



Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS



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

Part No: 819-2990-10
December 2006, Revision A

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

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

This distribution may include materials developed by third parties.

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

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

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

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 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées 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.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Contents

Preface	5
Installing and Configuring Sun Cluster HA for Sun Java System Web Server	9
Planning the Installation and Configuration	10
Overview of the Installation and Configuration Process for Sun Cluster HA for Sun Java System Web Server	11
Installing and Configuring a Sun Java System Web Server	12
▼ How to Install a Sun Java System Web Server	12
▼ How to Configure a Sun Java System Web Server	13
Installing the Sun Cluster HA for Sun Java System Web Server Packages	16
▼ How to Install the Sun Cluster HA for Sun Java System Web Server Packages	16
Registering and Configuring Sun Cluster HA for Sun Java System Web Server	18
Setting Sun Cluster HA for Sun Java System Web Server Extension Properties	18
Monitoring Arbitrary URIs	18
Setting IP Addresses and Port Numbers for the Sun Java System Web Server Application	19
▼ How to Register and Configure Sun Cluster HA for Sun Java System Web Server in a Failover Configuration	19
▼ How to Register and Configure Sun Cluster HA for Sun Java System Web Server in a Scalable Configuration	23
How to Configure SUNW.HAStoragePlus Resource Type	28
Tuning the Sun Cluster HA for Sun Java System Web Server Fault Monitor	28
Operations by the Fault Monitor During a Probe	28
Upgrading the Sun Cluster HA for Sun Java System Web Server Resource Type	31
Information for Registering the New Resource Type Version	31
Information for Migrating Existing Instances of the Resource Type	31
A Sun Cluster HA for Sun Java System Web Server Extension Properties	33
Index	37

Preface

Sun Cluster Data Service for Sun Java System Web Server Guide for Solaris OS explains how to install and configure Sun™ Cluster HA for Sun Java System Web Server for Solaris OS on both SPARC® based systems and x86 based systems.

Note – This Sun Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, and AMD64. In this document, the label x86 refers to systems that use the AMD64 family of processor architectures.

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 book assume knowledge of the Solaris™ Operating System (Solaris OS) and expertise with the volume-manager software that is used with Sun Cluster software.

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

Using UNIX Commands

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

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>aabcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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

Related Documentation

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

Topic	Documentation
Data service administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides
Concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Overview	<i>Sun Cluster Overview for Solaris OS</i>
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
System administration	<i>Sun Cluster System Administration Guide for Solaris OS</i>
Hardware administration	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Data service development	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
Error messages	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
Command and function reference	<i>Sun Cluster Reference Manual for Solaris OS</i>

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

Related Third-Party Web Site References

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

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

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation](http://www.sun.com/documentation/) (<http://www.sun.com/documentation/>)
- [Support](http://www.sun.com/support/) (<http://www.sun.com/support/>)
- [Training](http://www.sun.com/training/) (<http://www.sun.com/training/>)

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the Solaris Operating System (for example, Solaris 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

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>SPARC: prtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</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 Java System Web Server

This chapter provides the procedures to install and configure Sun Cluster HA for Sun Java System Web Server.

This chapter contains the following sections.

- “Planning the Installation and Configuration” on page 10
- “Overview of the Installation and Configuration Process for Sun Cluster HA for Sun Java System Web Server” on page 11
- “Installing and Configuring a Sun Java System Web Server” on page 12
- “Installing the Sun Cluster HA for Sun Java System Web Server Packages” on page 16
- “Registering and Configuring Sun Cluster HA for Sun Java System Web Server” on page 18
- “Tuning the Sun Cluster HA for Sun Java System Web Server Fault Monitor” on page 28
- “Upgrading the Sun Cluster HA for Sun Java System Web Server Resource Type” on page 31

You can configure Sun Cluster HA for Sun Java System Web Server as a failover or scalable data service. See Chapter 1, “Planning for Sun Cluster Data Services,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* and the *Sun Cluster Concepts Guide for Solaris OS* for general information about data services, resource groups, resources, and other related topics.

