

**Oracle® Solaris Cluster Data Service
for Oracle JD Edwards EnterpriseOne
Enterprise Server Guide**

SPARC Platform Edition

ORACLE®

Part No: E64637
September 2015

Part No: E64637

Copyright © 2000, 2015, 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, then the following notice is applicable:

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

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

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

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

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support 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.

Référence: E64637

Copyright © 2000, 2015, Oracle et/ou ses affiliés. Tous droits réservés.

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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

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

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 des applications dangereuses.

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 de 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

- Using This Documentation 13

- 1 Installing and Configuring the HA for JD Edwards EnterpriseOne Enterprise Server Data Service 15**
 - HA for JD Edwards EnterpriseOne Enterprise Server Overview 16
 - Overview of Installing and Configuring HA for JD Edwards EnterpriseOne Enterprise Server 17
 - Planning the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration 18
 - Configuration Restrictions for JD Edwards EnterpriseOne Enterprise Server Software 18
 - Configuration Requirements for HA for JD Edwards EnterpriseOne Enterprise Server 19
 - HA for JD Edwards EnterpriseOne Enterprise Server Configurations 21
 - Preparing the Nodes and Disks 30
 - ▼ How to Prepare the Nodes for a Failover Configuration 30
 - ▼ How to Prepare the Nodes for a Multi-Instance Configuration 32
 - ▼ How to Prepare the Nodes for a Multi-Master Configuration 33
 - Installing and Configuring JD Edwards EnterpriseOne Enterprise Server Software 35
 - ▼ How to Install and Configure JD Edwards EnterpriseOne Enterprise Server Software 35
 - Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration 37
 - ▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Failover Services 37
 - ▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Multi-Master and Multi-Instance Services 39
 - Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package 40

▼ How to Install the HA for JD Edwards EnterpriseOne Enterprise Server Package	40
Registering and Configuring HA for JD Edwards EnterpriseOne Enterprise Server	41
▼ How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server for Failover	42
▼ How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Data Service	45
▼ How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Instance Data Service	48
Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration	51
▼ How to Verify the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Failover Configuration	52
▼ How to Verify the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Multi-Master or Multi-Instance Configuration	53
Tuning the HA for JD Edwards EnterpriseOne Enterprise Server Fault Monitors	54
Resource Properties	54
Probing Algorithm and Functionality	54
Operation of the HA for JD Edwards EnterpriseOne Enterprise Server Probe	55
Debugging HA for JD Edwards EnterpriseOne Enterprise Server	55
▼ How to Activate Debugging for HA for JD Edwards EnterpriseOne Enterprise Server	56
▼ How to Deactivate Debugging for HA for JD Edwards EnterpriseOne Enterprise Server	56
A HA for JD Edwards EnterpriseOne Enterprise Server Extension Properties	59
ORCL.JDE_Enterprise_Server Extension Properties	59
Index	63

Figures

- FIGURE 1** HA for JD Edwards EnterpriseOne Enterprise Server Configured for Failover with Traditional File Storage 23
- FIGURE 2** JD Edwards EnterpriseOne Enterprise Server Configured for Failover with NAS 25
- FIGURE 3** HA for JD Edwards EnterpriseOne Enterprise Server Configured as a Multi-Instance Application With Single-Node Resource Groups 27
- FIGURE 4** HA for JD Edwards EnterpriseOne Enterprise Server in a Multi-Master Configuration Using a Scalable Resource Group 29

Tables

TABLE 1	JD Edwards EnterpriseOne System Components and Their Data Services	16
TABLE 2	Tasks for Installing and Configuring HA for JD Edwards EnterpriseOne Enterprise Server	17

Examples

- EXAMPLE 1** Configuring HA for JD Edwards EnterpriseOne Enterprise Server as a Failover Service on a Two-Node Cluster With Traditional File Storage 43
- EXAMPLE 2** Registering HA for JD Edwards EnterpriseOne Enterprise Server as a Failover Service on a Two-Node Cluster With NAS 44
- EXAMPLE 3** Configuring HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Service 46
- EXAMPLE 4** Configuring HA for JD Edwards EnterpriseOne Enterprise Server in a Multi-Instance Service 49

Using This Documentation

- **Overview** – Describes how to install and configure the Oracle Solaris Cluster HA for JD Edwards EnterpriseOne Enterprise Server data service.
- **Audience** – Technicians, system administrators, and authorized service providers.
- **Required knowledge** – Advanced experience troubleshooting and replacing hardware.

