



Sun Cluster 3.1 Data Service for Sun ONE Web Server

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

Part No: 817-1528
May 2003, Revision A

Copyright 2003 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, registered trademarks, or service marks 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.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 2003 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 Etats-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, ou marques de service, de Sun Microsystems, Inc. aux Etats-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 Etats-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.



030314@5533



Contents

Preface 5

Installing and Configuring Sun Cluster HA for Sun ONE Web Server 11

Planning the Installation and Configuration 12

Installing and Configuring Sun Cluster HA for Sun ONE Web Server 13

Installing and Configuring a Sun ONE Web Server 14

▼ How to Install a Sun ONE Web Server 14

▼ How to Configure a Sun ONE Web Server 16

Installing Sun Cluster HA for Sun ONE Web Server Packages 18

▼ How to Install Sun Cluster HA for Sun ONE Web Server Packages 18

Registering and Configuring Sun Cluster HA for Sun ONE Web Server 19

▼ How to Register and Configure Sun Cluster HA for Sun ONE Web Server 19

▼ How to Configure `SUNW.HAStoragePlus` Resource Type 26

Configuring Sun Cluster HA for Sun ONE Web Server Extension Properties 27

Sun Cluster HA for Sun ONE Web Server Fault Monitor 29

Index 33

Preface

The Sun Cluster 3.1 Data Service for Sun ONE Web Server contains procedures to install and configure the Sun Cluster data services.

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 environment and expertise with the volume manager software used with Sun Cluster.

UNIX Commands

This document contains information on commands specific to installing and configuring Sun Cluster data services. It might not contain information on basic UNIX® commands and procedures, such as shutting down the system, booting the system, and configuring devices. For that information, see one or more of the following:

- Online documentation for the Solaris software environment
- Solaris operating environment man pages
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	<i>machine_name</i> %
C shell superuser	<i>machine_name</i> #
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Installation	<i>Sun Cluster 3.1 Software Installation Guide</i>	816-3388
Data Services Installation and Configuration	<i>Sun Cluster 3.1 Data Service Planning and Administration Guide</i>	817-1526
	<i>Sun Cluster 3.1 Data Service for Apache</i>	817-1532
	<i>Sun Cluster 3.1 Data Service for BroadVision One-To-One Enterprise</i>	817-1542
	<i>Sun Cluster 3.1 Data Service for Domain Name Service (DNS)</i>	817-1533
	<i>Sun Cluster 3.1 Data Service for Netbackup</i>	817-1538
	<i>Sun Cluster 3.1 Data Service for Network File System (NFS)</i>	817-1534
	<i>Sun Cluster 3.1 Data Service for Oracle</i>	817-1527
	<i>Sun Cluster 3.1 Data Service for Oracle Parallel Server/Real Application Clusters</i>	817-1535
	<i>Sun Cluster 3.1 Data Service for SAP</i>	817-1536
	<i>Sun Cluster 3.1 Data Service for SAP liveCache</i>	817-1539
	<i>Sun Cluster 3.1 Data Service for Siebel</i>	817-1540
	<i>Sun Cluster 3.1 Data Service for Sun ONE Application Server</i>	817-1530
	<i>Sun Cluster 3.1 Data Service for Sun ONE Directory Server</i>	817-1529
	<i>Sun Cluster 3.1 Data Service for Sun ONE Message Queue</i>	817-1531
	<i>Sun Cluster 3.1 Data Service for Sun ONE Web Server</i>	817-1528
	<i>Sun Cluster 3.1 Data Service for Sybase ASE</i>	817-1537
<i>Sun Cluster 3.1 Data Service for DHCP</i>	817-1716	
<i>Sun Cluster 3.1 Data Service for Samba</i>	817-1715	
<i>Sun Cluster 3.1 Data Service for WebSphere MQ</i>	817-1714	
<i>Sun Cluster 3.1 Data Service for WebSphere MQ Integrator</i>	817-1713	

Application	Title	Part Number
Hardware	<i>Sun Cluster 3.x Hardware Administration Manual</i>	817-0168
	Sun Cluster 3.x Data Service Collection at http://docs.sun.com/db/coll/1024.1	
API development	<i>Sun Cluster 3.1 Data Services Developer's Guide</i>	816-3385
Administration	<i>Sun Cluster 3.1 5/03 System Administration Guide</i>	816-3384
Cluster concepts	<i>Sun Cluster 3.1 Concepts Guide</i>	816-3383
Error Messages	<i>Sun Cluster 3.1 Error Messages Guide</i>	816-3382
Man Pages	<i>Sun Cluster 3.1 Reference Manual</i>	816-5251
Release notes	<i>Sun Cluster 3.1 Data Service 5/03 Release Notes</i>	817-1790
	<i>Sun Cluster 3.1 Release Notes</i>	816-5317
	<i>Sun Cluster 3.1 Release Notes Supplement</i>	816-3381

Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation on the Web. You can browse the docs.sun.com archive or search for a specific book title or subject at <http://docs.sun.com>.