Note – You can use Sun Cluster Manager to install and configure this data service. See the Sun Cluster 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 Java System Web Server depends on Sun Cluster HA for DNS, you must set up DNS first. See *Sun Cluster Data Service for DNS Guide for Solaris OS* for details. The Solaris Operating System 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 Java System Web Server except by using the cluster administration command `clresourcegroup(1CL)`. See the man page for details. After the Sun Java System Web Server is started, the Sun Cluster software controls it.

Planning the Installation and Configuration

Answer the following questions before you start your installation.

- Will you run Sun Cluster HA for Sun Java System Web Server as a failover or as a scalable data service? See the *Sun Cluster Concepts Guide for Solaris OS* 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 Java System 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 Appendix B, “Standard Properties,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

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 Appendix B, “Standard Properties,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 Appendix B, “Standard Properties,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 Java System Web Server to bind to `INADDR_ANY`, if you plan to run multiple instances of the Sun Java System 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 Java System Web Server in a Failover Configuration](#)” on page 19 for details.
- Determine if you will utilize the `Monitor_Uri_List` extension property. This extension property enables you to monitor an arbitrary list of URIs. Monitoring arbitrary URIs is beneficial if you will configure other data services that are accessible over the web. Use of the `Monitor_Uri_List` extension property is not supported with secure instances of Sun Java System Web Server. You must install release 3.1 10/03 or 3.2 Sun Cluster HA for Sun Java System Web Server to use this property. If you are upgrading Sun Cluster HA for Sun Java System Web Server from a previous version, you must perform the resource type upgrade procedure to use the new property. For instructions, see “[Upgrading a Resource Type](#)” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*. See “[Monitoring Arbitrary URIs](#)” on page 18 for detailed information about this option and example usage of `Monitor_Uri_List`.
- Determine where to place logs and error files on the local file system.
- Determine where to place the contents on the cluster file system.

Overview of the Installation and Configuration Process for Sun Cluster HA for Sun Java System Web Server

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

TABLE 1 Task Map: Installing and Configuring Sun Cluster HA for Sun Java System Web Server

Task	For Instructions
Install Sun Java System Web Server	“ Installing and Configuring a Sun Java System Web Server ” on page 12
Install the Sun Cluster HA for Sun Java System Web Server package	“ How to Install the Sun Cluster HA for Sun Java System Web Server Packages ” on page 16
Register Sun Cluster HA for Sun Java System Web Server and configure the cluster for the data service	“ Registering and Configuring Sun Cluster HA for Sun Java System Web Server ” on page 18
Tune the Sun Cluster HA for Sun Java System Web Server fault monitor	“ Tuning the Sun Cluster HA for Sun Java System Web Server Fault Monitor ” on page 28

Installing and Configuring a Sun Java System Web Server

This section describes the steps to perform the following tasks:

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

Note – The Sun Java System Web Server can be configured to run in a whole root or a sparse root non-global zone, if required.

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 Data Services Developer’s Guide for Solaris OS* to put the server under Sun Cluster control.

▼ How to Install a Sun Java System 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 Data Services Planning and Administration Guide for Solaris OS*.
- 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 Java System 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 On a cluster member, become superuser or assume a role that provides `solaris.cluster.admin` RBAC authorization.

2 Start the installation.

- If you are installing the Sun Java System Web Server on Solaris 8, execute the setup command for Sun Java System Web Server from the install directory on the CD.

- If you are installing the Sun Java System Web Server packaged with Solaris 9, follow the instructions on the installation CD.

Note – If you are installing the Sun Java System Web Server packaged with Solaris 9, do *not* enable the auto start of the web server at system restart.

3 When prompted, enter the location where the Sun Java System 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, install the web server on all 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 Java System 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 Java System 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 Java System install script selects for the administration server. You might want to use this default value later when you use the administration server to configure an instance of the Sun Java System Web Server. Otherwise, you can specify a different port number when you configure the Sun Java System 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.

▼ How to Configure a Sun Java System Web Server

This procedure describes how to configure an instance of the Sun Java System 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.

- Sun Cluster HA for Sun Java System Web Server now supports Sun Java System Proxy Server. For information about the Sun Java System Proxy Server product, see <http://docs.sun.com/db/prod/s1.webproxys>. For Sun Java System Proxy Server installation and configuration information, see http://docs.sun.com/db/coll/S1_ipwebproxysrvr36.
- 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 If you will configure Sun Java System Web Server as a scalable data service, create a directory on the local disk of all the nodes to hold the logs and error files that Sun Java System Web Server manages.

