number \  
solaris-zone-resource-group
```

2 Enable the scalable resource group.

```
# clresourcegroup online -M solaris-zone-resource-group
```

▼ How to Install a Zone and Perform the Initial Internal Zone Configuration

Perform this task on each node that is to host the zone.

Note – For complete information about installing a zone, see [System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones](#).

Before You Begin

Consult “[Configuration Restrictions](#)” on [page 12](#) and then determine the following requirements for the deployment of the zone with Solaris Cluster:

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

Ensure that the zone is configured.

If the zone that you are installing is to run in a failover configuration, configure the zone's zone path to specify a highly available local file system. The file system must be managed by the SUNW.HAStoragePlus resource that you created in “[How to Enable a Zone to Run in Failover Configuration](#)” on [page 15](#).

For detailed information about configuring a zone before installation of the zone, see the following documentation:

- Chapter 17, “Non-Global Zone Configuration (Overview),” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*
- Chapter 18, “Planning and Configuring Non-Global Zones (Tasks),” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*

1 If the zone is to run in a failover configuration, ensure that the zone's zone path can be created on the zone's disk storage.

If the zone is to run in a multiple-masters configuration, omit this step.

- a. On the node where you are installing the zone, bring online the resource group that contains the resource for the zone's disk storage.**

```
# clresourcegroup switch -n node solaris-zone-resource-group
```

- b. If the zone's zone path already exists on the zone's disk storage, remove the zone path.**

The zone's zone path already exists on the zone's disk storage if you have previously installed the zone on another node.



Caution – If the zone is to run in a failover configuration, each node being able to host that zone must have the exact same zone configuration for that zone. After installing the zone on the first node, the zone's zone path already exists on the zones' disk storage. Therefore, the zone path must be removed on the next node prior to successfully creating and installing the zone. Otherwise, the next two steps will fail. Only a zone's zone path that is created on the last node will be kept as the final zone path for the HA container. For that reason, perform any configuration and customization within the HA container only after the HA container is known to all nodes that should be able to host it.

2 Create the zone.

```
# zonecfg -z zone
```

For more detailed information about creating a zone, see “Configuring, Verifying, and Committing a Zone” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*.

3 Install the zone.

```
# zoneadm -z zone install
```

For more detailed information about installing a zone, see “How to Install a Configured Zone” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*.

- 4 **Perform the initial internal zone configuration. If the zone is to run in a failover configuration, perform this step on the last node.**

- a. **Log in to the zone's console.**

```
# zlogin -C zone
```

You are prompted to configure the zone.

- b. **Follow the prompts to configure the zone.**

- c. **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:

```
# ~.
```

Verifying the Installation and Configuration of a Zone

Before you install the 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 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 15](#)

- 1 **Start the zone.**

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

- 2 **Log in to the zone.**

```
# zlogin zone
```

- 3 **Perform the required task depending upon the brand type of the zone.**

- **For a native brand type zone, 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

```

- For a **lx** brand type zone, confirm that the run level is 3.

```

# runlevel
N 3

```

- For a **solaris8** or **solaris9** brand type zone, confirm that the legacy run level is 3.

```

# who -r
run-level 3  Sep 10 23:49      3      0  S

```

4 Stop the zone.

```

# zoneadm -z zone halt

```

Next Steps Go to [“Registering and Configuring HA for Solaris Containers” on page 21.](#)

Installing the HA for Solaris Containers Packages

If you did not install the HA for Solaris Containers packages during your initial Oracle Solaris Cluster installation, perform this procedure to install the packages. To install the packages, use the installer program.

Note – You need to install the HA for Solaris Containers packages in the global cluster and not in the zone cluster.

▼ How to Install the HA for Solaris Containers Packages

Perform this procedure on each cluster node where you want the HA for Solaris Containers software to run.

You can run the installer program with a command-line interface (CLI) or with a graphical user interface (GUI). The content and sequence of instructions in the CLI and the GUI are similar.

Before You Begin Ensure that you have the Oracle Solaris Cluster installation media.

If you intend to run the installer program with a GUI, ensure that your **DISPLAY** environment variable is set.

1 On the cluster node where you are installing the data service packages, become superuser.

2 Load the Oracle Solaris Cluster installation media into the DVD-ROM drive.

If the Volume Management daemon `vold(1M)` is running and configured to manage DVD-ROM devices, the daemon automatically mounts the DVD-ROM on the `/cdrom` directory.