Help

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

- Your name and E-mail address (if available)
- Your company name, address, and phone number
- The model and serial numbers of your systems
- The release number of the operating environment (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.
<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 Sun™ ONE Web Server

This chapter provides the procedures to install and configure Sun Cluster HA for Sun ONE Web Server. This data service was formerly known as Sun Cluster HA for Netscape™ HTTP and Sun Cluster HA for iPlanet Web Server. Some error messages from the application might use the name Netscape, but the messages refer to Sun ONE Web Server. The application name on the Sun Cluster Agents CD-ROM might still be iPlanet Web Server.

This chapter contains the following procedures.

- “How to Install a Sun ONE Web Server” on page 14
- “How to Configure a Sun ONE Web Server” on page 16
- “How to Install Sun Cluster HA for Sun ONE Web Server Packages” on page 18
- “How to Register and Configure Sun Cluster HA for Sun ONE Web Server” on page 19
- “How to Configure SUNW.HAStoragePlus Resource Type” on page 26

You can configure Sun Cluster HA for Sun ONE Web Server as a failover or scalable data service. See “Planning for Sun Cluster Data Services” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* and the *Sun Cluster 3.1 Concepts Guide* document for general information about data services, resource groups, resources, and other related topics.

Note – You can use SunPlex Manager to install and configure this data service. See the SunPlex Manager online help for details.

Note – If you run multiple data services in your Sun Cluster configuration, you can set up the data services in any order, with the following exception. If Sun Cluster HA for Sun ONE Web Server depends on Sun Cluster HA for DNS, you must set up DNS first. See *Sun Cluster 3.1 Data Service for Domain Name Service (DNS)* for details. The Solaris operating environment includes the DNS software. If the cluster is to obtain the DNS service from another server, then configure the cluster to be a DNS client first.

Note – After installation, do not manually start and stop the Sun ONE Web Server except by using the cluster administration command `scswitch(1M)`. See the man page for details. After the Sun ONE Web Server is started, the Sun Cluster software controls it.

Planning the Installation and Configuration

Use the following section in conjunction with your configuration worksheets as a checklist before you install and configure Sun Cluster HA for Sun ONE Web Server.

Consider the following questions before you start your installation.

- Will you run Sun Cluster HA for Sun ONE Web Server as a failover or as a scalable data service? See the *Sun Cluster 3.1 Concepts Guide* document for information on the two types of services. For scalable services, consider the following questions.
 - What nodes will host the scalable service? In most cases, you will want to scale across all nodes. You can, however, limit the set of nodes that host the service.
 - Will your Sun ONE Web Server instances require sticky IP? Sticky IP is a resource property setting, `Load_balancing_policy`, which stores the client state in memory so that return traffic from the same node always goes to the same cluster node. You can choose from several load balancing policies, as described in the table on resource properties in “Standard Properties” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

Exercise caution when you change `Load_balancing_weights` for an online scalable service that has `Load_balancing_policy` set to `LB_STICKY` or `LB_STICKY_WILD`. Changing those properties while the service is online can cause existing client affinities to be reset, and hence a different node might service a subsequent client request even if another cluster member had previously serviced the client.

Similarly, when a new instance of the service is started on a cluster, existing client affinities might be reset.

- Where will the Web server root reside?
- Does the Web server serve data for another highly available application? If so, resource dependencies might exist between the resources so that one starts or stops before the other. See “Standard Properties” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for a description of the resource property `Resource_dependencies` that sets these dependencies.
- Determine the resource groups to use for network addresses and application resources and the dependencies between them. See “Standard Properties” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for a description of the resource group property `RG_dependencies` that sets these dependencies.
- Provide the logical hostname (for failover services) or shared address (for scalable services) for clients to use to access the data service.
- Because you can configure Sun ONE Web Server to bind to `INADDR_ANY`, if you plan to run multiple instances of the Sun ONE Web Server data service or multiple data services on the same node, each instance must bind to a unique network address and port number.
- Determine the entries for the `Confdir_list` and `Port_list` properties. For failover services, both of these properties can have only one entry. For scalable services, they can have multiple entries. The number of entries, however, must be the same and must map to each other in the order specified. See “How to Register and Configure Sun Cluster HA for Sun ONE Web Server” on page 19 for details.
- Determine where to place logs, error files, and the PID file on the local file system.
- Determine where to place the contents on the cluster file system.

Installing and Configuring Sun Cluster HA for Sun ONE Web Server

The following table lists the sections that describe the installation and configuration tasks.

TABLE 1-1 Task Map: Installing and Configuring Sun Cluster HA for Sun ONE Web Server

Task	For Instructions, Go To
Install Sun ONE Web Server	“Installing and Configuring a Sun ONE Web Server” on page 14

TABLE 1-1 Task Map: Installing and Configuring Sun Cluster HA for Sun ONE Web Server (Continued)

Task	For Instructions, Go To
Install the Sun Cluster HA for Sun ONE Web Server packages	"Installing Sun Cluster HA for Sun ONE Web Server Packages" on page 18
Configure Sun Cluster HA for Sun ONE Web Server	"Registering and Configuring Sun Cluster HA for Sun ONE Web Server" on page 19
Configure resource extension properties	"Configuring Sun Cluster HA for Sun ONE Web Server Extension Properties" on page 27
View fault monitor information	"Sun Cluster HA for Sun ONE Web Server Fault Monitor" on page 29

Installing and Configuring a Sun ONE Web Server

This section describes the steps to use the `setup` command to perform the following tasks.