For the scalable configuration to work correctly, these files must be located on each node of the cluster, not on the cluster file system. Use shared storage only if you will configure Sun Java System Web Server as a failover data service.

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, 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 `admin-server`, then start the Sun Java System admin server.

```
# cd admin-server  
# ./startserv
```

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

The URL consists of the logical hostname and port number that the Sun Java System 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 Java System 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 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 NSS database password, protect the file with the appropriate permissions.

7 If you do not want to save the password in the `keypass` file, add the following to the `server.xml` file located in the server root directory.

```
<pkcs11>
  <enabled>true</enabled>
  <token>
    <name>internal</name>
    <pin>admin123</pin>
  </token>
</pkcs11>
```

Installing the Sun Cluster HA for Sun Java System Web Server Packages

If you did not install the Sun Cluster HA for Sun Java System Web Server packages during your initial Sun Cluster installation, perform this procedure to install the packages. To install the packages, use the Sun Java Enterprise System Common Installer.

▼ How to Install the Sun Cluster HA for Sun Java System Web Server Packages

Perform this procedure on each cluster node where you are installing the Sun Cluster HA for Sun Java System Web Server packages.

You can run the Sun Java Enterprise System Common Installer 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.

Note – Even if you plan to configure this data service to run in non-global zones, install the packages for this data service in the global zone. The packages are propagated to any existing non-global zones and to any non-global zones that are created after you install the packages.

Before You Begin Ensure that you have the Sun Java™ Availability Suite DVD-ROM.

If you intend to run the Sun Java Enterprise System Common Installer 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 Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

If the Volume Management daemon `vol(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 Sun Java Enterprise System Common Installer 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 Sun Java Enterprise System Common Installer.

```
# ./installer
```

5 When you are prompted, accept the license agreement.

If any Sun Java Enterprise System components are installed, you are prompted to select whether to upgrade the components or install new software.

6 From the list of Sun Cluster agents under Availability Services, select the data service for Sun Java System Web Server.**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 Sun Java Enterprise System Common Installer 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 Sun Java Enterprise System Common Installer with the CLI, omit this step

11 Exit the Sun Java Enterprise System Common Installer.**12 Unload the Sun Java Availability Suite DVD-ROM 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
```

Next Steps See [“Registering and Configuring Sun Cluster HA for Sun Java System Web Server”](#) on page 18 to register Sun Cluster HA for Sun Java System Web Server and to configure the cluster for the data service.

Registering and Configuring Sun Cluster HA for Sun Java System Web Server

You can configure Sun Cluster HA for Sun Java System Web Server as a failover data service or as a scalable data service. You must include some additional steps to configure Sun Java System 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.

Setting Sun Cluster HA for Sun Java System Web Server Extension Properties

The sections that follow contain instructions for registering and configuring Sun Cluster HA for Sun Java System Web Server resources. For information about the extension properties, see [Appendix A](#). The Tunable entry indicates when you can update a property.

See Appendix B, “Standard Properties,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for details on all the Sun Cluster properties.

To set an extension property of a resource, include the following option in the `cl resource(1CL)` command that creates or modifies the resource:

`-p property=value`

`-p property`

Identifies the extension property that you are setting.

`value`

Specifies the value to which you are setting the extension property.

You can also use the procedures in Chapter 2, “Administering Data Service Resources,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* to configure resources after the resources are created.

Monitoring Arbitrary URIs

Set the `Monitor Uri List` extension property if you want the web server fault monitor to probe an arbitrary list of applications (URIs) served by the web server. This extension property provides extended probing functionality and is useful if you are layering services in addition to your web server. The `Monitor Uri List` extension property is not supported with a secure Sun Java System Web Server instance. If you do not set the `Monitor Uri List` extension property, the fault monitor will perform basic probing. See “[Tuning the Sun Cluster HA for Sun Java System Web Server Fault Monitor](#)” on page 28 for details. The following examples show how to set the `Monitor Uri List` extension property when you add the Sun Java System Web Server instance to your configuration.

EXAMPLE 1 Setting `Monitor_Uri_List` for Scalable Sun Java System Web Server Instance

(Add an insecure Sun Java System Web Server instance with default load balancing.)