Product Documentation Library

Documentation and resources for this product and related products are available at <http://www.oracle.com/pls/topic/lookup?ctx=E56676-01>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Installing and Configuring the HA for JD Edwards EnterpriseOne Enterprise Server Data Service

This chapter provides an overview of the Oracle Solaris Cluster HA for JD Edwards EnterpriseOne Enterprise Server (HA for JD Edwards EnterpriseOne Enterprise Server) data service and explains how to install and configure it.

This chapter contains the following sections:

- [“HA for JD Edwards EnterpriseOne Enterprise Server Overview” on page 16](#)
- [“Overview of Installing and Configuring HA for JD Edwards EnterpriseOne Enterprise Server” on page 17](#)
- [“Planning the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 18](#)
- [“Preparing the Nodes and Disks” on page 30](#)
- [“Installing and Configuring JD Edwards EnterpriseOne Enterprise Server Software” on page 35](#)
- [“Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 37](#)
- [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package” on page 40](#)
- [“Registering and Configuring HA for JD Edwards EnterpriseOne Enterprise Server” on page 41](#)
- [“Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 51](#)
- [“Tuning the HA for JD Edwards EnterpriseOne Enterprise Server Fault Monitors” on page 54](#)
- [“Debugging HA for JD Edwards EnterpriseOne Enterprise Server” on page 55](#)

HA for JD Edwards EnterpriseOne Enterprise Server Overview

The HA for JD Edwards EnterpriseOne Enterprise Server data service provides orderly startup, shutdown, fault monitoring, and automatic failover of the Enterprise Server. Use the information in this section to understand how to make the Enterprise Server highly available.

The HA for JD Edwards EnterpriseOne Enterprise Server data service provides fault monitoring and automatic failover for the Enterprise Server to eliminate single points of failure in a JD Edwards EnterpriseOne system. [Table 1, “JD Edwards EnterpriseOne System Components and Their Data Services,” on page 16](#) lists the data service that protects each of the JD Edwards EnterpriseOne components in an Oracle Solaris Cluster configuration.

Note - You can install and configure this data service to run in either the global zone or a zone cluster. For updated information about supported configurations of this data service, see the [Oracle Solaris Cluster 4 Compatibility Guide](#).

TABLE 1 JD Edwards EnterpriseOne System Components and Their Data Services

JD Edwards EnterpriseOne System Component	Data Service
Database server	<p>The data service for the database that you are using:</p> <ul style="list-style-type: none"> ■ For Oracle Database, the data service is explained in the Oracle Solaris Cluster Data Service for Oracle Database Guide. ■ For Oracle RAC, the data service is explained in the Oracle Solaris Cluster Data Service for Oracle Real Application Clusters Guide. ■ For Oracle External Proxy, the data service is explained in the Oracle Solaris Cluster Data Service for Oracle External Proxy Guide.
Application server	The data service is HA for JD Edwards EnterpriseOne Enterprise Server. The resource type is ORCL.JDE_Enterprise_Server. The data service is explained in Oracle Solaris Cluster Data Service for JD Edwards EnterpriseOne Enterprise Server Guide .
Web server	The data service is explained in the Oracle Solaris Cluster Data Service for Oracle WebLogic Server Guide .

High availability is provided for the JD Edwards EnterpriseOne Enterprise Server component of the JD Edwards EnterpriseOne system.

The JD Edwards EnterpriseOne Enterprise Server component can be configured as one of the following data services: failover, multiple-master (mastered by multiple nodes), or multiple-instance (more than one instance of the same application provide an aggregation of services).

Note - HA for JD Edwards EnterpriseOne Enterprise Server does not support JD Edwards EnterpriseOne Enterprise Server in an active-active load balancing configuration.

For conceptual information about failover and multi-master services, see the [Oracle Solaris Cluster 4.3 Concepts Guide](#).

HA for JD Edwards EnterpriseOne Enterprise Server requires that a functioning cluster with the initial cluster framework is already installed. See the [Oracle Solaris Cluster 4.3 Software Installation Guide](#) for details on initial installation of clusters and data service software. You register HA for JD Edwards EnterpriseOne Enterprise Server after you successfully install the basic components of Oracle Solaris Cluster and JD Edwards EnterpriseOne software.

For information about which version of JD Edwards EnterpriseOne Enterprise Server software is compatible, see the [Oracle Solaris Cluster Compatibility Guide](#) available at the [Oracle Solaris Cluster Technical Resources](#) page.