- Install the Sun ONE Web Server.
- Enable the Sun ONE Web Server to run as Sun Cluster HA for Sun ONE Web Server.

Note – You must follow certain conventions when you configure URL mappings for the Web server. For example, to preserve availability when setting the CGI directory, you must locate the mapped directories on the cluster file system. In this example, you map your CGI directory to `/global/pathname/cgi-bin`.

In situations where the CGI programs access "back-end" servers, such as an RDBMS, ensure that the Sun Cluster software also controls the "back-end" server. If the server is an RDBMS that the Sun Cluster software supports, use one of the highly available RDBMS packages. Alternatively, you can use the APIs documented in the *Sun Cluster 3.1 Data Services Developer's Guide* to put the server under Sun Cluster control.

▼ How to Install a Sun ONE Web Server

To perform this procedure, you need the following information about your configuration.

- The server root directory (the path to the application binaries). You can install the binaries on the local disks or on the cluster file system. For a discussion of the advantages and disadvantages of each location, see “Configuration Guidelines for Sun Cluster Data Services” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.
- The logical hostname (for failover services) or shared address (for scalable services) that clients use to access the data service. You must configure these addresses, and they must be online.

Note – If you run Sun Cluster HA for Sun ONE Web Server and another HTTP server and they use the same network resources, configure them to listen on different ports. Otherwise, a port conflict might occur between the two servers.

1. **Become superuser on a cluster member.**
2. **Run the `setup` command for iPlanet Web Server from the install directory on the CD.**
3. **When prompted, enter the location where the Sun ONE Web Server binaries will be installed.**

You can specify a location on the cluster file system or on local disks for the location of the install. If you choose to install on local disks, run the `setup` command on all of the cluster nodes that are potential primaries of the network resource (logical hostname or shared address) that the next step specifies.

4. **When prompted for a machine name, enter the logical hostname on which the Sun ONE Web Server depends and the appropriate DNS domain name.**

A full logical hostname is of the format *network-resource.domainname*, such as `schost-1.sun.com`.

Note – For Sun Cluster HA for Sun ONE Web Server to fail over correctly, you must use either the logical hostname or shared address resource name (rather than the physical hostname) here and everywhere else that you are asked.

5. **Select Run Admin Server as Root when you are asked.**

Note the port number that the Sun ONE install script selects for the administration server. You might want to use this default value later when you use the admin server to configure an instance of the Sun ONE Web Server. Otherwise, you can specify a different port number when you configure the Sun ONE server instance.

6. **Type a Server Administrator ID and a chosen password when you are asked.**
Follow the guidelines for your system.

When a message displays that the admin server will be started, your installation is ready for configuration.

Where to Go From Here

To configure the Sun ONE Web Server, see “How to Configure a Sun ONE Web Server” on page 16.

▼ How to Configure a Sun ONE Web Server

This procedure describes how to configure an instance of the Sun ONE Web server to be highly available. Use the Netscape browser to interact with this procedure.

Consider the following points before you perform this procedure.

- Before you start, ensure that you have installed the browser on a machine that can access the network on which the cluster resides. You can install the browser on a cluster node or on the administrative workstation for the cluster.
- Your configuration files can reside on either a local file system or on the cluster file system.
- Any certificates that are installed for the secure instances must be installed from all cluster nodes. This installation involves running the admin console on each node. Thus, if a cluster has nodes *n1*, *n2*, *n3*, and *n4*, the installation steps are as follows.
 1. Run the admin server on node *n1*.
 2. From your Web browser, connect to the admin server as `http://n1.domain:port`—for example, `http://n1.eng.sun.com:8888`—or whatever you specified as the admin server port. The port is typically 8888.
 3. Install the certificate.
 4. Stop the admin server on node *n1* and run the admin server from node *n2*.
 5. From the Web browser, connect to the new admin server as `http://n2.domain:port`, for example, `http://n2.eng.sun.com:8888`.
 6. Repeat these steps for nodes *n3* and *n4*.

After you have considered the preceding points, complete the following steps.

- 1. Create a directory on the local disk of all the nodes to hold the logs, error files, and PID file that Sun ONE Web Server manages.**

For Sun ONE to work correctly, these files must be located on each node of the cluster, not on the cluster file system.

Choose a location on the local disk that is the same for all the nodes in the cluster. Use the `mkdir -p` command to create the directory. Make nobody the owner of this directory.

The following example shows how to complete this step.

```
phys-schost-1# mkdir -p /var/pathname/http-instance/logs/
```

Note – If you anticipate large error logs and PID files, do not put them in a directory under `/var` because they will overwhelm this directory. Rather, create a directory in a partition with adequate space to handle large files.