```
# clresource create -g resource-group-1 -t SUNW.iws \
-p Confdir_List=/opt/SUNWbsrv/https-Sun-app-insecure-1 \
-p Scalable=True -p Network_resources_used=schost-1 \
-p Port_list=8000/tcp \
-p Monitor_Uri_list=http://schost-1:8000/servlet/monitor web-not-secure-1
```

EXAMPLE 2 Setting `Monitor_Uri_List` for Failover Sun Java System Web Server Instance

(Add an insecure Sun Java System Web Server application resource instance.)

```
# clresource create -g resource-group-1 -t SUNW.iws \
-p Confdir_list=/opt/SUNWbsrv/conf -p Scalable=False \
-p Network_resources_used=schost-1 -p Port_list=80/tcp \
-p Monitor_Uri_list=http://schost-1:80/servlet/monitor web-not-secure-1
```

Setting IP Addresses and Port Numbers for the Sun Java System Web Server Application

The `Resource_dependencies` setting on the Sun Java System Web Server application resource determines the set of IP addresses that Sun Java System Web Server uses. The `Port_list` setting on the resource determines the list of port numbers that Sun Java System Web Server uses. The fault monitor assumes that the Sun Java System Web Server daemon is listening on all combinations of IP and port. You can configure the `server.xml` file for the Sun Java System Web Server to listen on different port numbers (in addition to port 80). The fault monitor attempts to probe all such combinations and starts to fail if the Sun Java System Web Server is not listening on a particular IP address-port combination. If the Sun Java System Web Server does not serve all IP address-port combinations, you must break the Sun Java System Web Server into separate instances that do.

▼ How to Register and Configure Sun Cluster HA for Sun Java System Web Server in a Failover Configuration

This procedure describes how to use the Sun Cluster maintenance commands to register and configure Sun Cluster HA for Sun Java System Web Server in a failover configuration.

Note – Other options also enable you to register and configure the data service. See “Tools for Data Service Resource Administration” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* 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 Java System 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 Java System 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 Data Services Planning and Administration Guide for Solaris OS* for a discussion of the advantages and disadvantages of each location.

Note – The `Network_resources_used` setting on the Sun Java System application resource determines the set of IP addresses that Sun Java System Web Server uses. The `Port_list` setting on the resource determines the list of port numbers that Sun Java System Web Server uses. The fault monitor assumes that the Sun Java System Web Server daemon is listening on all combinations of IP and port. You can configure the `server.xml` file for the Sun Java System Web Server to listen on different port numbers (in addition to port 80). The fault monitor attempts to probe all such combinations and starts to fail if the Sun Java System Web Server is not listening on a particular IP address-port combination. If the Sun Java System Web Server does not serve all IP address-port combinations, you must break the Sun Java System Web Server into separate instances that do.

Note – Perform this procedure on any cluster member.

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

2 Register the `SUNW.iws` resource type for Sun Cluster HA for Sun Java System Web Server.

```
# clresourcetype register SUNW.iws
```

3 Create a failover resource group.

You can optionally select the set of nodes or zones on which the data service can run with the `-n` option.

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

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.

`-n node-zone-list`

Specifies a comma-separated, ordered list of zones that can master this resource group. The format of each entry in the list is *node*. In this format, *node* specifies the node name and *zone* specifies the name of a non-global Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only *node*.

This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

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 Software Installation Guide for Solaris OS* 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 logical hostname to the failover resource group.

```
# clreslogicalhostname create -g resource-group \  
-h logical-hostname, . . . \  
[-p auxnodelist=node] [-N netiflist] resource
```

`-g resource-group`

Specifies the name of the failover resource group.

`-h logical-hostname, . . .`

Specifies a comma-separated list of logical hostnames that this resource is to make available.

`-p 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.

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.

`-N netiflist`

Specifies an optional, comma-separated list that identifies the IP Networking Multipathing groups that are on each node or zone. The format of each entry in the list is *netif@node*. The replaceable items in this format are as follows:

netif Specifies an IPMP group name, such as `sc_ipmp0`, or a public network interface card (NIC). If you specify a public NIC, Sun Cluster attempts to create the required IPMP groups.

node Specifies the name or ID of a node. To specify the global zone, or to specify a node without non-global zones, specify only *node*.

This list is optional. If you omit this list, Sun Cluster attempts to create the required IPMP groups.

Note – If you require a fully qualified hostname, you must specify the fully qualified name with the `-h` option and you cannot use the fully qualified form in the resource name.

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

resource

Specifies the name of the resource.

6 Create an application resource in the failover resource group.

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

