

Oracle® Solaris Cluster Data Service for Oracle WebLogic Server Guide

Copyright © 2010, 2012, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Contents

Preface	5
1 Installing and Configuring Oracle Solaris Cluster HA for WebLogic Server	9
Oracle Solaris Cluster HA for WebLogic Server Overview	9
Overview of the Installation and Configuration Process for Oracle Solaris Cluster HA for WebLogic Server	10
Planning the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration	11
Configuration Restrictions	11
Configuration Requirements	12
Standard Data Service Configurations	14
Preparing the Nodes and Disks	18
▼ How to Prepare the Nodes for a Failover Configuration	18
▼ How to Prepare the Nodes for a Multi-Master Configuration	19
Installing and Configuring the WebLogic Server Application	20
▼ How to Install WebLogic Server	20
Installing a Database	22
Installing a Web Server	22
Verifying the WebLogic Server Installation and Configuration	22
▼ How to Verify the WebLogic Server Installation and Configuration for Failover Services	22
▼ How to Verify the WebLogic Server Installation and Configuration for Multi-Master Services	23
Installing the HA for WebLogic Server Package	24
▼ How to Install the HA for WebLogic Server Package	24
Registering and Configuring the Oracle Solaris Cluster HA for WebLogic Server	25
Setting Oracle Solaris Cluster HA for WebLogic Server Extension Properties	25
▼ How to Register and Configure the Oracle Solaris Cluster HA for WebLogic Server as a Web Proxy Server	26

▼ How to Register and Configure Oracle Solaris Cluster HA for WebLogic Server as a Failover Data Service	27
▼ How to Register and Configure Oracle Solaris Cluster HA for WebLogic Server as a Multi-Master Data Service	28
Setting Up the WebLogic Server for HAStoragePlus Configuration	29
▼ How to Set Up the WebLogic Server on Zone Clusters for HAStoragePlus Configuration	29
Verifying the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration	31
▼ How to Verify the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration for Failover Services	31
▼ How to Verify the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration in a Multi-Master Configuration	32
Alternate Oracle Solaris Cluster HA for WebLogic Server Configurations	33
Creating a Simple WebLogic Server Resource	33
Creating a Resource With a Server Name Argument	33
Creating a Resource That Shuts Down Smoothly	34
Creating a Resource That Probes the Database	35
Creating a Resource That Monitors URIs	35
Using a Non-Clustered Managed Server Instance as a Proxy Server	36
Operation of the Oracle Solaris Cluster HA for WebLogic Server Fault Monitor	36
Probing Algorithm and Functionality	36
A HA for WebLogic Server Extension Properties	39
Index	43

Preface

Oracle Solaris Cluster Data Service for WebLogic Server Guide explains how to install and configure Oracle Solaris Cluster data services.

Note – This Oracle Solaris Cluster release supports systems that use the SPARC and x86 families of processor architectures. In this document, “x86” refers to the larger family of x86 compatible products. Information in this document pertains to all platforms unless otherwise specified.

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

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

Bash is the default shell for Oracle Solaris 11. Machine names shown with the Bash shell prompt are displayed for clarity.

Using UNIX Commands

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

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-2 Shell Prompts

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

Related Documentation

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

Topic	Documentation
Hardware installation and administration	<i>Oracle Solaris Cluster 4.0 Hardware Administration Manual</i> Individual hardware administration guides
Concepts	<i>Oracle Solaris Cluster Concepts Guide</i>
Software installation	<i>Oracle Solaris Cluster Software Installation Guide</i>
Data service installation and administration	<i>Oracle Solaris Cluster Data Services Planning and Administration Guide</i> and individual data service guides
Data service development	<i>Oracle Solaris Cluster Data Services Developer's Guide</i>
System administration	<i>Oracle Solaris Cluster System Administration Guide</i> <i>Oracle Solaris Cluster Quick Reference</i>
Software upgrade	<i>Oracle Solaris Cluster Upgrade Guide</i>
Error messages	<i>Oracle Solaris Cluster Error Messages Guide</i>
Command and function references	<i>Oracle Solaris Cluster Reference Manual</i> <i>Oracle Solaris Cluster Data Services Reference Manual</i> <i>Oracle Solaris Cluster Geographic Edition Reference Manual</i> <i>Oracle Solaris Cluster Quorum Server Reference Manual</i>

Access to Oracle Support

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

Getting Help

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

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the operating environment (for example, Oracle Solaris 11)
- The release number of Oracle Solaris Cluster (for example, Oracle Solaris Cluster 4.0)

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

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

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

Installing and Configuring Oracle Solaris Cluster HA for WebLogic Server

This chapter explains how to install and configure Oracle Solaris Cluster HA for WebLogic Server.

This chapter contains the following sections.

- “Oracle Solaris Cluster HA for WebLogic Server Overview” on page 9
- “Overview of the Installation and Configuration Process for Oracle Solaris Cluster HA for WebLogic Server” on page 10
- “Planning the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration” on page 11
- “Preparing the Nodes and Disks” on page 18
- “Installing and Configuring the WebLogic Server Application” on page 20
- “Verifying the WebLogic Server Installation and Configuration” on page 22
- “Installing the HA for WebLogic Server Package” on page 24
- “Registering and Configuring the Oracle Solaris Cluster HA for WebLogic Server” on page 25
- “Setting Up the WebLogic Server for HAStoragePlus Configuration” on page 29
- “Verifying the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration” on page 31
- “Alternate Oracle Solaris Cluster HA for WebLogic Server Configurations” on page 33
- “Operation of the Oracle Solaris Cluster HA for WebLogic Server Fault Monitor” on page 36

Oracle Solaris Cluster HA for WebLogic Server Overview

Use the information in this section to understand how to make the WebLogic Server application highly available.

WebLogic Server running on Oracle Solaris Cluster systems delivers a highly available platform for developing and deploying mission-critical e-commerce applications across distributed, heterogeneous application environments.

The Oracle Solaris Cluster HA for WebLogic Server provides fault monitoring and high availability for the WebLogic Server application. High availability is provided for the WebLogic Administration Server and WebLogic Managed Servers.

Oracle Solaris Cluster HA for WebLogic Server enables the Managed Server component of the WebLogic Server clustering solution to be highly available in an Oracle Solaris Cluster system. To achieve this, you must configure the data service to be mastered by multiple nodes.

For conceptual information on failover and multi-master services, see the [Oracle Solaris Cluster Concepts Guide](#).

TABLE 1-1 Protection of WebLogic Server Components

WebLogic Server Component	Protected by
WebLogic Server	Oracle Solaris Cluster HA for WebLogic Server The resource type is SUNW.wls.
WebLogic Server database	All databases supported by WebLogic Server are supported on Oracle Solaris Cluster
HTTP servers	All HTTP servers supported by WebLogic Server are supported on Oracle Solaris Cluster

Overview of the Installation and Configuration Process for Oracle Solaris Cluster HA for WebLogic Server

The table below lists the tasks for installing and configuring Oracle Solaris Cluster HA for WebLogic Server. Perform these tasks in the order that they are listed.