2. From the administrative workstation or a cluster node, start the Netscape browser.
3. On one of the cluster nodes, go to the directory `https-admserv`, then start the Sun ONE admin server.

```
# cd https-admserv
# ./start
```

4. Enter the URL of the Sun ONE admin server in the Netscape browser.

The URL consists of the physical hostname and port number that the Sun ONE installation script established in Step 4 of the server installation procedure, for example, `n1.eng.sun.com:8888`. When you perform Step 2 of this procedure, the `./start` command displays the admin URL.

When prompted, use the user ID and password you specified in Step 6 of the server installation procedure to log in to the Sun ONE administration server interface.

5. Using the administration server where possible and manual changes otherwise, complete the following:
 - Verify that the server name is correct.
 - Verify that the server user is set as superuser.
 - Change the bind address field to one of the following addresses.
 - A logical hostname or shared address if you use DNS as your name service
 - The IP address associated with the logical hostname or shared address if you use NIS as your name service
 - Update the ErrorLog, PidLog, and Access Log entries to reflect the directory created in Step 1 of this section.
 - Save your changes.
6. Create a file that contains the secure key password you need to start this instance, and place this file under the server root directory. Name this file `keypass`.

Note – Because this file contains the key database password, protect the file with the appropriate permissions.

Installing Sun Cluster HA for Sun ONE Web Server Packages

You can use the `scinstall(1M)` utility to install `SUNWschtt`, the Sun Cluster HA for Sun ONE Web Server package, on a cluster. Do not use the `-s` option to non-interactive `scinstall` to install all data service packages on the CD.

If you installed the data service packages during your initial Sun Cluster installation, proceed to “Registering and Configuring Sun Cluster HA for Sun ONE Web Server” on page 19. Otherwise, use the following procedure to install the `SUNWschtt` package.

▼ How to Install Sun Cluster HA for Sun ONE Web Server Packages

You need the Sun Cluster Agents CD-ROM to complete this procedure. Perform this procedure on all the cluster nodes that will run Sun Cluster HA for Sun ONE Web Server.

- 1. Load the Sun Cluster Agents CD-ROM into the CD-ROM drive.**
- 2. Run the `scinstall` utility with no options.**

This step starts the `scinstall` utility in interactive mode.
- 3. Choose the menu option, Add Support for New Data Service to This Cluster Node.**

The `scinstall` utility prompts you for additional information.
- 4. Provide the path to the Sun Cluster Agents CD-ROM.**

The utility refers to the CD as the “data services cd.”
- 5. Specify the data service to install.**

The `scinstall` utility lists the data service that you selected and asks you to confirm your choice.

Note – The application name on the CD-ROM might be iPlanet Web Server.

- 6. Exit the `scinstall` utility.**
- 7. Unload the CD from the drive.**

Where to Go From Here

See “Registering and Configuring Sun Cluster HA for Sun ONE Web Server” on page 19 to register Sun Cluster HA for Sun ONE Web Server and configure the cluster for the data service.

Registering and Configuring Sun Cluster HA for Sun ONE Web Server

You can configure Sun Cluster HA for Sun ONE Web Server as a failover data service or as a scalable data service. You must include some additional steps to configure Sun ONE Web Server as a scalable data service. In the first procedure in this section, these additional steps begin with a notation that they are required for scalable services only. Individual examples of a failover service and a scalable service follow the procedure.

▼ How to Register and Configure Sun Cluster HA for Sun ONE Web Server

This procedure describes how to use the `scrgadm(1M)` command to register and configure Sun Cluster HA for Sun ONE Web Server.

Note – Other options also enable you to register and configure the data service. See “Tools for Data Service Resource Administration” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for details about these options.

To perform this procedure, you must have the following information.

- The name of the resource type for Sun Cluster HA for Sun ONE Web Server. This name is `SUNW.iws`.
- The names of the cluster nodes that master the data service. For a failover service, only one node can master a data service at a time.
- The logical hostname (for failover services) or shared address (for scalable services) that clients use to access the data service.
- The path to the Sun ONE binaries. You can install the binaries on the local disks or the cluster file system. See “Configuration Guidelines for Sun Cluster Data Services” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for a discussion of the advantages and disadvantages of each location.

Note – The `Network_resources_used` setting on the Sun ONE application resource determines the set of IP addresses that Sun ONE Web Server uses. The `Port_list` setting on the resource determines the list of port numbers that Sun ONE Web Server uses. The fault monitor assumes that the Sun ONE Web Server daemon is listening on all combinations of IP and port. If you have customized your `magnus.conf` file for the Sun ONE Web Server to listen on different port numbers (in addition to port 80), your resultant `magnus.conf` file must contain all possible combinations of IP address and ports. The fault monitor attempts to probe all such combinations and starts to fail if the Sun ONE Web Server is not listening on a particular IP address-port combination. If the Sun ONE Web Server does not serve all IP address-port combinations, you must break the Sun ONE Web Server into separate instances that do.

Note – Perform this procedure on any cluster member.