3 Change to the installation wizard directory of the DVD-ROM.

- If you are installing the data service packages on the SPARC platform, type the following command:

```
# cd /cdrom/cdrom0/Solaris_sparc
```

- If you are installing the data service packages on the x86 platform, type the following command:

```
# cd /cdrom/cdrom0/Solaris_x86
```

4 Start the installation wizard.

```
# ./installer
```

5 When you are prompted, accept the license agreement.**6 From the list of Oracle Solaris Cluster agents under Availability Services, select the data service for Solaris Zones.****7 If you require support for languages other than English, select the option to install multilingual packages.**

English language support is always installed.

8 When prompted whether to configure the data service now or later, choose Configure Later.

Choose Configure Later to perform the configuration after the installation.

9 Follow the instructions on the screen to install the data service packages on the node.

The installation wizard displays the status of the installation. When the installation is complete, the wizard displays an installation summary and the installation logs.

10 (GUI only) If you do not want to register the product and receive product updates, deselect the Product Registration option.

The Product Registration option is not available with the CLI. If you are running the installation wizard with the CLI, omit this step.

11 Exit the installation wizard.

12 Unload the installation media from the DVD-ROM drive.

a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.

b. Eject the DVD-ROM.

```
# eject cdrom
```

Registering and Configuring HA for Solaris Containers

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

Use the configuration and registration files in the following directories to register the 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 HA for Solaris Containers components. For information about these dependencies, see [“Dependencies Between HA for Solaris Containers Components” on page 14](#)

Registering and configuring 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 21](#)
2. [“Writing Scripts for the Zone Script Resource” on page 26](#)
3. [“Specifying Configuration Parameters for the Zone Script Resource” on page 27](#)
4. [“Writing a Service Probe for the Zone SMF Resource” on page 29](#)
5. [“Specifying Configuration Parameters for the Zone SMF Resource” on page 30](#)
6. [“How to Create and Enable Resources for the Zone Boot Component” on page 32](#)
7. [“How to Create and Enable Resources for the Zone Script Component” on page 32](#)
8. [“How to Create and Enable Resources for the Zone SMF Component” on page 33](#)

Specifying Configuration Parameters for the Zone Boot Resource

HA for Solaris Containers provides the script `sczbt_register`, which automates the process of configuring the zone boot resource. By default 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, copy the `sczbt_config` file to a different

filename and amend it as described below. It is recommended to keep this file as a future reference. The register script provides option `-f` to specify the fully qualified filename to the copied configuration file.

The directory you specify for the `PARAMETERDIR` can exist on the local file system of all nodes, the global file system, or the failover file system.

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 41](#). 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
Zonebrand=zone-brand-type
Zonebootopt=zone-boot-options
Milestone=zone-boot-milestone
LXrunlevel=linux-runlevel
SLrunlevel=solaris-legacy-runlevel
Mounts=list-of-mountpoints
```

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. This directory can exist on the local file system of all nodes, the global file system, or on a failover file system.

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 `zoncfg` utility.

```
SC_NETWORK=false
SC_LH=
```

- If only 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 by triggering a failover 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 zone's addresses that are not to be under the control of the `SUNW.LogicalHostName` resource.

```
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-lh-rs`

Specifies the name of the `SUNW.LogicalHostName` resource for the zone boot resource. Refer to [“Restrictions for Zone Network Addresses” on page 12](#) 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.HASStoragePlus` resource for the zone boot resource. This name must be the `SUNW.HASStoragePlus` resource name you assigned when you created the resource in [“How to Enable a Zone to Run in Failover Configuration” on page 15](#). 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.

Zonebrand=zone-brand-type

Specifies the brand type of the zone. The options that are currently supported are native (default), lx, solaris8, or solaris9. 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 successfully booted. This option is only used for the native brand type. You must specify a value for this keyword if you set the Zonebrand option to native.

LXrunlevel=linux-runlevel

Specifies the runlevel that needs to be attained before the zone is considered booted. This option is used only for the lx brand type. You must specify a value for this keyword if you set the Zonebrand option to lx.

SLrunlevel=solaris-legacy-runlevel

Specifies the legacy run level that needs to be attained before the zone is considered booted. This option is only used for the solaris8 and solaris9 brand types. You must specify a value for this keyword, if you set the Zonebrand option to solaris8 or solaris9.