TABLE 1-2 Task Map: Installing and Configuring Oracle Solaris Cluster HA for WebLogic Server

Task	For Instructions
Plan the WebLogic Server installation	“Planning the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration” on page 11
Prepare the nodes and disks	“How to Prepare the Nodes for a Failover Configuration” on page 18
Install and configure WebLogic Server	“How to Install WebLogic Server” on page 20 “Installing a Database” on page 22 “Installing a Web Server” on page 22

TABLE 1-2 Task Map: Installing and Configuring Oracle Solaris Cluster HA for WebLogic Server
(Continued)

Task	For Instructions
Verify WebLogic Server installation and configuration	<ul style="list-style-type: none"> ■ For failover services, see “How to Verify the WebLogic Server Installation and Configuration for Failover Services” on page 22 ■ For multi-master services, see “How to Verify the WebLogic Server Installation and Configuration for Multi-Master Services” on page 23
Install Oracle Solaris Cluster HA for WebLogic Server packages	“How to Install the HA for WebLogic Server Package” on page 24
Register and configure Oracle Solaris Cluster HA for WebLogic Server	<ul style="list-style-type: none"> ■ For failover services, see “How to Register and Configure Oracle Solaris Cluster HA for WebLogic Server as a Failover Data Service” on page 27 ■ For multi-master services, see “How to Register and Configure Oracle Solaris Cluster HA for WebLogic Server as a Multi-Master Data Service” on page 28
Verify Oracle Solaris Cluster HA for WebLogic Server installation and configuration	<ul style="list-style-type: none"> ■ For failover services, see “How to Verify the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration for Failover Services” on page 31 ■ For multi-master services, see “How to Verify the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration in a Multi-Master Configuration” on page 32
View Oracle Solaris Cluster HA for WebLogic Server Fault Monitor	“Operation of the Oracle Solaris Cluster HA for WebLogic Server Fault Monitor” on page 36

Planning the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration

This section contains the information you need to plan your Oracle Solaris Cluster HA for WebLogic Server installation and configuration.

Configuration Restrictions



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

This section provides a list of software and hardware configuration restrictions that apply to Oracle Solaris Cluster HA for WebLogic Server. Use the restrictions in this section to plan the installation and configuration of Oracle Solaris Cluster HA for WebLogic Server.

For restrictions that apply to all data services, see the [Oracle Solaris Cluster 4.0 Release Notes](#).

- Oracle Solaris Cluster HA for WebLogic Server provides high availability to stand alone WebLogic Server instances (instances not using WebLogic Server clustering).
- For a failover file system, all WebLogic Server instances using that file system must be configured in a single resource group.
- For a multi-master configuration, the WebLogic software must be installed in the local file system on all Oracle Solaris Cluster nodes. Installing the software locally avoids data corruption or loss when the WebLogic Server instances access the file system simultaneously.
- If you want to configure different instances in different resource groups, the WebLogic Server installation must be on a global file system.
- If the WebLogic Server cluster directory and other configuration files are installed on a global file system and you want to configure WebLogic Server in a zone cluster, use the `SUNW.HASStoragePlus` resource type to make a global file system available to zone cluster nodes. For information on the `HASStoragePlus` configuration for a zone cluster, see [Oracle Solaris Cluster Software Installation Guide](#).

Configuration Requirements



Caution – Your data service configuration might not be supported if you do not adhere to these requirements.

Use the requirements in this section to plan the installation and configuration of Oracle Solaris Cluster HA for WebLogic Server. These requirements apply to Oracle Solaris Cluster HA for WebLogic Server only. You must meet these requirements before you proceed with your Oracle Solaris Cluster HA for WebLogic Server installation and configuration.

For requirements that apply to all data services, see “[Configuration Guidelines for Oracle Solaris Cluster Data Services](#)” in [Oracle Solaris Cluster Data Services Planning and Administration Guide](#).

- If the WebLogic Server home directory and other configuration files are installed on a failover file system using `HASStoragePlus`, then all WebLogic Server resources using that file system should be configured in the same resource group.
- If the WebLogic Server depends on any database such as Oracle or Sybase, the administrator must make sure that the database is highly available before configuring and starting the WebLogic Server under Oracle Solaris Cluster HA for WebLogic Server.
- In WebLogic Server, the Managed Servers can be configured to start independently in “Managed Server Independence mode.” For Oracle Solaris Cluster HA for WebLogic Server, the administrator must configure all Managed Servers to be able to start independently.

To configure the Managed Server in the independent mode, see the relevant WebLogic Server documentation for the version you are using. For example, for WebLogic Server 10.3, the documentation is located at http://download.oracle.com/docs/cd/E12840_01/wls/docs103/sitemap.html.

- You can configure separate scripts to start the WebLogic Server Managed Servers. If separate scripts are configured, the Administration Server URL and the Managed Server name must be set in the START script. The scripts should be able to start each WebLogic Server instance without arguments.
- If you intend to use a single START script to start several Managed Servers, the Administration URL must be specified in the START script and only the Managed Server name should be taken as the argument to the script. For example, if the script used to start the Managed Server is `startManagedWeblogic.sh`, then the Managed Server `petstore_server` should be able to start as follows:

```
node1>cd /global/beahome/weblogic700/domain/petstore
node1>./startManagedWeblogic.sh petstore_server
```

- The WebLogic Server START script needs a user name and password. You configure the `$DOMAIN_DIR/boot.properties` file to store the user name and password which are encrypted after the first startup of the WebLogic Server. The `$DOMAIN_DIR/boot.properties` file also enables a smooth shutdown when used with a resource for which the extension property `smooth_shutdown` is set to `TRUE`.

Note – Starting with WebLogic Server version 10.3, the `boot.properties` file must reside in both the `$DOMAIN_DIR` and `$DOMAIN_DIR/servers/server-name/security` directories to bypass the need to prompt for user credentials when starting and stopping instances of WebLogic Server. If the `$DOMAIN_DIR/servers/server-name/security` directory does not exist or does not include a `boot.properties` file, create the directory and copy the `boot.properties` file from the `$DOMAIN_DIR`. For more information about the `boot.properties` file for WebLogic Server version 10.3, see http://download.oracle.com/docs/cd/E12840_01/wls/docs103/server_start/overview.html.

- All required environment variables must be set in the WebLogic Server environment file before you start the server. For details about the WebLogic Server environment file, see the relevant WebLogic Server documentation for the version you are using. For example, for WebLogic Server 10.3, the documentation is located at http://download.oracle.com/docs/cd/E12840_01/wls/docs103/sitemap.html. To set the path to the WebLogic Server environment file, use the `Environment_File` property described in [Appendix A, “HA for WebLogic Server Extension Properties.”](#)
- To enable the discovery of running Managed Servers after a failover and restart of an Administration Server, make sure the Administration server START script does not include `-Dweblogic.management.discover+=false`. This setting disables the Administration Server from discovering its running Managed Servers. For details, see the relevant

WebLogic Server documentation for the version you are using. For example, for WebLogic Server 10.3, the documentation is located at http://download.oracle.com/docs/cd/E12840_01/wls/docs103/sitemap.html.

- Make sure that the host names used in the `server_url` and `monitor_uri_list` are resolvable. If you use Fully Qualified Domain Names (FQDNs), then DNS must be enabled and `/etc/nsswitch.conf` must have the correct entries to resolve the host names using DNS.

Note – To avoid any failures due to delays in name-service lookup, the host name used in the `server_url` property should be added to the `/etc/hosts` file. Configure name-service mapping in the `/etc/nsswitch.conf` file on the servers to first check the local files before trying to access other name services.

The following requirements apply specifically to the multi-master configuration:

- The WebLogic Server managed server instances must be configured to listen on "All Local Addresses" while configuring the domain. All WebLogic Server instances must listen on the same port. While configuring the WebLogic Server domain, configure machines by entering each node name as a machine name and assigning each WebLogic Server server to each machine.

Standard Data Service Configurations

Use the standard configurations in this section to plan the installation and configuration of Oracle Solaris Cluster HA for WebLogic Server. Oracle Solaris Cluster HA for WebLogic Server supports the standard configurations in this section.

In addition to the standard configurations provided in this section, the following WebLogic Server WebLogic Server configurations are possible.

- Separate START script for each WebLogic Server (Administration and Managed)
- Single START script for all Managed Servers which takes the Managed Server name as the input
- Setting the user name and password in the `$(DOMAIN_DIR)/boot.properties` file
This configuration enables a smooth shutdown.

Note – Starting with WebLogic Server version 10.3, the `boot.properties` file must reside in both the `$DOMAIN_DIR` and `$DOMAIN_DIR/servers/server-name/security` directories to bypass the need to prompt for user credentials when starting and stopping instances of WebLogic Server. If the `$DOMAIN_DIR/servers/server-name/security` directory does not exist or does not include a `boot.properties` file, create the directory and copy the `boot.properties` file from the `$DOMAIN_DIR`.

Note – The WebLogic Server can be configured either as a failover data service or as a multi-master data service.