```
# clresource create -g resource-group \  
-t resource-type \  
-p Confdir_list=config-directory \  
-p Network_resources_used=logical-hostname-list, . . . \  
-p Port_list=port-number/protocol resource
```

-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.

-p Confdir_list=config-directory

Specifies the location of the Sun Java System 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 Java System Web Server requires this extension property.

-p Network_resources_used=logical-hostname-list,...

Specifies a comma-separated list of network resources that identify the logical hosts that the data service uses.

-p 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`.

resource

Specifies the name of the resource to add.

The resource is created in the enabled state.

Note – Optionally, you can set additional extension properties that belong to the Sun Cluster HA for Sun Java System Web Server to override the default values of the properties. See [Appendix A](#) for a list of these properties.

7 Bring the failover resource group online.

```
# clresourcegroup online resource-group
```

resource-group

Specifies the name of the failover resource group.

Example 3 Registering a Failover Sun Cluster HA for Sun Java System Web Server 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), *Sun-app-insecure-1* (insecure application application resource), *Sun-app-secure-1* (secure application resource)

(Add the resource group to contain all resources.)

```
# clresourcegroup create resource-group-1
```

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

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

(Register the resource type for the Sun Cluster HA for Sun Java System Web Server.)

```
# clresourcetype register SUNW.iws
```

(Add an insecure application resource instance.)

```
# clresource create -g resource-group-1 -t SUNW.iws \
-p Confdir_list=/opt/SUNWbsrv/conf -p Scalable=False \
-p Network_resources_used=schost-1 -p Port_list=80/tcp Sun-app-insecure-1 \
```

(Add a secure application resource instance.)

```
# clresource create -g resource-group-1 -t SUNW.iws \
-p Confdir_List=/opt/SUNWbsrv/https-Sun-app-secure-1 -p Scalable=False \
-p Network_resources_used=schost-1 -p Port_list=443/tcp Sun-app-secure-1 \
```

(Bring the failover resource group online.)

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

▼ How to Register and Configure Sun Cluster HA for Sun Java System Web Server in a Scalable Configuration

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

- 2 Register the `SUNW.iws` resource type for Sun Cluster HA for Sun Java System Web Server.

```
# clresourcetype register SUNW.iws
```

- 3 Create a failover resource group for the shared address resource.

You can optionally select the set of nodes or zones on which the data service can run with the `-n` option.

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

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.

-n node-zone-list

Specifies a comma-separated, ordered list of zones that can master this resource group. The format of each entry in the list is *node*. In this format, *node* specifies the node name and *zone* specifies the name of a non-global Solaris zone. To specify the global zone, or to specify a node without non-global zones, specify only *node*.

This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

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 Software Installation Guide for Solaris OS* 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 shared address resource to the failover resource group.

```
# clressharedaddress create -g resource-group \  
-h shared-address,... \  
[-p auxnodelist=node] [-N netiflist] resource
```

-g resource-group

Specifies the name of the failover resource group.

-h shared-address,...

Specifies a comma-separated list of shared addresses that this resource is to make available.

-p 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.

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.

-N netiflist

Specifies an optional, comma-separated list that identifies the IP Networking Multipathing groups that are on each node or zone. The format of each entry in the list is *netif@node*. The replaceable items in this format are as follows:

- netif* Specifies an IPMP group name, such as `sc_ipmp0`, or a public network interface card (NIC). If you specify a public NIC, Sun Cluster attempts to create the required IPMP groups.
- node* Specifies the name or ID of a node. To specify the global zone, or to specify a node without non-global zones, specify only *node*.

This list is optional. If you omit this list, Sun Cluster attempts to create the required IPMP groups.

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

resource

Specifies the name of the resource.

6 Create a scalable resource group to run on all the desired cluster nodes.

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.