Mounts=list-of-mountpoints

Specifies a space separated list of directories with their mount options, which will automatically get lofs mounted from the global zone into the booted zone. The mount point used in the global zone can be different to the mount point in the booted zone. Specifying a value for this keyword is optional.

The Mounts keyword format is as follows:

```
Mounts="/global-zone-dir:/local-zone-dir:mount-options <next entry>"
```

While *mount-options* can be a comma separated list of file system mount options.

The only required entry when setting this keyword is the */global-zone-dir* part of the colon separated variable. The */local-zone-dir* and *mount-options* part can be omitted.

Omitting the */local-zone-dir* part will make the zone's mount point the same as the global zone directory.

Omitting the *mount-options* part will not provide any mount options except the default options from the mount command.

Note – If you are omitting the */local-zone-dir* or the *mount-options*, you must also omit the “:” as delimiter.

Note – You must manually create any mount point directories within the booted zone that will be used within the Mounts keyword, before registering this resource within Oracle Solaris Cluster.

Note – If the file system of the source mount point in the global zone is mounted by a SUNW.HAStoragePlus resource, you must specify a strong resource dependency from the sczbt resource to this SUNW.HAStoragePlus resource.

EXAMPLE 1-1 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.
- The brand type of the zone is native.
- Indicates that the zone boot resource's boot option is null.
- Indicates that the zone boot resource's milestone is multi-user-server.
- Defines that /global/app/bin from the global zone gets mounted read-only within zone zone1 under mount point /app/bin.
- Defines that /app/data from the global zone gets mounted read-write within zone zone1 under mount point /app/data.
- Defines that /logs from the global zone gets mounted with default mount options within zone zone1 under mount point /logs.

```
RS=zone1 - rs
RG=zone1 - rg
PARAMETERDIR=/global/zones/pfiles
SC_NETWORK=true
```

EXAMPLE 1-1 Sample sczbt_config File (Continued)

```
SC_LH=zone1-lh
FAILOVER=true
HAS_RS=zone1-has
Zonename=zone1
Zonebrand=native
Zonebootopt=
Milestone=multi-user-server
Mounts="/global/app/bin:/app/bin:ro /app/data:rw /logs"
```

Writing Scripts for the Zone Script Resource

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 1-3 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 1-2 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 and being executable within the zone.

```
# cat /var/tmp/probe-apache2
#!/usr/bin/ksh
if echo "GET; exit" | mconnect -p 80 > /dev/null 2>&1
```

EXAMPLE 1-2 Zone Probe Script for Apache2 (Continued)

```

then
    exit 0
else
    exit 100
fi

# chmod 755 /var/tmp/probe-apache2

```

Specifying Configuration Parameters for the Zone Script Resource

HA for Solaris Containers provides the script `sczsh_register`, which automates the process of configuring zone script resource. By default 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, copy the `sczsh_config` file to a different filename and amend it as described below. It is recommended to keep this file as a future reference. The register script provides option `-f` to specify the fully qualified filename to the copied configuration 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 46](#). 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 1-3 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 httpd.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 1-2](#) and must exist in `zone1`.

```
RS="zone1-script-rs"
RG="zone1-rg"
SCZBT_RS="zone1-rs"
```

EXAMPLE 1-3 Sample sczsh_config File (Continued)

```
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 a Service Probe for the Zone SMF Resource

The zone SMF resource provides the ability to enable, disable, and probe an SMF service within a zone that is of brand type native. 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 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 1-4 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 1-4 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 and being executable within the zone.

```
# cat /var/tmp/probe-apache2
#!/usr/bin/ksh
if echo "GET; exit" | mconnect -p 80 > /dev/null 2>&1
then
    exit 0
```

EXAMPLE 1-4 Zone SMF Probe Script for Apache2 (Continued)

```
else
    exit 100
fi

# chmod 755 /var/tmp/probe-apache2
```

Specifying Configuration Parameters for the Zone SMF Resource

HA for Solaris Containers provides the script `sczsmf_register`, which automates the process of configuring the zone SMF resource. By default 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, copy the `sczsmf_config` file to a different filename and amend it as described below. It is recommended to keep this file as a future reference. The register script provides option `-f` to specify the fully qualified filename to the copied configuration 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 47](#). 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. Specifying a value for this keyword is optional.

EXAMPLE 1-5 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 1-4](#) 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 or a copy of it to specify configuration parameters for the HA for Solaris Containers zone boot component. For more information, see [“Specifying Configuration Parameters for the Zone Boot Resource” on page 21.](#)

- 1 **Become superuser on one of the nodes in the cluster that will host the zone.**
- 2 **Register the `SUNW.gds` resource type.**