1. **Become superuser on a cluster member.**
2. **Register the resource type for Sun Cluster HA for Sun ONE Web Server.**

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

-a Adds the data service resource type.

-t SUNW.iws Specifies the predefined resource type name for your data service.

3. **Create a failover resource group to hold the network and application resources.**

For failover services, this resource group also holds the application resources. You can optionally select the set of nodes on which the data service can run with the `-h` option.

```
# scrgadm -a -g resource-group [-h nodelist]
```

-g *resource-group* Specifies the name of the failover resource group. This name can be your choice but must be unique for resource groups within the cluster.

-h *nodelist* An optional comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the nodes are considered as primary during failover.

Note – Use `-h` to specify the order of the node list. If all the nodes in the cluster are potential masters, you do not need to use the `-h` option.

4. **Verify that all of the network addresses that you use have been added to your name service database.**

You should have performed this verification during the Sun Cluster installation. See the planning chapter in the *Sun Cluster 3.1 Software Installation Guide* for details.

Note – To avoid any failures because of name service lookup, ensure that all logical hostnames and shared addresses are present in the server’s and client’s `/etc/inet/hosts` file. Configure name service mapping in `/etc/nsswitch.conf` on the servers to first check the local files before trying to access NIS or NIS+.

5. **Add a network resource (logical hostname or shared address) to the failover resource group.**

```
# scrgadm -a {-S | -L} -g resource-group \  
-l network-resource, ... [-j resource] \  
[-X auxnodelist=node, ...] [-n netiflist]
```

- S | -L You use -S for shared address resources or -L for logical hostname resources.
- g *resource-group* Specifies the name of the failover resource group.
- l *network-resource, ...* Specifies a comma-separated list of network resources to add. You can use the -j option to specify a name for the resources. If you do not do so, the network resources have the name of the first entry on the list.
- j *resource* Specifies an optional resource name. If you do not supply this name, the name of the network resource defaults to the first name that is specified after the -l option.
- X *auxnodelist=node, ...* Specifies an optional comma-separated list of physical node IDs that identify cluster nodes that can host the shared address but never serve as a primary if failover occurs. These nodes are mutually exclusive with the nodes identified in *nodelist* for the resource group, if specified.
- g *resource-group* Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster.
- h *nodelist* Specifies an optional comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the nodes are considered as primary during failover.
- n *netiflist* Specifies an optional, comma-separated list that identifies the IP Networking Multipathing groups that are on each node. Each element in *netiflist* must be in the form of

netif@node. netif can be given as an IP Networking Multipathing group name, such as sc_ipmp0. The node can be identified by the node name or node ID, such as sc_ipmp0@1 or sc_ipmp@phys-schost-1.

Note – Sun Cluster does not currently support using the adapter name for netif.

6. For scalable services only – Create a scalable resource group to run on all of the desired cluster nodes.

If you run Sun Cluster HA for Sun ONE Web Server as a failover data service, do not perform this step—go to Step 8.

Create a resource group to hold a data service application resource. You must specify the maximum and desired number of primary nodes, as well as a dependency between this resource group and the failover resource group that you created in Step 3. This dependency ensures that in the event of failover, the resource manager starts the network resource before starting any data services that depend on the network resource.

```
# scrgadm -a -g resource-group \  
-y Maximum primaries=m -y Desired primaries=n \  
-y RG_dependencies=resource-group
```

-y Maximum primaries=m

Specifies the maximum number of active primary nodes allowed for this resource group. If you do not assign a value to this property, the default is 1.

-y Desired primaries=n

Specifies the desired number of active primary nodes allowed for this resource group. If you do not assign a value to this property, the default is 1.

-y RG_dependencies= resource-group

Identifies the resource group that contains the shared address resource on which the resource group being created depends.

7. For scalable services only – Create an application resource in the scalable resource group.

If you run Sun Cluster HA for Sun ONE Web Server as a failover data service, do not perform this step—go to Step 8.

You can repeat this step to add multiple application resources (such as secure and insecure versions) to the same resource group.

You might also want to set load balancing for the data service. To do so, use the two standard resource properties Load_balancing_policy and Load_balancing_weights. See “Standard Properties” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for a description of these properties. Additionally, see the examples that follow this section.

```

# scrgadm -a -j resource -g resource-group \
-t resource-type -y Network_resources_used=network-resource, ... \
-y Port_list=port-number/protocol, ... -y Scalable=True \
-x Confdir_list=config-directory, ...

-j resource
    Specifies the name of the resource to add.

-g resource-group
    Specifies the name of the scalable resource group into which the resources are to be
    placed.

-t resource-type
    Specifies the type of the resource to add.

-y Network_resources_used= network-resource, ...
    Specifies a comma-separated list of network resources that identify the shared
    addresses that the data service uses.

-y Port_list=port-number/protocol, ...
    Specifies a comma-separated list of port numbers and protocol to be used, for
    example, 80/tcp, 81/tcp.

-y Scalable=True
    Specifies a Boolean that is required for scalable services.

-x Confdir_list=config-directory, ...
    Specifies a comma-separated list of the locations of the Sun ONE configuration files.
    Sun Cluster HA for Sun ONE Web Server requires this extension property.