```
# clresourcegroup create \  
-p Maximum primaries=m -p Desired primaries=n \  
-p RG_dependencies=resource-group scalable-resource-group
```

-p *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.

-p *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.

-p *RG_dependencies= resource-group*

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

scalable-resource-group

Specifies the scalable resource group.

7 Create an application resource in the scalable resource group.

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

To set load balancing for the data service, use the two standard resource properties `Load_balancing_policy` and `Load_balancing_weights`. See Appendix B, “Standard Properties,” in

Sun Cluster Data Services Planning and Administration Guide for Solaris OS for a description of these properties. Additionally, see the examples that follow this section.

```
# clresource create -g scalable-resource-group \  
-t resource-type -p Confdir_list=config-directory,... \  
-p Resource_dependencies=shared-address,... \  
-p Port_list=port-number/protocol,... \  
-p Scalable=True resource
```

-g *scalable-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.

-p *Confdir_list=config-directory, ...*

Specifies a comma-separated list of the locations of the Sun Java System configuration files. Sun Cluster HA for Sun Java System Web Server requires this extension property.

-p *Resource_dependencies=shared-address,...*

Specifies a comma-separated list of network resources that identify the shared addresses that the data service uses.

-p *Port_list=port-number/protocol,...*

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

-p *Scalable=True*

Specifies a Boolean that is required for scalable services.

resource

Specifies the name of the resource to add.

The resource is created in the enabled state.

Note – A one-to-one mapping applies for *Confdir_List* and *Port_List*. Each value in one list must correspond to the value in the other list in the specified order.

8 Bring the failover resource group online.

```
# clresourcegroup online resource-group
```

resource-group

Specifies the name of the failover resource group.

9 Bring the scalable resource group online.

```
# clresourcegroup online resource-group
```

resource-group

Specifies the name of the scalable resource group.

Example 4 Registering a Scalable Sun Cluster HA for Sun Java System 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 application resources)**Resources: schost-1 (shared address), Sun-app-insecure-1 (insecure application resource), Sun-app-secure-1 (secure application resource)**(Add a failover resource group to contain shared addresses.)***# clresourcegroup create sa-resource-group-1***(Add the shared address resource to the failover resource group.)***# clressharedaddress create -g sa-resource-group-1 -h schost-1 sa-resource***(Add a scalable resource group.)***# clresourcegroup create -p Maximum primaries=2 \
-p Desired primaries=2 \
-p RG_dependencies=sa-resource-group-1 iws-resource-group-1***(Register the resource type for the Sun Cluster HA for Sun Java System Web Server.)***# clresourcetype register SUNW.iws***(Add an insecure application instance with default load balancing.)***# clresource create -g iws-resource-group-1 -t SUNW.iws \
-p Confdir_List=/opt/SUNWbsrv/https-Sun-app-insecure-1 \
-p Scalable=True -p Resource_dependencies=schost-1 \
-p Port_list=80/tcp Sun-app-insecure-1***(Add a secure application instance with sticky IP load balancing.)***# clresource create -g iws-resource-group-1 -t SUNW.iws \
-p Confdir_List=/opt/SUNWbsrv/https-Sun-app-secure-1 \
-p Scalable=True -p Resource_dependencies=schost-1 \
-p Port_list=443/tcp -p Load_balancing_policy=LB_STICKY \
-p Load_balancing_weights=40@1,60@2 Sun-app-secure-1***(Bring the failover resource group online.)***# clresourcegroup online sa-resource-group-1***(Bring the scalable resource group online.)***# clresourcegroup online iws-resource-group-1**

How to Configure SUNW.HASStoragePlus Resource Type

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

SUNW.HASStoragePlus also has an additional feature to make a local file system highly available.

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

Tuning the Sun Cluster HA for Sun Java System Web Server Fault Monitor

The Sun Cluster HA for Sun Java System Web Server fault monitor is contained in the resource that represents Sun Java System Web Server. You create this resource when you register and configure Sun Cluster HA for Sun Java System Web Server. For more information, see [“Registering and Configuring Sun Cluster HA for Sun Java System Web Server” on page 18](#).

System properties and extension properties of this resource control the behavior of the fault monitor. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Sun Cluster installations. Therefore, you should tune the Sun Cluster HA for Sun Java System Web Server fault monitor *only* if you need to modify this preset behavior.

For more information, see the following sections.

- “Tuning Fault Monitors for Sun Cluster Data Services” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
- “Changing Resource Type, Resource Group, and Resource Properties” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
- Appendix B, “Standard Properties,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*

Operations by the Fault Monitor During a Probe

The probe for Sun Cluster HA for Sun Java System 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 Java System Web Server.

- Secure instance
- Insecure instance

If the web server is in secure mode and if the probe cannot communicate with the secure ports, the 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 `Network_resources_used` resource-property setting on the Sun Java System 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 Java System Web Server uses. The fault monitor assumes that the web server is listening on all combinations of IP and port. 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.

If the probe fails to connect to the web server using a specified IP address and port combination, a complete failure occurs. The probe records the failure and takes appropriate action.

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 Java System 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 were not set correctly when you created 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 Java System 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. The following error message is sent, where %s indicates the host name and %d indicates the port number.