Oracle Solaris Cluster HA for WebLogic Server might support additional configurations. However, you must contact your Oracle service provider for information about additional configurations.

WebLogic Server Configured as a Failover Resource

Simple Configuration

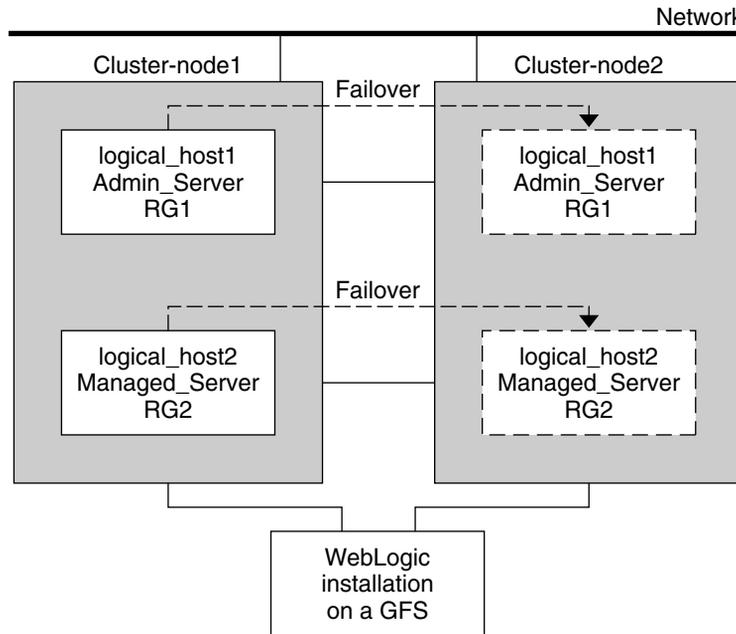
This simple configuration for Oracle Solaris Cluster HA for WebLogic Server uses two resource groups configured as follows.

- RG1 containing `logical-host-1` and `WLS-Admin-Server-resource`
- RG2 containing `logical-host-2` and `WLS-Managed-Server-resource`

RG1 and RG2 are failover resource groups.

This configuration is possible only if the WebLogic Server is installed on a global file system. See [Figure 1–1](#) for an illustration of this configuration.

FIGURE 1-1 Oracle Solaris Cluster HA for WebLogic Server Simple Configuration



Configuration 1

In this configuration, RG1, RG2, RG3, and RG4 are failover resource groups configured as follows.

- RG1 containing logical-host-1 and WLS-Admin-resource
- RG2 containing logical-host-2 and WLS-Managed-Server-resource2
- RG3 containing logical-host-3 and WLS-Managed-Server-resource3
- RG4 containing logical-host-4 and WLS-Managed-Server-resource4

This configuration is possible only if the WebLogic Server is installed on a global file system.

Configuration 2

In this configuration, RG1, RG2, and RG3 are failover resource groups configured as follows.

- RG1 containing logical-host-1, WLS-Admin-resource, and WLS-Managed-Server-resource1
- RG2 containing logical-host-2, WLS-Managed-Server-resource2, WLS-Managed-Server-resource3, and WLS-Managed-Server-resource4
- RG3 containing logical-host-3, WLS-Managed-Server-resource5, WLS-Managed-Server-resource6, WLS-Managed-Server-resource7

This configuration is possible only if the WebLogic Server is installed on a global file system.

Configuration 3

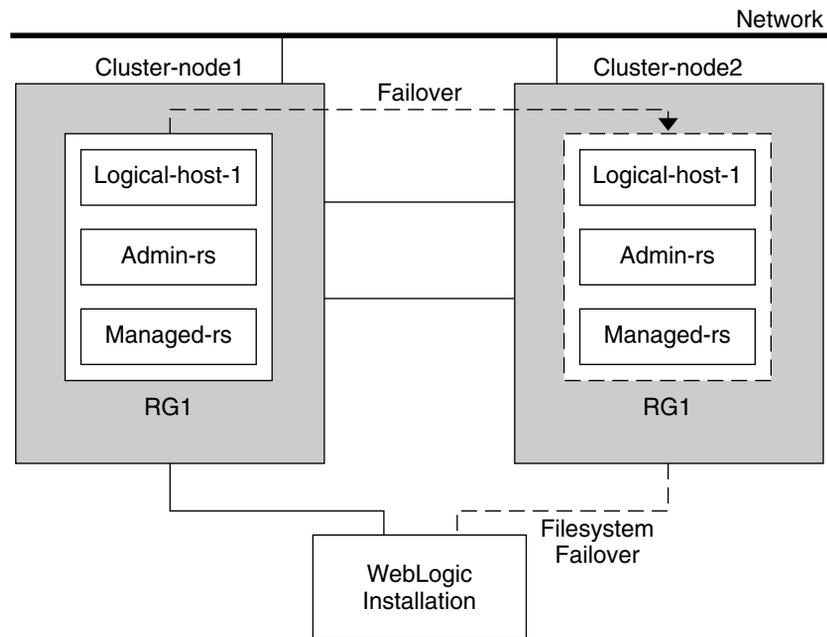
If the WebLogic Server is installed on a failover file system by using HAStoragePlus, all the WebLogic Server instances must be configured in the same failover resource group that has a dependency on the failover file system.

RG1 contains logical-host-resources and all Administration and Managed-Server resources.

For information about how to register a resource with dependency on an HAStoragePlus resource, see “[Enabling Highly Available Local File Systems](#)” in *Oracle Solaris Cluster Data Services Planning and Administration Guide*.

See [Figure 1–2](#) for an illustration of this configuration.

FIGURE 1–2 Oracle Solaris Cluster HA for WebLogic Server File System Failover Configuration

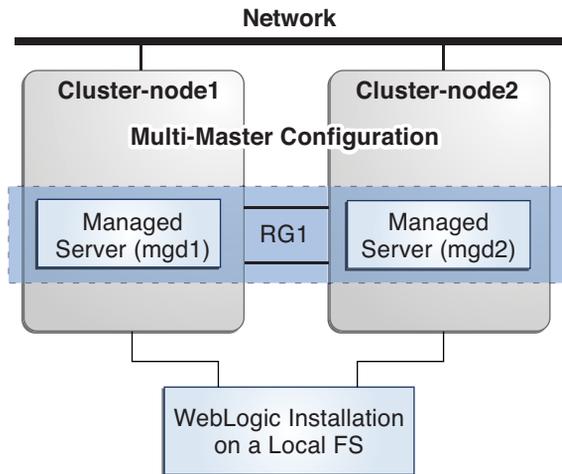


WebLogic Server Configured as a Multi-Master Resource

This example illustrates a multi-master configuration for Oracle Solaris Cluster HA for Oracle WebLogic Server (configured as a cluster). In this configuration, a single resource group is configured as follows:

- RG1 containing Managed Server mgd1 and Managed Server mgd2

FIGURE 1-3 Oracle WebLogic Server Configured as a Multi-Master Resource



Preparing the Nodes and Disks

This section contains the procedures that you need to prepare the nodes and disks.

Note – The following steps are based on the configurations that are described in “[Standard Data Service Configurations](#)” on page 14.

▼ How to Prepare the Nodes for a Failover Configuration

Use this procedure to prepare for the installation and configuration of WebLogic Server in a failover configuration.

- 1 **Create a failover resource group to hold the logical hostname and the WebLogic Administration Server resources.**

```
# clresourcegroup create [-n nodelist] RG1
```

- 2 **Add the logical host name resource.**

The *logical-host-1* is the host name the WebLogic Administration Server will be listening on. This host name can be set from the WebLogic console.

```
# clreslogicalhostname create -g RG1 logical-host-1
```

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

3 Bring the resource group online.

```
# clresourcegroup online -M RG1
```

4 Create a failover resource group to hold the logical host name and the WebLogic Managed Server resources.

Note – Depending on the number of WebLogic Servers and on the configuration planned, create the resource groups and the *logical_hostname* resources. Refer to “[Standard Data Service Configurations](#)” on page 14.

```
# clresourcegroup create [-n nodelist] RG2
```

5 Add the logical host name resource.

The *logical-host-2* is the host name on which the WebLogic Administration Server will be listening.

```
# clreslogicalhostname create -g RG2 logical-host-2
```

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

6 Bring the resource group online.

```
# clresourcegroup online -M RG2
```

▼ How to Prepare the Nodes for a Multi-Master Configuration

Use this procedure to prepare for the installation and configuration of WebLogic Server in a multi-master configuration.