```

Note – A one-to-one mapping applies for Confdir_List and Port_List, that is, each of the values in one list must correspond to the values in the other list in the order specified.

8. For failover services only – Create an application resource in the failover resource group.

Perform this step only if you run Sun Cluster HA for Sun ONE Web Server as a failover data service. If you run Sun Cluster HA for Sun ONE Web Server as a scalable service, you must have performed Step 6 and Step 7 previously and must now go to Step 10.

You can repeat this step to add multiple application resources (such as secure and insecure versions) to the same resource group.

```

# scrgadm -a -j resource -g resource-group \
-t resource-type -y Network_resources_used=logical-hostname-list \
-y Port_list=port-number/protocol \
-x Confdir_list=config-directory

-j resource
    Specifies the name of the resource to add.

```

- g *resource-group*
Specifies the name of the failover resource group into which the resources are to be placed.
- t *resource-type*
Specifies the type of the resource to add.
- y Network_resources_used=*network-resource*, ...
Specifies a comma-separated list of network resources that identify the logical hosts that the data service uses.
- y Port_list=*port-number/protocol*
Specifies the port number and protocol to use, for example, 80/tcp. Port_list for failover services must have exactly one entry only because of the one-to-one mapping rule between Port_list and Confdir_list.
- x Confdir_list=*config-directory*
Specifies the location of the Sun ONE configuration files. The Confdir_list file for failover services must have exactly one entry only. The *config-directory* must contain a directory called config. Sun Cluster HA for Sun ONE Web Server requires this extension property.

Note – Optionally, you can set additional extension properties that belong to the Sun Cluster HA for Sun ONE Web Server to override the default values of the properties. See Table 1-2 for a list of these properties.

9. Bring the failover resource group online.

- ```
scswitch -Z -g resource-group
```
- Z Enables the network resource and fault monitoring, switches the resource group into a MANAGED state, and brings the resource group online.
  - g *resource-group* Specifies the name of the failover resource group.

### 10. For scalable services only – Bring the scalable resource group online.

- ```
# scswitch -Z -g resource-group
```
- Z Enables the resource and monitor, moves the resource group to the MANAGED state, and brings the resource group online.
 - g *resource-group* Specifies the name of the scalable resource group.

Example – Registering Scalable Sun Cluster HA for Sun ONE Web Server

The following example shows how to register a scalable Sun Cluster HA for Sun ONE Web Server.

Cluster Information

Node names: *phys-schost-1, phys-schost-2*

Shared address: *schost-1*

Resource groups: *sa-resource-group-1* (for shared addresses),
iws-resource-group-1 (for scalable Sun ONE application resources)

Resources: *schost-1* (shared address), *SunONE-insecure-1* (insecure SunONE application resource), *SunONE-secure-1* (secure Sun ONE application resource)

(Add a failover resource group to contain shared addresses.)

```
# scrgadm -a -g sa-resource-group-1
```

(Add the shared address resource to the failover resource group.)

```
# scrgadm -a -S -g sa-resource-group-1 -l schost-1
```

(Add a scalable resource group.)

```
# scrgadm -a -g iws-resource-group-1 -y Maximum primaries=2 \  
-y Desired primaries=2 -y RG_dependencies=sa-resource-group-1
```

(Register the Sun ONE resource type.)

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

(Add an insecure Sun ONE instance with default load balancing.)

```
# scrgadm -a -j Sun ONE-insecure-1 -g iws-resource-group-1 -t SUNW.iws \  
-x Confdir_List=/opt/SunONE/https-SunONE-insecure-1 \  
-y Scalable=True -y Network_resources_used=schost-1 -y Port_list=80/tcp
```

(Add a secure SunONE instance with sticky IP load balancing.)

```
# scrgadm -a -j SunONE-secure-1 -g iws-resource-group-1 -t SUNW.iws \  
-x Confdir_List=/opt/SunONE/https-SunONE-secure-1 \  
-y Scalable=True -y Network_resources_used=schost-1 \  
-y Port_list=443/tcp -y Load_balancing_policy=LB_STICKY \  
-y Load_balancing_weights=40@1,60@2
```

(Bring the failover resource group online.)

```
# scswitch -Z -g sa-resource-group-1
```

(Bring the scalable resource group online.)

```
# scswitch -Z -g iws-resource-group-1
```

Example – Registering Failover Sun Cluster HA for Sun ONE Web Server

The following example shows how to register a failover Sun ONE service on a two-node cluster.

Cluster Information

Node names: *phys-schost-1, phys-schost-2*

Logical hostname: *schost-1*

Resource group: *resource-group-1* (for all resources)

Resources: *schost-1* (logical hostname), *SunONE-insecure-1* (insecure SunONE application resource), *SunONE-secure-1* (secure SunONE application resource)

(Add the resource group to contain all resources.)

```
# scrgadm -a -g resource-group-1
```

(Add the logical hostname resource to the resource group.)

```
# scrgadm -a -L -g resource-group-1 -l schost-1
```

(Register the SunONE resource type.)