```
# clresourcetype register SUNW.gds
```
- 3 **Go to the directory that contains the script for creating the HA for Solaris Containers boot resource.**

```
# cd /opt/SUNWsczone/sczbt/util
```
- 4 **Run the script that creates the zone boot resource.**
 Copy the `$PARAMETERDIR/sczbt_rs_name` from the node where this registration step was executed to `$PARAMETERDIR/` of all other nodes that can master the resource group if `$PARAMETERDIR` is a local file system.

```
# ./sczbt_register -f /mypath/sczbt_config
```
- 5 **Bring online the zone boot resource.**

```
# clresource enable sczbt-rs
```

Next Steps Go to [“Verifying the HA for Solaris Containers and Configuration” on page 33.](#)

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

Before You Begin Ensure you have edited the `sczsh_config` file or a copy of it to specify configuration parameters for the HA for Solaris Containers zone script component. For more information, see [“Specifying Configuration Parameters for the Zone Script Resource” on page 27.](#)

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

```
# cd /opt/SUNWsczone/sczsh/util
```
- 2 **Run the script that creates the zone script resource.**

```
# ./sczsh_register -f /mypath/sczsh_config
```


- 3 Bring online the zone script resource.

```
# clresource enable sczsh-rs
```

Next Steps Go to [“Verifying the HA for Solaris Containers and Configuration”](#) on page 33.

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

Before You Begin Ensure you have edited the `sczsmf_config` file or a copy of it to specify configuration parameters for the HA for Solaris Containers zone SMF component. For more information, see [“Specifying Configuration Parameters for the Zone SMF Resource”](#) on page 30.

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

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

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

```
# ./sczsmf_register -f /mypath/sczsmf_config
```

- 3 Bring online the zone SMF resource.

```
# clresource enable sczsmf-rs
```

Next Steps Go to [“Verifying the HA for Solaris Containers and Configuration”](#) on page 33.

Verifying the HA for Solaris Containers and Configuration

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

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

- 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.

```
# cluster status -t rg,rs
```
 - b. If the resource is not online, bring online the resource.

```
# clresource enable solaris-zone-resource
```
- 3 For a failover service configuration, switch the zone resource group to another cluster node, such as *node2*.

```
# clresourcegroup switch -n node2 solaris-zone-resource-group
```
- 4 Confirm that the resource is now online on *node2*.

```
# cluster status -t rg,rs
```

Patching the Global Zone and Non-Global Zones

The procedure that follows is required only if you are applying a patch to the global zone *and* to non-global zones. If you are applying a patch to only the global zone, follow the instructions in [Chapter 11, “Patching Oracle Solaris Cluster Software and Firmware,”](#) in *Oracle Solaris Cluster System Administration Guide*.

▼ How to Patch to the Global Zone and Non-Global Zones

This task applies to both nonrebooting patches and rebooting patches.

Perform this task on all nodes in the cluster.

- 1 Ensure that the node that you are patching can access the zone paths of all zones that are configured on the node.
Some zones might be configured to run in a failover configuration. In this situation, bring online on the node that you are patching the resource group that contains the resources for the zones' disk storage.