1 Create a scalable resource group for managed servers.

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

Use the following information:

- `-p Maximum primaries=m` specifies the maximum number of active primary nodes permitted for this resource group. The default is 1.
- `-p Desired primaries=n` specifies the desired number of active primary nodes for this resource group. The default is 1.
- `scalable-resource-group` is the resource group name.

2 Bring the resource group online

```
# clresourceoup online -M scalable-resource-group
```

Installing and Configuring the WebLogic Server Application

This section contains the procedures that you need to install and configure the WebLogic Server application.

▼ How to Install WebLogic Server

Note – This procedure is based on the configuration that is described in [“Simple Configuration” on page 15](#). If you are planning a different configuration, you must configure the WebLogic Server accordingly.

- 1 **Follow the WebLogic Server documentation to install the WebLogic Server on the global file system or a failover file system.**
- 2 **Follow the *WebLogic Server Administration Guide* to configure the WebLogic Domain.**

Note – For failover services, type the logical host name or its IP address when configuring the WebLogic Servers.

For multi-master configurations, ensure that you are listening on All Local Addresses when you configure the WebLogic Servers.

- 3 **After completing all necessary configuration for the domain, shut down the Administration Server.**
- 4 **Edit the Administration Server START script and the Managed Server START script by setting the WLS_USER and WLS_PW.**

The START scripts should be able to run without any arguments.

Note – To perform a smooth shutdown of the WebLogic Server, the `boot.properties` file with a valid user name and password is required. You do not need to set the `WLS_USER` and `WLS_PW` variables for a smooth shutdown if the `boot.properties` file is available. This waiver applies only to WebLogic Server 8.1 or later versions.

Note – Starting with WebLogic Server version 10.3, the `boot.properties` file must reside in both the `$DOMAIN_DIR` and `$DOMAIN_DIR/servers/server-name/security` directories to bypass the need to prompt for user credentials when starting and stopping instances of WebLogic Server. If the `$DOMAIN_DIR/servers/server-name/security` directory does not exist or does not include a `boot.properties` file, create the directory and copy the `boot.properties` file from the `$DOMAIN_DIR`. For more information about the `boot.properties` file for WebLogic Server version 10.3, see http://download.oracle.com/docs/cd/E12840_01/wls/docs103/server_start/overview.html.

5 In the Managed Server START script, set the ADMIN_URL and the SERVER_NAME variables.

This script should be able to run without any arguments.

6 Set the necessary environment variables in the setWLSEnv.sh or the startWLS.sh scripts.

These environment variables include any additional `CLASSPATH` or `LD_LIBRARY_PATH`, etc. These scripts contain all necessary environmental variables such as `CLASSPATH` and `LD_LIBRARY_PATH`.

Note – Starting with WebLogic Server version 10.3, the `startWLS.sh` script has been removed. Therefore, these environment settings must be made only in the `setWLSEnv.sh` script.

7 For versions of WebLogic Server before 10.3, verify that the config.xml file is in the same directory as the START scripts.

8 For a multi-master configuration, replicate the domain and Managed Servers across the cluster.

- a. Use the `pack.sh` command to create a template of the domain and Managed Servers that you created.
- b. Copy the template to other Oracle Solaris Cluster nodes.
- c. Use the `unpack.sh` command to unpack the template to create the domain and Managed Servers.

For more information, see http://download.oracle.com/docs/cd/E12840_01/wls/docs103/server_start/overview.html.

Installing a Database

If you are using a database with WebLogic Server, install the database by using its installation documentation. If the database that you are using has a highly available agent on Oracle Solaris Cluster, install that version.

Installing a Web Server

If you are using a web server with WebLogic Server, install the web server by using its installation documentation. If the web server that you are using has a highly available agent on Oracle Solaris Cluster, install that version.

Verifying the WebLogic Server Installation and Configuration

This section contains the procedure that you need to verify the WebLogic Server installation and configuration.

▼ How to Verify the WebLogic Server Installation and Configuration for Failover Services

Use this procedure to verify the WebLogic Server Administration Server, WebLogic Managed Server, and database installation and configuration. This procedure does not verify that your application is highly available because you have not installed the data service yet.

Note – The following steps are for the failover configuration that is described in [“Standard Data Service Configurations” on page 14](#). If you are planning a different configuration, you must configure the WebLogic Server accordingly.

- 1 **Verify that the *logical-host-1* and *logical-host-2* configured in RG1 and RG2 are online on *cluster-node1*.**
- 2 **If you are using a database, verify that the database is running.**
- 3 **Go to the directory that contains the Administration Server and the Managed Server START scripts configured in the section [“Installing and Configuring the WebLogic Server Application” on page 20](#).**
- 4 **Start the Administration Server by running the Administration Server START script.**

```
example:  
cluster-node1>cd /global/wls/user_projects/petstore  
cluster-node1>./startWebLogic.sh
```

- 5 After the Administration Server starts successfully, start the Managed Server.

```
example:
cluster-node1>cd /global/wls/user_projects/petstore
cluster-node1>./startManagedWebLogic.sh
```

- 6 Make sure that you can successfully connect to the *logical-host-1* and *logical-host-2* from a client.

- 7 Verify that the Administration Server and the Managed Servers are in the RUNNING mode on *cluster-node1*.

- 8 Shut down the Administration Server and the Managed Server.

- 9 Switch the Resource Groups RG1 and RG2 to *cluster-node2*.

This action brings down *logical-host-1* and *logical-host-2* on *cluster-node1* and start them on *cluster-node2*.

```
clresourcegroup switch -n cluster-node2 RG1
clresourcegroup switch -n cluster-node2 RG2
```

- 10 Repeat [Step 4](#), [Step 5](#), and [Step 6](#) on *cluster-node2*.

- 11 Verify that the Administration Server and the Managed Servers are in the RUNNING mode on *cluster-node2*.

- 12 If the clients can successfully connect to the Administration Server and the Managed Server on each of the potential master nodes, you have successfully configured the WebLogic Servers to work with the Oracle Solaris Cluster HA for WebLogic Server.

- 13 Stop the Administration Server and the Managed Servers and proceed to the next section to install the Oracle Solaris Cluster HA for WebLogic Server.

▼ How to Verify the WebLogic Server Installation and Configuration for Multi-Master Services

Use this procedure to verify the WebLogic Server Administration Server, WebLogic Managed Server, and database installation and configuration for multi-master services. This procedure does not verify that your application is highly available because you have not installed the data service yet.

Note – The following steps are for the multi-master configuration that is described in [“Standard Data Service Configurations” on page 14](#). If you are planning a different configuration, you must configure the WebLogic Server accordingly.

- 1 If you are using a database, verify that the database is running.

- 2 **Go to the directory that contains the Administration Server and the Managed Server START scripts.**

- 3 **Start the Administration Server by running the Administration Server START script.**

For example:

```
cluster-node1>cd /mnt/user_projects/domains/base_domain
cluster-node1>./startWebLogic.sh
```

- 4 **After the Administration Server starts successfully, start the Managed Server mgd1 on phys-host-1.**

For example:

```
cluster-node1>cd /mnt/user_projects/domains/base_domain_mgd1
cluster-node1>./startManagedWebLogic.sh mgd1
```

- 5 **Verify that the Administration Server and the Managed Servers are in the RUNNING mode on phys-node-1.**

- 6 **Go to the directory containing the domain directory on phys-host-2 and start the Managed Server mgd2.**

For example:

```
cluster-node1>cd /mnt/user_projects/domains/base_domain_mgd1
cluster-node1>./startManagedWebLogic.sh mgd2
```

- 7 **Make sure that the managed servers are listening on all local addresses on the configured port on both the nodes.**

- 8 **If the clients can successfully connect to the Administration Server and the Managed Server on each of the potential master nodes, you have successfully configured the WebLogic Servers to work with the Oracle Solaris Cluster HA for WebLogic Server in a multi-master configuration.**

- 9 **Stop the Administration Server and the Managed Servers, and proceed to the next task to install the Oracle Solaris Cluster HA for WebLogic Server software.**

Installing the HA for WebLogic Server Package

If you did not install the HA for WebLogic Server package during your initial Oracle Solaris Cluster installation, perform this procedure to install the package.