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

(Add an insecure SunONE application resource instance.)

```
# scrgadm -a -j SunONE-insecure-1 -g resource-group-1 -t SUNW.iws \  
-x Confdir_list=/opt/SunONE/conf -y Scalable=False \  
-y Network_resources_used=schost-1 -y Port_list=80/tcp\  

```

(Add a secure SunONE application resource instance.)

```
# scrgadm -a -j SunONE-secure-1 -g resource-group-1 -t SUNW.iws \  
-x Confdir_List=/opt/SunONE/https-SunONE-secure-1 -y Scalable=False \  
-y Network_resources_used=schost-1 -y Port_list=443/tcp \  

```

(Bring the failover resource group online.)

```
# scswitch -Z -g resource-group-1
```

Where to Go From Here

To configure the `SUNW.HAStorage` resource type, see “How to Configure `SUNW.HAStoragePlus` Resource Type” on page 26.

▼ How to Configure `SUNW.HAStoragePlus` Resource Type

The `SUNW.HAStoragePlus` resource type was introduced in Sun Cluster 3.0 5/02. This new resource type performs the same functions as `SUNW.HAStorage`, and synchronizes actions between HA storage and the data service.

SUNW.HAStoragePlus also has an additional feature to make a local file system highly available. Sun Cluster HA for Sun ONE Web Server is scalable, and therefore you should set up the SUNW.HAStoragePlus resource type.

See the SUNW.HAStoragePlus(5) man page and “Relationship Between Resource Groups and Disk Device Groups” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for background information. See “Synchronizing the Startups Between Resource Groups and Disk Device Groups” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for the procedure. (If you are using a Sun Cluster 3.0 version prior to 5/02, you must set up SUNW.HAStorage instead of SUNW.HAStoragePlus. See “Synchronizing the Startups Between Resource Groups and Disk Device Groups” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for New Resources for the procedure.)

Configuring Sun Cluster HA for Sun ONE Web Server Extension Properties

This section describes the Sun Cluster HA for Sun ONE Web Server extension properties. For failover, the data service enforces that the size of `Confdir_list` is one. If you want multiple configuration files (instances), make multiple failover resources, each with one `Confdir_list` entry.

Typically, you use the command line `scrgadm -x parameter=value` to configure extension properties when you create the Sun ONE Web Server resource. You can also use the procedures described in Chapter 18, *Administering Data Service Resources* to configure them later. See “Standard Properties” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for details on all Sun Cluster properties.

Table 1–2 describes extension properties that you can configure for the Sun ONE server. The only required extension property when you create an Sun ONE server resource is the `Confdir_list` property. You can update some extension properties dynamically. You can update others, however, only when you create the resource. The Tunable entries indicate when you can update each property

TABLE 1-2 Sun Cluster HA for Sun ONE Web Server Extension Properties

Name/Data Type	Default
Confdir_list (string array)	<p>A pointer to the server root directory for a particular Sun ONE Web Server instance. If the Sun ONE Web Server is in secure mode, the path name must contain a file named <code>keypass</code>, which contains the secure key password needed to start this instance.</p> <p>Default: None</p> <p>Range: None</p> <p>Tunable: At creation</p>
Monitor_retry_count (integer)	<p>The number of times the process monitor facility (PMF) restarts the fault monitor during the time window that the <code>Monitor_retry_interval</code> property specifies. Note that this property refers to restarts of the fault monitor itself rather than to the resource. The system-defined properties <code>Retry_interval</code> and <code>Retry_count</code> control restarts of the resource.</p> <p>Default: 4</p> <p>Range:0 - 2, 147, 483, 641</p> <p>-1 indicates an infinite number of retry attempts.</p> <p>Tunable: Any time</p>
Monitor_retry_interval (integer)	<p>The time (in minutes) over which failures of the fault monitor are counted. If the number of times the fault monitor fails exceeds the value specified in the extension property <code>Monitor_retry_count</code> within this period, the PMF does not restart the fault monitor.</p> <p>Default: 2</p> <p>Range:0 - 2, 147, 483, 641</p> <p>-1 indicates an infinite retry interval.</p> <p>Tunable: Any time</p>

TABLE 1-2 Sun Cluster HA for Sun ONE Web Server Extension Properties (Continued)

Name/Data Type	Default
Probe_timeout (integer)	The time-out value (in seconds) that the fault monitor uses to probe an Sun ONE Web Server instance. Default: 90 Range: 0 – 2, 147, 483, 641 Tunable: Any time

Sun Cluster HA for Sun ONE Web Server Fault Monitor

The probe for Sun Cluster HA for Sun ONE Web Server uses a request to the server to query the health of that server. Before the probe actually queries the server, a check is made to confirm that network resources are configured for this Web server resource. If no network resources are configured, an error message (*No network resources found for resource*) is logged, and the probe exits with failure.

The probe must address the following two configurations of Sun ONE Web Server.