```
# clresourcegroup switch -n node solaris-zone-resource-group
```

Note – This step might also start any applications managed within the resource group *solaris-zone-resource-group*. Verify if you need to stop any application prior to install the patches. If the applications need to be stopped, disable the corresponding resources before proceeding to the next step.

Note – If the patches need to get applied in single-user mode, it will not be possible to start the resource group as described. Instead the corresponding zone paths need to get mounted manually.

2 Apply the patch(es) to the node.

For detailed instructions, see [Chapter 11, “Patching Oracle Solaris Cluster Software and Firmware,” in *Oracle Solaris Cluster System Administration Guide*](#) in *Oracle Solaris Cluster System Administration Guide*.

Tuning the HA for Solaris Containers Fault Monitors

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

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

Each 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 HA for Solaris Containers. For more information, see [“Registering and Configuring 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 Oracle Solaris Cluster installations. Therefore, you should tune the HA for Solaris Containers fault monitor *only* if you need to modify this preset behavior.

Tuning the 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 Oracle Solaris Cluster Data Services” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*](#).

Operation of the HA for Solaris Containers Parameter File

The 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 corresponding `zsched` process for the zone 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 logical hostname that is managed by a `SUNW.LogicalHostname` resource is operational.
If the logical hostname is not operational, the fault monitor fails over the resource group that contains resource for the zone boot component.
- The specified milestone for the native zone brand type is either online or degraded. The specified run level for the `lx`, `solaris8` or `solaris9` zone brand type is active.
If the milestone is not online or degraded, or if the specified run level is not active, 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 or run level, 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 resource for the zone boot component.

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 1–3](#).

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 the resource for the zone SMF component.

If the service is not disabled, the fault monitor runs the SMF service probe that you can 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 1–4](#).

Tuning the HA for Solaris Containers Stop_timeout property

The HA for Solaris Containers components consist all of the resource type `SUNW.gds(5)`. As described in “[Stop_command Property](#)” in *Oracle Solaris Cluster Data Services Developer’s Guide* the value for the `Stop_timeout` should be chosen so that the `Stop_command` can successfully return within 80% of its value.

Choosing the Stop_timeout value for the Zone Boot Component

The stop method for the zone boot component spends 60% of the value for the `Stop_timeout` performing a complete shutdown `-y -g0 -i0` within the zone. If that failed, the next 20% of the value for the `Stop_timeout` will be spent halting the zone performing a `"zoneadm -z zonename" halt` and perform some additional cleanup steps in order to force the zone into the state installed. Therefore the `Stop_timeout` value for the zone boot component should be computed so that 60% is enough to successfully shutdown the zone.

Choosing the Stop_timeout value for the Zone Script Component

The stop method for the zone script component calls the command or script configured for the `ServiceStopCommand` keyword. Therefore the `Stop_timeout` value for the zone script component should be computed so that 80% is enough for the configured `ServiceStopCommand` to succeed.

Choosing the Stop_timeout value for the Zone SMF Component

The stop method for the zone SMF component spends 60% of the value for the Stop_timeout using svcadm to disable the configured SMF service in the zone. If that failed, the next 20% of the value for the Stop_timeout will be spent to first send SIGTERM then SIGKILL to the processes associated with this SMF service. Therefore the Stop_timeout value for the zone SMF component should be computed so that 60% is enough to successfully disable the configured SMF service in the zone.

Denying Cluster Services for a 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.

For instructions, see “[How to Deny Cluster Services For a Non-Global Zone](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide* and “[How to Allow Cluster Services For a Non-Global Zone](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

Debugging HA for Solaris Containers

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

Each component of 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 HA for Solaris Containers

1 Determine whether debugging for 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 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.

```
# svcadm restart system-log
```

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 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident "@(#)config      1.1      06/02/22 SMI"
#
# Usage:
#      DEBUG=<RESOURCE_NAME> or ALL
#
DEBUG=ALL
#
```

Note – To deactivate debugging, reverse the preceding steps.

Files for Configuring HA for Solaris Containers Resources

The /opt/SUNWsczone/zone component/util directory contains files that automate the process of configuring HA for Solaris Containers resources. Listings of these files are provided in the following sections:

- [“Listing of sczbt_config” on page 41](#)
- [“Listing of sczsh_config” on page 46](#)
- [“Listing of sczsmf_config” on page 47](#)