▼ How to Install the HA for WebLogic Server Package

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

- 1 On the cluster node where you are installing the data service package, become superuser.
- 2 Ensure that the `solaris` and `ha-cluster` publishers are valid.

```
# pkg publisher
PUBLISHER          TYPE    STATUS  URI
solaris            origin online  solaris-repository
ha-cluster        origin online  ha-cluster-repository
```

For information about setting the `solaris` publisher, see [“Set the Publisher Origin To the File Repository URI”](#) in *Copying and Creating Oracle Solaris 11 Package Repositories*.

- 3 Install the HA for WebLogic Server software package.

```
# pkg install ha-cluster/data-service/weblogic
```

- 4 Verify that the package installed successfully.

```
$ pkg info ha-cluster/data-service/weblogic
```

Installation is successful if output shows that State is Installed.

- 5 Perform any necessary updates to the Oracle Solaris Cluster software.

For instructions on updating single or multiple packages, see [Chapter 11, “Updating Your Software,”](#) in *Oracle Solaris Cluster System Administration Guide*.

Registering and Configuring the Oracle Solaris Cluster HA for WebLogic Server

This section contains the procedures that you need to configure Oracle Solaris Cluster HA for WebLogic Server.

Setting Oracle Solaris Cluster HA for WebLogic Server Extension Properties

The sections that follow contain instructions for registering and configuring Oracle Solaris Cluster HA for WebLogic Server resources. For information about the extension properties, see [Appendix A, “HA for WebLogic Server Extension Properties.”](#) The Tunable entry indicates when you can update a property.

For details about all the Oracle Solaris Cluster properties, see the following man pages: `cluster(1CL)`, `rt_properties(5)`, `r_properties(5)`, and `rg_properties(5)`.

To set an extension property of a resource, include the following option in the `clresource(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 *Oracle Solaris Cluster Data Services Planning and Administration Guide*](#) to configure resources after the resources are created.

▼ How to Register and Configure the Oracle Solaris Cluster HA for WebLogic Server as a Web Proxy Server

You can configure your WebLogic Server running as a web proxy server in a failover configuration. The procedure below assumes that you have completed the following steps:

- You are using one of the configurations described in [“Standard Data Service Configurations” on page 14](#).
- You have configured the WebLogic Servers as described in [“Installing and Configuring the WebLogic Server Application” on page 20](#)
- You have created a resource group for running WebLogic Server as a web proxy server.
- You have started an administration server and a managed server either inside or outside the agent’s control.

This configuration does not enable a smooth shutdown. The Oracle Solaris Cluster HA for WebLogic Server stop method forces the shutdown of the WebLogic Server, and kills the process. See [“Creating a Resource That Shuts Down Smoothly” on page 34](#) for instructions on enabling a smooth shut down. For more information, see [“STOP Method” on page 37](#).

1 Install and configure the Oracle Solaris Cluster HA for WebLogic Server as a web proxy server.

Follow the instructions on the [Oracle download site](#).

2 Add the WebLogic Server resource type.

```
# clresourcetype register SUNW.wls
```

3 Create a web proxy server resource for starting the WebLogic Server as a proxy server.

```
# clresource create -g wls-proxy-rg -t SUNW.wls
-p Confdir_list=/global/bea/beahome/weblogic103
-p Server_url=http://logical-host-2:portnumber
-p Start_script=/global/bea/beahome/user_projects/petstore/startManagedWebLogic.sh
-p resource_dependencies=wls-mgd1-rs,wls-mgd2-rs.... wls-proxy-rs
```

Note – If you are running managed server instances outside agent control, ensure that the managed server instances are running. If you are running the managed server instances under the agent control, add the managed server resources that are configured in the web proxy server configuration under the `resource_dependencies` property. This ensures that the managed server resources are online before bringing the proxy server resource online.

▼ How to Register and Configure Oracle Solaris Cluster HA for WebLogic Server as a Failover Data Service

Use this procedure to configure Oracle Solaris Cluster HA for WebLogic Server as a failover data service. The following steps assume that you are using the failover configuration that is described in [“Standard Data Service Configurations” on page 14](#) and have configured the WebLogic Servers as mentioned in [“Installing and Configuring the WebLogic Server Application” on page 20](#). These steps assume that you have already created resource groups RG1 and RG2 and have also created `logical-host-1` and `logical-host-2` resources in these resource groups. Refer to [“How to Prepare the Nodes for a Failover Configuration” on page 18](#) for procedures to create logical hostname resources. If you are planning a different configuration, you must configure the WebLogic Server resources accordingly.

Note – Before you start Oracle Solaris Cluster HA for WebLogic Server, if you are using a database, check that your database is accessible.

1 Add the WebLogic Server resource type.

```
# clresourcetype register SUNW.wls
```

2 Create a WebLogic Administration Server resource in the RG1 resource group.

```
# clresource create -g RG1 -t SUNW.wls \
-p Confdir_list=/global/beam/beahome/weblogic700 \
-p Server_url=http://logical-host-1:7001 \
-p Start_script=/global/beam/beahome/user_projects/petstore/startWebLogic.sh admin-rs
```

The resource is created in the enabled state.

Note – This configuration does not enable smooth shutdown. The Oracle Solaris Cluster HA for WebLogic Server stop method forces the shutdown of the WebLogic Administration Server (kills the process). Refer to [“Creating a Resource That Shuts Down Smoothly” on page 34](#) for the configuration to enable smooth shutdown. For more information, see [“STOP Method” on page 37](#).

3 Create a WebLogic Managed Server resource in the RG2 resource group.

```
# clresource create -g RG2 -t SUNW.wls \
-p Confdir_list=/global/BEA/beahome/weblogic700 \
-p Server_url=http://logical-host-2:7004 \
-p Start_script=/global/BEA/beahome/user_projects/petstore/startManagedWebLogic.sh managed-rs
```

The resource is created in the enabled state.

▼ How to Register and Configure Oracle Solaris Cluster HA for WebLogic Server as a Multi-Master Data Service

Use this procedure to configure Oracle Solaris Cluster HA for WebLogic Server in a multi-master configuration. The following steps assume that you are using the configuration that is described in [“WebLogic Server Configured as a Multi-Master Resource” on page 17](#) and have configured the WebLogic Servers as mentioned in [“Installing and Configuring the WebLogic Server Application” on page 20](#). These steps assume that you have already created a resource group and started the administration server either outside or inside the agent control. If you are planning a different configuration, you must configure the WebLogic Server resources accordingly.

Note – This configuration does not enable smooth shutdown. The Oracle Solaris Cluster HA for WebLogic Server stop method forces the shutdown of the WebLogic Server (kills the process). Refer to [“Creating a Resource That Shuts Down Smoothly” on page 34](#) for the configuration to enable smooth shutdown in a multi-master configuration. For more information, see [“STOP Method” on page 37](#).

1 Add the WebLogic Server resource type.

```
# clresourcetype register SUNW.wls
```

2 Create a multi-master resource for starting the managed server instances simultaneously.

Note – The `Server_url` must be set to `localhost:port-number` because the managed server instances are listening on all local addresses.