```
Failed to connect to %s port %d
```

- Timeout (exceeding the resource-property timeout `Probe_timeout`) after trying to connect to the server.
- Failure to successfully send the probe string to the server. The following error message is sent, where the first `%s` indicates the host name, `%d` indicates the port number, and 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.

- Timeout (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. The following error message is sent, where the first `%s` indicates the host name, `%d` indicates the port number, and the second `%s` indicates further details about the error.

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

The probe connects to the Sun Java System Web Server server and performs an HTTP 1.1 GET check by sending a HTTP request to each of the URIs in `Monitor Uri List`. If the HTTP server return code is 500 (Internal Server Error) or if the connect fails, the probe will take action.

The result of the HTTP requests is either failure or success. If all of the requests successfully receive a reply from the Sun Java System Web Server server, the probe returns and continues the next cycle of probing and sleeping.

Heavy network traffic, heavy system load, and misconfiguration can cause the HTTP GET probe to fail. Misconfiguration of the `Monitor Uri List` property can cause a failure if a URI in the `Monitor Uri List` includes an incorrect port or hostname. For example, if the web server instance is listening on logical host `schost - 1` and the URI was specified as `http://schost - 2/servlet/monitor`, the probe will try to contact `schost - 2` to request `/servlet/monitor`.

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 "Tuning Fault Monitors for Sun Cluster Data Services" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Upgrading the Sun Cluster HA for Sun Java System Web Server Resource Type

Upgrade the SUNW.iws resource type if the following conditions apply:

- You are upgrading from an earlier version of the Sun Cluster HA for Sun Java System Web Server data service.
- You need to use the new features of this data service.

For general instructions that explain how to upgrade a resource type, see “Upgrading a Resource Type” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*. The information that you require to complete the upgrade of the SUNW.iws resource type is provided in the subsections that follow.

Information for Registering the New Resource Type Version

The relationship between a resource type version and the release of Sun Cluster data services is shown in the following table. The release of Sun Cluster data services indicates the release in which the version of the resource type was introduced.

Resource Type Version	Sun ClusterData Services Release
1	1.0
3.1	3.1 5/03
4	3.1 10/03
5	3.2

To determine the version of the resource type that is registered, use the `clresourcetype show` command.

The resource type registration (RTR) file for this resource type is `/opt/SUNWschtt/etc/SUNW.iws`.

Information for Migrating Existing Instances of the Resource Type

The information that you require to edit each instance of the SUNW.iws resource type is as follows:

- You can perform the migration at any time.
- If you need to use the new features of the Sun Cluster HA for Sun Java System Web Server data service, the required value of the `Type_version` property is 4.

- If you need to monitor deployed applications, set the `Monitor Uri List` extension property to a single URI or a list of URIs to specify the locations of the applications that are to be probed.

The following example shows a command for modifying an instance of the `SUNW.iws` resource type.

EXAMPLE 5 Migrating Instances of the `SUNW.iws` Resource Type