Listing of sczbt_config

```
#
# Copyright 2008 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident      "@(#)sczbt_config    1.5    08/04/17 SMI"
#
# 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
#
#    NOTE: If the ip-type keyword for the non-global zone is set
#
#        to "exclusive", only "false" is allowed for SC_NETWORK
#
#    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=Name of the SC Logical Hostname resource
#
#    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=Name of the SC Logical Hostname resource
#
#    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 HAStoragePlus 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
# Zonebrand     Brand of the zone. Current supported options are
#               "native" (default), "lx", "solaris8" or "solaris9"
# Zonebootopt   Zone boot options ("-s" requires that Milestone=single-user)
# Milestone     SMF Milestone which needs to be online before the zone is
#               considered booted. This option is only used for the
#               "native" Zonebrand.
# LXrunlevel    Runlevel which needs to get reached before the zone is
#               considered booted. This option is only used for the "lx"
#               Zonebrand.
# SLrunlevel    Solaris legacy runlevel which needs to get reached before the
#               zone is considered booted. This option is only used for the
#               "solaris8" or "solaris9" Zonebrand.
```

```

# Mounts      Mounts is a list of directories and their mount options,
#              which are loopback mounted from the global zone into the
#              newly booted zone. The mountpoint in the local zone can
#              be different to the mountpoint from the global zone.
#
#              The Mounts parameter format is as follows,
#
#              Mounts="/global zone directory:/local zone directory:mount options"
#
#              The following are valid examples for the "Mounts" variable
#
#              Mounts="/globalzone-dir1:/localzone-dir1:rw"
#              Mounts="/globalzone-dir1:/localzone-dir1:rw /globalzone-dir2:rw"
#
#              The only required entry is the /global zone directory, the
#              /local zone directory and mount options can be omitted.
#
#              Omitting /local zone directory will make the local zone
#              mountpoint the same as the global zone directory.
#
#              Omitting mount options will not provide any mount options
#              except the default options from the mount command.
#
#              Note: You must manually create any local zone mountpoint
#                     directories that will be used within the Mounts variable,
#                     before registering this resource within Solaris Cluster.
#
Zonename=""

```

```
Zonebrand="native"

Zonebootopt=""

Milestone="multi-user-server"

LXrunlevel="3"

SLrunlevel="3"

Mounts=""
```

Listing of sczsh_config

```
#
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident      "@(#)sczsh_config    1.1    06/02/22 SMI"
#
# 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 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
# ident    "@(#)sczsmf_config    1.1    06/02/22 SMI"
#
# 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
#         the 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
 clresource, 15, 32, 33
 clresourcegroup, 16
 clresourcetype, 15–16, 32
 svcs, 18
 zlogin, 18
 zoneadm, 17, 18, 19
 zonecfg, 17
configuration files, 21
 zone boot resource (sczbt), 21
 zone script resource (sczsh), 21
 zone SMF resource (sczsmf), 21
create logical host, clreslogicalhostname, 16
create multiple-master resource group,
 clresourcegroup, 16
create resource group, clresourcegroup, 15

D

debugging, HA for Solaris Containers, 38–39
denying cluster services, non-global zone, 38
dependencies, 14

E

enable resource group, 16
extension properties, effect on fault monitors, 35

F

fault monitors, tuning, 35–37
files
 config, 38
 /opt/SUNWsczone/sczbt/etc/config, 38
 /opt/SUNWsczone/sczsh/etc/config, 38
 /opt/SUNWsczone/sczsmf/etc/config, 38
 sczbt_config
 listing, 41–46
 sczsh_config
 listing, 46
 sczsmf_config
 listing, 47–48

H

HA for Solaris Containers
 debugging, 38–39
 fault monitors, 35–37
 installing, 19–21
 software packages, installing, 19–21
 Stop_timeout, 37–38
help, 8

I

installing, HA for Solaris Containers, 19–21

N

Non-global zone, turning off cluster functionality, 38

O

/opt/SUNWsczone/sczbt/etc/config file, 38
/opt/SUNWsczone/sczsh/etc/config file, 38
/opt/SUNWsczone/sczsmf/etc/config file, 38
overview, 9–10

P

packages, 19–21
parameter file directory, 14–15
patching, 34–35

R

register
 sczbt_register, 32
 sczsh_register, 32
 sczsmf_register, 33
requirements, 14–15
resource types, fault monitors, 35–37
resources
 Solaris Zones application
 debugging, 38–39
restrictions, 12–13
return codes
 zone script resource (sczsh), 26–27
 zone SMF resource (sczsmf), 29–30

S

sczbt_config file, listing, 41–46
sczsh_config file, listing, 46
sczsmf_config file, 47–48
software packages, 19–21
Solaris Zones application, fault monitor, 35–37
Stop_timeout
 tuning, 37–38

Stop_timeout (*Continued*)

 zone boot resource (sczbt), 37
 zone script resource (sczsh), 37
 zone SMF resource (sczsmf), 38
system properties, effect on fault monitors, 35

T

technical support, 8
tuning
 fault monitors, 35–37
 Stop_timeout, 37–38

Z

zone boot resource (sczbt), Stop_timeout, 37
zone script resource (sczsh)
 return codes, 26–27
 Stop_timeout, 37
zone SMF resource (sczsmf)
 return codes, 29–30
 Stop_timeout, 38