The following example starts the WebLogic Server managed server instances `mgd1` on `node1` and `mgd2` on `node2`.

```
# clresource create -t SUNW.wls -g mdg-rg \
-p Confdir_list=/local/BEA/beahome/wlserver_10.3 \
-p Server_url=http://localhost:port-number \
-p Start_script=/local/BEA/beahome/user_projects/domains/base_domain_mgd1 \
/bin/startManagedWebLogic.sh \
-p "Server_name{node1}"=mgd1 \
-p "Server_name{node2}"=mgd2 \
-p smooth_shutdown=true managed-rs
```

Tip – You can create the resource in the global cluster or in a zone cluster. To create a resource in the specific zone cluster from the global cluster, use the `-Z` option to specify the name of the zone cluster.

Setting Up the WebLogic Server for HAStoragePlus Configuration

Use the following procedure to configure an HAStoragePlus resource.

▼ How to Set Up the WebLogic Server on Zone Clusters for HAStoragePlus Configuration

Use the following procedure to configure an HAStoragePlus resource for zone clusters.

Note – The examples shown in the following procedure assume that you perform all steps from inside the zone cluster. If you choose to create resources and resource groups from outside the cluster, you need to add the `-Z sczone` option to all commands to specify the name of the zone cluster.

Before You Begin

- The entries in the `/etc/vfstab` file for cluster file systems should contain the global keyword in the mount options.
- The cluster file systems can be configured for zone clusters in the HAStoragePlus resources using the loopback mount mechanism. The `SUNW.HAStoragePlus` resource type makes the cluster file system available to a zone cluster by mounting the file system in the global cluster and then performing a loopback mount in the zone cluster node.
- The cluster file systems configured in the HAStoragePlus resource type for zone clusters should be authorized for use in zone clusters using the `clzonecluster` command. For more information, see “[How to Add a Cluster File System to a Zone Cluster](#)” in *Oracle Solaris Cluster Software Installation Guide*.
- In Zone Cluster, file systems that are used by different resources in different resource groups must reside in a single HAStoragePlus resource that resides in a scalable resource group. The node list of the scalable HAStoragePlus resource group must be a superset of the node lists of the application resource groups that have resources which depend on the file systems. These application resources that depend on the file systems must have a strong resource dependency set to the HAStoragePlus resource.

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

2 Create the scalable resource group that contains the HASStoragePlus resource.

```
# clresourcegroup create -p Maximum primaries=m \
-p Desired primaries=n [-n node-zone-list] hasp-rg
```

```
-p Maximum primaries=m
```

Specifies the maximum number of active primaries for the resource group.

```
-p Desired primaries=n
```

Specifies the number of desired primaries on which the resource group should attempt to start.

```
-n node-zone-list
```

Specifies the list of node names as the node list of the HASStoragePlus resource group where the WebLogic Server instances can be brought online.

```
hasp-rg
```

Specifies the name of the scalable resource group to be added. This name must begin with an ASCII character.

3 Configure the cluster file system in the zone cluster.

For more information, see “[Configuring a Zone Cluster](#)” in *Oracle Solaris Cluster Software Installation Guide*.

4 Register the resource type for the HASStoragePlus resource.

```
# clresourcetype register HASStoragePlus
```

5 Create the HASStoragePlus resource *hasp-rs* and define the WebLogic Server file system mount points.

```
# clresource create -g hasp-rg -t SUNW.HASStoragePlus \
-p filesystemMountPoints=/global/wls hasp-rs
```

6 Bring the HASStoragePlus resource and resource group online.

```
# clresourcegroup online -M hasp-rg
```

7 Register the resource type for the WebLogic Server application.

```
# clresourcetype register SUNW.wls
```

8 Create a WebLogic Server resource group.

```
# clresourcegroup create -g wls-rg
```

9 Add the WebLogic Server resource to *wls-rg* and set the dependency to HASStoragePlus resource.

```
# clresource create -t SUNW.wls -g wls-rg \
-p Confdir_list=/global/BEA/beahome/wlserver_10.3 \
-p Server_url=http://localhost:port-number \
-p Start_script=/global/BEA/beahome/user_projects/domains/base_domain_mgd1/bin/startManagedWebLogic.sh \
-p "Server_name{node1}"= ManagedServer_1 \
```

```
-p "Server_name{node2}"= ManagedServer_2 \  
-p smooth_shutdown=true \  
-p resource_dependencies=hasp-rs wls-rs
```

10 Bring the failover resource group online.

```
# clresourcegroup online -M wls-rg
```

Verifying the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration

This section describes the procedure to verify if you have installed and configured your data service correctly.

▼ How to Verify the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration for Failover Services

To verify that you have installed and configured Oracle Solaris Cluster HA for WebLogic Server correctly for a failover configuration, perform the following steps:

1 Make sure that you can successfully connect to the WebLogic Servers from a client and that both the Administration Server and the Managed Server are in the RUNNING mode.

2 Switch the resource groups to *cluster-node2*.

```
# clresourcegroup switch cluster-node2 RG1  
# clresourcegroup switch cluster-node2 RG2
```

The WebLogic Administration and Managed Servers must go offline on *cluster-node1* and go online on *cluster-node2*. Ensure that the WebLogic Server processes on *cluster-node1* are stopped. Because smooth shutdown is not enabled, the WebLogic Server process will be killed by the Oracle Solaris Cluster HA for WebLogic Server STOP method instead of being shut down smoothly. Refer to [“Creating a Resource That Shuts Down Smoothly” on page 34](#) for the configuration that enables smooth shutdown. For more information, see [“STOP Method” on page 37](#) for the configuration that enables smooth shutdown.

3 Make sure that you can successfully connect to the WebLogic Servers from the client and that both the Administration Server and the Managed Server are in the RUNNING mode.

Note – The client sessions might be lost when the WebLogic Servers go offline on *cluster-node1* and come online on *cluster-node2*.

4 On *cluster-node2*, kill the Administration Server process.

The Fault Monitor must detect the killed process and restart the Administration Server.

5 Kill the Administration Server process repeatedly.

The Administration Server Resource group fails over after the number of retries exceeds the retry count within the retry interval.

6 On *cluster-node2*, kill the Managed Server process.

The Fault Monitor must detect the killed process and restart the Managed Server.

7 Kill the Managed Server process repeatedly.

The Managed Server Resource group fails over after the number of retries exceeds the retry count within the retry interval.

If the previous steps are successful, you have correctly configured Oracle Solaris Cluster HA for WebLogic Server.

Troubleshooting

After the WebLogic Administration and Managed servers are in production under Oracle Solaris Cluster control, do not start or stop the servers manually or through the Administration Console. If you need to perform maintenance on the servers, you must first disable their associated resources.

▼ **How to Verify the Oracle Solaris Cluster HA for WebLogic Server Installation and Configuration in a Multi-Master Configuration**

To verify that you have installed and configured Oracle Solaris Cluster HA for WebLogic Server correctly for a multi-master configuration, perform the following steps:

1 Ensure that you can successfully connect to the WebLogic Servers from a client and that all the Managed Servers are in the RUNNING mode.

2 On one of the cluster nodes, kill the Managed Server process.

The Fault Monitor must detect the killed process and restart the Managed Server.

3 On the WebLogic Server cluster, deploy a sample application that consists of the WebLogic servers.

4 Use a proxy web server to try to access the sample application.

5 Shut down one of the nodes and try to access the application using the proxy web server address.

You should still be able to access the application.

If the previous steps are successful, you have correctly configured Oracle Solaris Cluster HA for WebLogic Server in a multi-master configuration.

Troubleshooting After the WebLogic Server managed servers are in production under Oracle Solaris Cluster control, do not start or stop the servers manually or through the Administration Console. If you need to perform maintenance on the servers, you must first disable their associated resources.

Alternate Oracle Solaris Cluster HA for WebLogic Server Configurations

The installation, registration, and configuration procedures in this chapter assume that you are installing the configuration that is described in “[Standard Data Service Configurations](#)” on [page 14](#). This section provides examples of other possible Oracle Solaris Cluster HA for WebLogic Server configurations.

Creating a Simple WebLogic Server Resource

This configuration assumes that the START script, `startWebLogic.sh`, can start the WebLogic Server without any arguments to the script. The user name and password that are needed to start the WebLogic Server can be configured within this START script or in the `boot.properties` file.