- the secure instance
- the insecure instance

If the Web server is in secure mode and if the probe cannot get the secure ports from the configuration file, an error message (*Unable to parse configuration file*) is logged, and the probe exits with failure. The secure and insecure instance probes involve common steps.

The probe uses the time-out value that the resource property `Probe_timeout` specifies to limit the time spent trying to successfully probe Sun ONE Web Server. See “Standard Properties” in *Sun Cluster 3.1 Data Service Planning and Administration Guide* for details on this resource property.

The `Network_resources_used` resource-property setting on the Sun ONE Web Server resource determines the set of IP addresses that the Web server uses. The `Port_list` resource-property setting determines the list of port numbers that Sun ONE Web Server uses. The fault monitor assumes that the Web server is listening on all combinations of IP and port. If you customize your Web server configuration to listen on different port numbers (in addition to port 80), ensure that your resultant configuration (`magnus.conf`) file contains all possible combinations of IP addresses and ports. The fault monitor attempts to probe all such combinations and might fail if the Web server is not listening on a particular IP address and port combination.

The probe executes the following steps.

1. The probe uses the specified IP address and port combination to connect to the Web server. If the connection is unsuccessful, the probe concludes that a complete failure has occurred. The probe then records the failure and takes appropriate action.
2. If the probe successfully connects, the probe checks if the Web server is run in a secure mode. If so, the probe disconnects and returns with a success status. No further checks are performed for a secure Sun ONE Web Server.

However, if the Web server is running in insecure mode, the probe sends an HTTP 1.0 HEAD request to the Web server and waits for the response. The request can be unsuccessful for various reasons, including heavy network traffic, heavy system load, and misconfiguration.

Misconfiguration can occur when the Web server is not configured to listen on all IP address and port combinations that are being probed. The Web server should service every port for every IP address specified for this resource.

Misconfigurations can also result if the `Network_resources_used` and `Port_list` resource properties are not set correctly while you create the resource.

If the reply to the query is not received within the `Probe_timeout` resource time limit, the probe considers this probe a failure of Sun Cluster HA for Sun ONE Web Server. The failure is recorded in the probe's history.

A probe failure can be a complete or partial failure. The following probe failures are considered complete failures.

- Failure to connect to the server, as the following error message flags, with `%s` indicating the host name and `%d` the port number.

```
Failed to connect to %s port %d
```
- Running out of time (exceeding the resource-property timeout `Probe_timeout`) after trying to connect to the server.
- Failure to successfully send the probe string to the server, as the following error message flags, with the first `%s` indicating the host name and `%d` the port number. The second `%s` indicates further details about the error.

```
Failed to communicate with server %s port %d: %s
```

The monitor accumulates two such partial failures within the resource-property interval `Retry_interval` and counts them as one failure.

The following probe failures are considered partial failures.

- Running out of time (exceeding the resource-property timeout `Probe_timeout`) while trying to read the reply from the server to the probe's query.
- Failing to read data from the server for other reasons, as the following error message flags, with the first `%s` indicating the host name and `%d` the port number. The second `%s` indicates further details about the error.

Failed to communicate with server %s port %d: %s

3. Based on the history of failures, a failure can cause either a local restart or a failover of the data service. This action is further described in “Sun Cluster Data Service Fault Monitors” in *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

Index

A

adding
 See modifying

C

changing
 See modifying

commands

prtconf -v, 9
prtdiag -v, 9
psrinfo -v, 9
showrev -p, 9

configuring

Sun Cluster HA for Sun ONE Web Server, 19
Sun ONE Web Server, 16

D

deleting
 See modifying
disabling
 See stopping
displaying
 See listing

E

editing
 See modifying
enabling
 See starting
executing
 See starting
extension properties
 Sun Cluster HA for Sun ONE Web Server
 Confdir_list, 29
 Monitor_retry_count, 29
 Probe_timeout, 29

F

fault monitor, Sun Cluster HA for Sun ONE
Web Server, 29

I

installing

- Sun Cluster HA for Sun ONE Web Server, 18
- Sun ONE Web Server, 14

P

printing

See listing

- prtconf -v command, 9
- psrinfo -v command, 9
- pv command, 9

R

registering, Sun Cluster HA for Sun ONE Web Server, 19

running

See enabling

S

Sun Cluster HA for Sun ONE Web Server

See also Sun ONE Web Server

configuration

planning, 12

configuring, 19

extension properties

Confdir_list, 29

Monitor_retry_count, 29

Probe_timeout, 29

fault monitor, 29

installation

installing software, 18

planning, 12

task map, 13

registering, 19

SUNW.HAStoragePlus resource type, 26

task map, 13

Sun ONE Web Server

See also Sun Cluster HA for Sun ONE Web Server

configuring, 16

installation

installing software, 14

SUNW.HAStoragePlus resource type, Sun

Cluster HA for Sun ONE Web Server, 26

T

task map, Sun Cluster HA for Sun ONE Web Server, 13

turning off

See disabling

turning on

See enabling

U

/usr/cluster/bin/scinstall -pv, 9