Overview of Installing and Configuring HA for JD Edwards EnterpriseOne Enterprise Server

The JD Edwards EnterpriseOne Enterprise Server component is protected by the HA for JD Edwards EnterpriseOne Enterprise Server data service.

[Table 2, “Tasks for Installing and Configuring HA for JD Edwards EnterpriseOne Enterprise Server,” on page 17](#) summarizes the tasks for installing and configuring HA for JD Edwards EnterpriseOne Enterprise Server and provides cross-references to detailed instructions for performing these tasks. Perform the tasks in the order in which they are listed in the table.

TABLE 2 Tasks for Installing and Configuring HA for JD Edwards EnterpriseOne Enterprise Server

Task	Instructions
1. Plan the installation	“Planning the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 18
2. Install and configure the JD Edwards EnterpriseOne Enterprise Server	“Installing and Configuring JD Edwards EnterpriseOne Enterprise Server Software” on page 35
3. Verify the JD Edwards EnterpriseOne Enterprise Server installation and configuration	“Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 37
4. Register and configure HA for JD Edwards EnterpriseOne Enterprise Server resources	“Registering and Configuring HA for JD Edwards EnterpriseOne Enterprise Server” on page 41

Task	Instructions
5. Verify the HA for JD Edwards EnterpriseOne Enterprise Server installation and configuration	“Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 51
6. Tune the HA for JD Edwards EnterpriseOne Enterprise Server fault monitor	“Tuning the HA for JD Edwards EnterpriseOne Enterprise Server Fault Monitors” on page 54
7. Debug HA for JD Edwards EnterpriseOne Enterprise Server	“Debugging HA for JD Edwards EnterpriseOne Enterprise Server” on page 55

Planning the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration

This section contains the following information to plan your JD Edwards EnterpriseOne Enterprise Server installation and configuration:

- [“Configuration Restrictions for JD Edwards EnterpriseOne Enterprise Server Software” on page 18](#)
- [“Configuration Requirements for HA for JD Edwards EnterpriseOne Enterprise Server” on page 19](#)
- [“HA for JD Edwards EnterpriseOne Enterprise Server Configurations” on page 21](#)

Note - Before you begin, consult your JD Edwards EnterpriseOne documentation for configuration restrictions and requirements that are not stated in Oracle Solaris Cluster documentation or imposed by Oracle Solaris Cluster software. Read all JD Edwards EnterpriseOne notes that pertain to your release of JD Edwards EnterpriseOne Enterprise Server. For more information, see the [JD Edwards EnterpriseOne documentation](#).

Configuration Restrictions for JD Edwards EnterpriseOne Enterprise Server Software

The configuration restrictions in the subsections that follow apply only to the JD Edwards EnterpriseOne Enterprise Server.



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

- Configure the database resource to be in a different resource group than the JD Edwards EnterpriseOne Enterprise Server resource, if possible.
However, if you are using Oracle Database as the database, and if you must configure Oracle Database in the same resource group as the JD Edwards EnterpriseOne Enterprise Server resource, ensure that the value of the `Restart_type` extension property for Oracle Solaris Cluster HA for Oracle Database is set to `RESOURCE_RESTART`. This property setting indicates that if the response to a fault in the Oracle database is to restart the Oracle resource, only the Oracle database is restarted. If this property is set to `RESOURCE_GROUP_RESTART`, and if the response to a fault in the Oracle database is to restart the Oracle resource, all the resources in the resource group are restarted.
- Ensure that each JD Edwards EnterpriseOne Enterprise Server instance is installed as a different user. If the installation is on a local file system, you can use the same user ID for each instance.

For restrictions that apply to all data services, see the [Oracle Solaris Cluster 4.3 Release Notes](#).

Configuration Requirements for HA for JD Edwards EnterpriseOne Enterprise Server

Follow the requirements in this section to plan the installation and configuration of the JD Edwards EnterpriseOne Enterprise Server. These requirements apply only to the JD Edwards EnterpriseOne Enterprise Server. You must meet these requirements before you proceed with your JD Edwards EnterpriseOne Enterprise Server installation and configuration.

Information about how to install JD Edwards EnterpriseOne Enterprise Server is published at the [JD Edwards EnterpriseOne Tools documentation site](http://docs.oracle.com/cd/E17984_01/index.htm) (http://docs.oracle.com/cd/E17984_01/index.htm). See the [Oracle Solaris Cluster 4 Compatibility Guide](http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf) (<http://www.oracle.com/technetwork/server-storage/solaris-cluster/overview/solariscluster4-compatibilityguide-1429037.pdf>) for supported versions of JD Edwards EnterpriseOne Enterprise Server.