Note – Starting with WebLogic Server version 10.3, the `boot.properties` file must reside in both the `$DOMAIN_DIR` and `$DOMAIN_DIR/servers/server-name/security` directories to bypass the need to prompt for user credentials when starting instances of WebLogic Server.

```
# clresource create -g bea-rg -t SUNW.wls \
-p Confdir_list=/global/bea/beahome/weblogic700 \
-p Server_url=http://logical_host1:7001 \
-p Start_script=/global/bea/beahome/user_projects/petstore/startWebLogic.sh bea-rs
```

The resource is created in the enabled state.

Creating a Resource With a Server Name Argument

If the WebLogic Server START script `startManagedWebLogic.sh` uses a Managed Server name as an argument, then the extension property `Server_name` can be set. The START script

`startManagedWebLogic.sh` can be used to start several Managed Servers. The `ADMIN_URL` must be configured within the `START` script. To start the WebLogic Server manually, use the following commands.

```
cluster-node-1> cd /global/wls/user_projects/petstore
cluster-node-1> ./startManagedWebLogic.sh petstore1
```

To configure a Oracle Solaris Cluster HA for WebLogic Server resource whose start script takes the server name `petstore1`, use the following commands.

```
# clresource create -g bea-rg -t SUNW.wls \
-p Confdir_list=/global/bea/beahome/weblogic700 \
-p Server_url=http://logical_host1:7001 \
-p Start_script=/global/bea/beahome/user_projects/petstore/startWebLogic.sh \
-p Server_name=petstore1 bea-rs
```

The resource is created in the enabled state.

To configure another Oracle Solaris Cluster HA for WebLogic Server resource that uses the same `START` script for a different server, `petstore2`, use the following commands.

```
# clresource create -g bea-rg -t SUNW.wls \
-p Confdir_list=/global/bea/beahome/weblogic700 \
-p Server_url=http://logical_host1:7001 \
-p Start_script=/global/bea/beahome/user_projects/petstore/startWebLogic.sh \
-p Server_name=petstore2 bea-rs
```

The resource is created in the enabled state.

Creating a Resource That Shuts Down Smoothly

The default for the Oracle Solaris Cluster HA for WebLogic Server `STOP` method kills the WebLogic Server process to shut down the instance. To enable smooth shutdown of the WebLogic Server, set the `Smooth_shutdown` extension property to `TRUE`. If this extension property is enabled, the `DOMAIN_DIR/boot.properties` file must contain a valid user name and password.

Note – Starting with WebLogic Server version 10.3, the `boot.properties` file must reside in both the `DOMAIN_DIR` and `DOMAIN_DIR/servers/server-name/security` directories to bypass the need to prompt for user credentials when stopping instances of WebLogic Server. If the `DOMAIN_DIR/servers/server-name/security` directory does not exist or does not include a `boot.properties` file, create the directory and copy the `boot.properties` file from the `DOMAIN_DIR`.

To configure a resource which will shut down smoothly in a failover configuration, use the following commands.

```
# clresource create -g bea-rg -t SUNW.wls \
-p Confdir_list=/global/bean/beahome/weblogic700 \
-p Server_url=http://logical_host1:7001 \
-p Start_script=/global/bean/beahome/user_projects/petstore/startWebLogic.sh \
-p Server_name=petstore1 \
-p Smooth_shutdown=TRUE bea-rs
```

The resource is created in the enabled state.

To configure a resource that will shut down smoothly in a mult-master configuration, use the following commands:

```
# clresource create -g RG -t SUNW.wls \
-p Confdir_list=/mnt/wlserver_10.3/ \
-p Server_url=http://localhost:port-number \
-p Start_script=/mnt/user_projects/domains/base_domain_mgd1/startManagedWebLogic.sh \
-p "Server_name{phys-host-1}"=mgd1 \
-p "Server_name{phys-host-2}"=mgd2 \
-p Smooth_shutdown=true managed-rs
```

The resource is created in the enabled state.

Creating a Resource That Probes the Database

The Oracle Solaris Cluster HA for WebLogic Server does not probe the database health. If the database is down when Oracle Solaris Cluster HA for WebLogic Server starts, the data service fails to start. To have the Oracle Solaris Cluster HA for WebLogic Server probe the database, supply your own database probe script to Oracle Solaris Cluster HA for WebLogic Server as an extension property. This database probe must return zero for success and nonzero for failure.

To configure a resource with the DB_probe script extension property set, use the following commands.

```
# clresource create -g bea-rg -t SUNW.wls \
-p Confdir_list=/global/bean/beahome/weblogic700 \
-p Server_url=http://logical_host1:7001 \
-p Start_script=/global/bean/beahome/user_projects/petstore/startWebLogic.sh \
-p Server_name=petstore1 \
-p Smooth_shtudown=TRUE \
-p DB_probe_script=/global/bean/script/db_probe_script bea-rs
```

The resource is created in the enabled state.

Creating a Resource That Monitors URIs

To enable the Oracle Solaris Cluster HA for WebLogic Server resource to monitor URIs that are served by the WebLogic Server configured in the resource, use the following commands to set the extension property `Monitor_uri_list`.

```
# clresource create -g bea-rg -t SUNW.wls \  
-p Confdir_list=/global/bea/beahome/weblogic700 \  
-p Server_url=http://logical_host1:7001 \  
-p Start_script=/global/bea/beahome/user_projects/petstore/startWebLogic.sh \  
-p Server_name=petstore1 \  
-p Smooth_shutdown=TRUE \  
-p DB_probe_script=/global/bea/script/db_probe_script \  
-p Monitor_uri_list=http://logical_host1:7001/petstore,http://logical_host1:7001/bookstore bea-rs
```

The resource is created in the enabled state.

Using a Non-Clustered Managed Server Instance as a Proxy Server

You can configure a non-clustered managed server instance to run as a proxy server that hosts a servlet that directs the requests to the clustered managed server instances. In this configuration, you can define a failover configuration to make the WebLogic proxy server highly-available. When you define this configuration, make sure that you specify a resource dependency from the WebLogic Server proxy server to the managed server resource. Setting this dependency ensures that the proxy server resource will not try to come online unless the managed server resource is online.

Operation of the Oracle Solaris Cluster HA for WebLogic Server Fault Monitor

This section describes the Oracle Solaris Cluster HA for WebLogic Server Fault Monitor.

Probing Algorithm and Functionality

The Fault Monitor detects failures and takes action. If the monitor detects a failure in a WebLogic Server, it first restarts the WebLogic Server. If the WebLogic Server fails a certain number of times (configurable by the administrator) within a certain time window (configurable by the administrator), the resource group containing the WebLogic Server is failed over to another surviving cluster node and restarted.

The Fault Monitor method, by default, probes the server URL set in the extension property `Server_url`. The probe connects to the hostname and the port and then sends an HTTP GET request on the URL. If the connection fails, it is considered a complete failure and the resource group containing the WebLogic Server is restarted or failed over to another surviving cluster node and restarted. If the connection succeeds, but the HTTP response code is 500 (internal server error), it is also considered a complete failure and the resource group is restarted or failed over. All other HTTP response codes are considered a success.

If the `monitor_uri_list` extension property is set, the probe method connects to the URIs mentioned in the list and takes action if a failure occurs. The probe sends an HTTP GET on the specified URI or URIs.

Note – The `Monitor Uri List` extension property supports HTTP requests only. It does not support HTTPS requests.

If a complete failure (URL or URI probe) of the WebLogic Server instance is detected by the probe, and if a database probe script is specified in the extension property `db_probe_script`, the probe method will probe the database before taking any action on the WebLogic Server resource. If the database probe script returns success (database is up), action is taken on the WebLogic Server resource. If the database probe script returns a failure (database is down), the WebLogic Server probe will not take any action (restart or failover) until the database is up.

START Method

Before starting the WebLogic Server that is configured in the resource, the WebLogic Server configuration and the resource extension properties are validated. If the `db_probe_script` extension property is set, the database is probed by running the script that is set in the extension property. If the database is up, the WebLogic Server is started by running the START script that is configured in the extension property `Start_script` under `pmf`. If the database is not up, the START method returns success and lets the probe method handle the starting of the WebLogic Server. The probe method waits until the database is up to start the WebLogic Server, as explained in [“Probing Algorithm and Functionality” on page 36](#).