```
# clresource set \  
  -p Monitor Uri List=http://schost-1/test.html \  
  -p Type_version=4 webserver-rs
```

This command modifies the `SUNW.iws` resource named `webserver-rs` as follows:

- The `Type_version` property of this resource is set to 4.
- The fault monitor probe will monitor the URI `http://schost-1/test.html`.

Sun Cluster HA for Sun Java System Web Server Extension Properties

This section describes the extension properties for the resource type `SUNW.iws`. This resource type represents the Sun Java System Web Server application in a Sun Cluster configuration.

For details about system-defined properties, see the `r_properties(5)` man page and the `rg_properties(5)` man page.

The extension properties of the `SUNW.iws` resource type are as follows:

`Confdir_list`

A pointer to the server root directory for a particular Sun Java System Web Server instance. If the Sun Java System Web Server is in secure mode, the path name must contain a file named `keypass`, which contains the secure key password needed to start this instance.

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.

Data type	String array
Default	No default defined
Range	Not applicable
Tunable	At creation

`Failover_enabled`

A specification of whether the fault monitor fails over the Sun Java System Web Server resource if the number of attempts to restart exceeds `Retry_count` within the time that `Retry_interval` specifies. The possible values of this extension property are as follows:

- `True` – Specifies that the fault monitor fails over the Sun Java System Web Server resource
- `False` – Specifies that the fault monitor does *not* fail over the Sun Java System Web Server resource

Data type	Boolean
Default	<code>True</code>

Range Not applicable
Tunable When resource is disabled

Monitor_retry_count

The number of times that the process monitor facility (PMF) restarts the fault monitor during the time window that the `Monitor_retry_interval` property specifies. This property refers to restarts of the fault monitor itself rather than to the resource. The system-defined properties `Retry_interval` and `Retry_count` control restarting of the resource.

Data type Integer
Default 4
Range 0 - 2,147,483,641
-1 indicates an infinite number of retry attempts.
Tunable At any time

Monitor_retry_interval

The time (in minutes) over which failures of the fault monitor are counted. If the number of times that the fault monitor fails exceeds the value that is specified in the extension property `Monitor_retry_count` within this period, the PMF does not restart the fault monitor.

Data type Integer
Default 2
Range 0 - 2,147,483,641
-1 indicates an infinite retry interval.
Tunable At any time

Monitor Uri_List

A single URI or a list of URIs that can be used by the fault monitor to probe any deployed applications on the Sun Java System Web Server. Probe deployed applications by setting the property to one or more URIs that are serviced by applications deployed on the Sun Java System Web Server.

Introduced in release 3.1 10/03.

Data type String
Default Null
Range Not applicable
Tunable At any time

Probe_timeout

The timeout value (in seconds) that the fault monitor uses to probe a Sun Java System Web Server instance.

Data type	Integer
Default	90
Range	0 – 2,147,483,641
Tunable	At any time

Index

C

`clnode` command, 8
commands, node information, 8
configuring, Sun Java System Web Server, 13

E

error messages, cluster, 8
extension properties
 `Monitor Uri List`, 11
 `SUNW.iws` resource type, 33-35

F

fault monitor, Sun Cluster HA for Sun Java System Web Server, 28
files, RTR, 31

G

global zone, 16

H

help, 8

I

installing
 Sun Cluster HA for Sun Java System Web Server, 16-17
 Sun Java System Web Server, 12
iPlanet Web Server, *See* Sun Java System Web Server

L

local zones, *See* non-global zones

M

messages file, 8
`Monitor Uri List` extension property, 11

N

Netscape™ HTTP, *See* Sun Java System Web Server
non-global zones, 16

P

packages, 16-17
properties
 See also extension properties
 `Monitor Uri List`, 32
 `Type version`, 31
`prtconf -v` command, 8
`prtdiag -v` command, 8

psrinfo -v command, 8

R

resource type registration (RTR) file, 31

restrictions, zones, 16

S

show- rev subcommand, 8

showrev -p command, 8

software packages, 16-17

Sun Cluster HA for Sun Java System Web Server

configuration

planning, 10

fault monitor, 28

installation

planning, 10

task map, 11

installing, 16-17

resource type versions, 31

software packages

installing, 16-17

SUNW.HAStoragePlus resource type, 28

task map, 11

Sun Java System Web Server

configuring, 13

installation

installing software, 12

Sun ONE Web Server, *See* Sun Java System Web Server

SUNW.HAStoragePlus resource type, Sun Cluster HA for

Sun Java System Web Server, 28

SUNW.iws resource type, extension properties, 33-35

T

task map, Sun Cluster HA for Sun Java System Web
Server, 11

technical support, 8

Type_version property, 31

V

/var/adm/messages file, 8

versions, resource types, 31

Z

zones, 16