After starting the START script under `pmf`, the START method waits until the WebLogic Server is in RUNNING mode before declaring the START method successful. While waiting for the WebLogic Server to start, the probe method tries to connect to the server to check if it is up. Some messages are displayed on the console during startup. The message “Failed to connect to host logical-host-1 and port 7001: Connection refused” will continue to be displayed until the WebLogic Server starts completely. After the WebLogic Server is in the RUNNING mode, the START method sets the status to “Started Successfully.”

WebLogic Server console messages are stored at `/var/cluster/logs/hawls/wls-resource-name` on the cluster node where the resource is started. You can check this log file for startup messages of WebLogic Server instances.

STOP Method

The STOP method stops the WebLogic Server that is configured in the resource. By default, the STOP method kills the WebLogic Server by sending a SIGKILL to the WebLogic Server process. If the `smooth_shutdown` extension property is set to TRUE, the STOP method tries to bring down the WebLogic Server instance by running the following command:

```
java weblogic.Admin -url hostname:port -username $WLS_USER  
-password $WLS_PW SHUTDOWN
```

If this command fails, the WebLogic Server is shut down by using SIGKILL. Even if the command succeeds, the STOP method sends SIGKILL to ensure that the WebLogic Server process is stopped.

If the `smooth_shutdown` extension property is set to TRUE, ensure that the `boot.properties` file with valid user name and password is present in the DOMAIN directory. The WebLogic Server variables, `WLS_USER` and `WLS_PW` do not need to be set in the START script for smooth shutdown. This waiver applies only to WebLogic Server 8.1 or later versions.

HA for WebLogic Server Extension Properties

This section describes the extension properties for the resource type `SUNW.wls`. This resource type represents the Apache application in a Oracle Solaris 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.wls` resource type are as follows:

`Confdir_list`

The complete path to the WebLogic Server home directory, `$WL_HOME`.

Note – If `setWLSenv.sh` and `startWLS.sh` are not in the standard `$confdir_list/server/bin` location, then you must create links to these files from `$confdir_list/server/bin` directory. If you do not create this link, the WebLogic Server resource creation will fail.

Starting with WebLogic Server version 10.3, the `startWLS.sh` script has been removed. Therefore, you do not need to create a link to the `startWLS.sh` script for these versions.

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

`DB_Probe_Script`

The complete path to a Database probe script. The HA-WLS probe method probes only the WLS instances. The database probe script can be provided using this extension property if the Administrators want the HA_WLS to probe the Database also. This probe script must return 0 for success. The WebLogic Servers will be started only if this Database probe returns success. If an HA-WLS probe detects a failure in a WebLogic Server instance and if this extension

property is set, the HA -WLS probe will take action only if the Database probe succeeds. Refer to [“Probing Algorithm and Functionality” on page 36](#) for more details.

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

Environment_File

The absolute path to the WebLogic Server environment file. This file is typically installed in `$WEBLOGIC_HOME/server/bin` when WebLogic Server is configured by using the installation wizard.

Data type	String
Default	No default defined
Range	Not applicable
Tunable	When disabled

Monitor_Uri_List

A single URI, or a comma-separated list of URIs, which can be used by the Fault Monitor to test the functionality of the WebLogic Server by doing an HTTP GET on the URI. The `Server_Url` extension property is for simple probes on the host name and port. This extension property can be used to probe the WebLogic Server functionality by retrieving a Java servlet or making the WebLogic Server run an application and return an HTTP server code. If the HTTP server return code is 500 (Internal Server Error) or if the connect fails, the probe will take action. See the probe method for more details.

Make sure that the host names used in the `server_url` and `monitor_uri_list` are resolvable. If you use Fully Qualified Domain Names (FQDNs), then DNS must be enabled and `/etc/nsswitch.conf` must have the correct entries to resolve the host names using DNS.

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

Server_name

A single start script that can be used to start all the Managed Servers by passing the Managed Server name as an argument to the startup script. If Agent start methods have to pass the server name as an argument to the start script then, configure the Managed Server name in this extension property.

You can specify a different value of the `Server_name` extension property for each node that can master the resource.

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

Server_Url

The complete URL of the WebLogic Server, including the protocol used to connect to the server that this resource is configured to start. The probe method uses this URL to check the health of the server by performing an HTTP GET on the URL. The protocol in the URL must be set to `http`. The complete URL should be in the following form:

```
http://host:port
```

For more information about the `Server_Url`, see the relevant WebLogic Server documentation for the version you are using. For example, for WebLogic Server 10.3, the documentation is located at http://download.oracle.com/docs/cd/E12840_01/wls/docs103/sitemap.html.

Ensure sure that the host names used in the `server_url` and `monitor_uri_list` are resolvable. If you use Fully Qualified Domain Names (FQDNs), then DNS must be enabled and `/etc/nsswitch.conf` must have the correct entries to resolve the host names using DNS.

Note – To avoid any failures due to delays in name-service lookup, the host name used in the `server_url` property should be added to the `/etc/hosts` file. Configure name-service mapping in the `/etc/nsswitch.conf` file on the servers to first check the local files before trying to access other name services.

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

Smooth_shutdown

Indication of whether to enable smooth shutdown using `WebLogic.Admin` class. To perform a smooth shutdown before trying to kill the WebLogic Server process, set this extension property to `TRUE` and provide a valid user name and password in the `DOMAIN_DIR/boot.properties` file.

Data type	Boolean
------------------	---------

Default	FALSE
Range	Not applicable
Tunable	At any time

Start_Script

The complete path to the script that should be used to start the WebLogic Server instance (either Administration or Managed). This script is typically present under the domain name directory along with the `config.xml` file. A separate script or a single script can be configured for starting each WebLogic Server. For more information, see [“Configuration Requirements” on page 12](#).

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

Index

C

- configurations, alternate, 33–36
- configuring, Oracle Solaris Cluster HA for WebLogic Server, 27–28

E

- extension properties, SUNW.wls resource type, 39–42

F

- Fault Monitor, Oracle Solaris Cluster HA for WebLogic Server, 36–38

H

- HA for WebLogic Server
 - installing, 24–25
 - software package, installing, 24–25
- help, 8

I

- installing
 - database, 22
 - HA for WebLogic Server, 24–25
 - web server, 22
 - WebLogic Server, 20–22

O

- Oracle Solaris, publisher, 25
- Oracle Solaris Cluster, publisher, 25
- Oracle Solaris Cluster HA for WebLogic Server
 - See also* WebLogic Server
 - alternate configurations, 33–36
 - configuration
 - planning, 11–18
 - requirements, 12–14
 - restrictions, 11–12
 - standard, 14–18
 - Fault Monitor, 36–38
 - installation
 - planning, 11–18
 - verifying, 31–32
 - overview, 9–10
 - registering and configuring, 27–28
- overview, Oracle Solaris Cluster HA for WebLogic Server, 9–10

P

- package, 24–25
- probing algorithm and functionality, 36–38
 - START method, 37
 - STOP method, 37–38
- publisher
 - Oracle Solaris, 25
 - Oracle Solaris Cluster, 25

R

registering, Oracle Solaris Cluster HA for WebLogic Server, 27–28

S

setting up the WebLogic Server for HAStoragePlus
configuration, zone cluster, 29
software package, 24–25
START method, 37
STOP method, 37–38
SUNW.wls resource type, extension properties, 39–42

T

technical support, 8

V

verifying
Oracle Solaris Cluster HA for WebLogic Server, 31–32
WebLogic Server installation, 22–23, 23–24

W

WebLogic Server
See also Oracle Solaris Cluster HA for WebLogic Server
installation
installing a database, 22
installing a web server, 22
installing software, 20–22
installing WebLogic Server, 20–21
verifying, 22–23, 23–24

Z

zone cluster, setting up the WebLogic Server for
HAStoragePlus configuration, 29