For requirements that apply to all data services, see [Chapter 1, “Planning for Oracle Solaris Cluster Data Services”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide*.



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

The following configuration requirements apply:

- **UNIX user and group** – The UNIX user and group that are used to install, operate, and manage the JD Edwards EnterpriseOne Enterprise Server must exist on all cluster nodes where the corresponding resource for the Enterprise Server is configured to come online. When you run the Oracle Universal Installer, a user is created on that node, for example, `jde910`. This user starts the application and the user must exist on every cluster node. You can choose to manually create a `jde` user on every node, or run the installer which creates the user on the node.
- **File systems** – The file systems used to store the required binaries and data for the JD Edwards EnterpriseOne Enterprise Server must be configured on highly available local file systems. If you choose to install the binaries on local storage, install and keep them identical on all the cluster nodes. The `JDE.INI` file must reside on a highly available local file system, which must be accessible where the corresponding resource for the JD Edwards EnterpriseOne Enterprise Server comes online.

If you are setting up the JD Edwards EnterpriseOne Enterprise Server in a multi-instance configuration, a local file system, or a global file system, the Enterprise Server installation should be located on network attached storage (NAS) that is accessible to all nodes running the server.

If you are setting up multiple instances of the JD Edwards EnterpriseOne Enterprise Server in a multi-instance or a multi-master configuration, each instance should be installed as a different user. For JD Edwards EnterpriseOne Enterprise Server version 9.1.x installed through the Platform Pack installer, you can choose to create a user account for the `jde910` user. If your configuration is on a local file system, the installation user can be the same for all instances.

- **Logical hostnames** – Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.
- **Environment variables** – In addition to the required environment variables that are explained in the *JD Edwards EnterpriseOne Installation Guide*, you must set up some variables after you configure the Enterprise Server. For instructions, see Step 5 of [“How to Install and Configure JD Edwards EnterpriseOne Enterprise Server Software”](#) on page 35.
- **Database tier dependency** – If the database tier is deployed on the same global cluster, the resource for the JD Edwards EnterpriseOne Enterprise Server must define a strong resource dependency to the resources for the database instance and database listener. This ensures that the JD Edwards EnterpriseOne Enterprise Server will try to start only when the corresponding database is already operational. This configuration is required for a successful startup of the JD Edwards EnterpriseOne Enterprise Server. If the Oracle database is not managed through Oracle Solaris Cluster, you can also configure the database tier dependency by using the data service delivered as part of the Oracle External Proxy resource.

- **Database client network connection** – The database client used by the JD Edwards EnterpriseOne Enterprise Server configuration must be configured to connect to the network address that is managed by the cluster framework for the corresponding database server.
- **JDE.INI attribute changes** – The following attributes must be changed in the JDE . INI file:
 - Base Datasource should have the logical host resource hostname configured for the JD Edwards EnterpriseOne Enterprise Server resource. For a multiple-master configuration, this is the per-node logical host resource hostname.

For example, the setting in the JDE . INI file looks similar to the following, where *LH-rs-hostname* is the hostname of the per-node logical host resource:

```
Base Datasource = LH-rs-hostname - 910 Server Map
```

- The `listenOnSpecificIP` setting in the JDE . INI file must be set to 1:

```
listenOnSpecificIP=1
```

- The `SecurityServer` attribute under the [SECURITY] section must match the hostname or IP address of the logical host resource configured for the JD Edwards EnterpriseOne Enterprise Server resource. For a multiple-master configuration, this is the per-node logical host resource hostname or IP address.

For example, the setting in the JDE . INI file looks similar to the following, where the IP address of the per-node logical host resource hostname configured for the JD Edwards EnterpriseOne Enterprise Server is 192 . 1 . 2 . 36:

```
SecurityServer=192.1.2.36
```

- The `PrimaryNode` attribute under the [CLUSTER] section must match the hostname or IP address of the logical host resource configured for the JD Edwards EnterpriseOne Enterprise Server resource. For a multiple-master configuration, this is the per-node logical host resource hostname or IP address.

For example, the setting in the JDE . INI file looks similar to the following, where the IP address of the per-node logical host resource hostname configured for the JD Edwards EnterpriseOne Enterprise Server resource is 192 . 1 . 2 . 36:

```
PrimaryNode=192.1.2.36
```

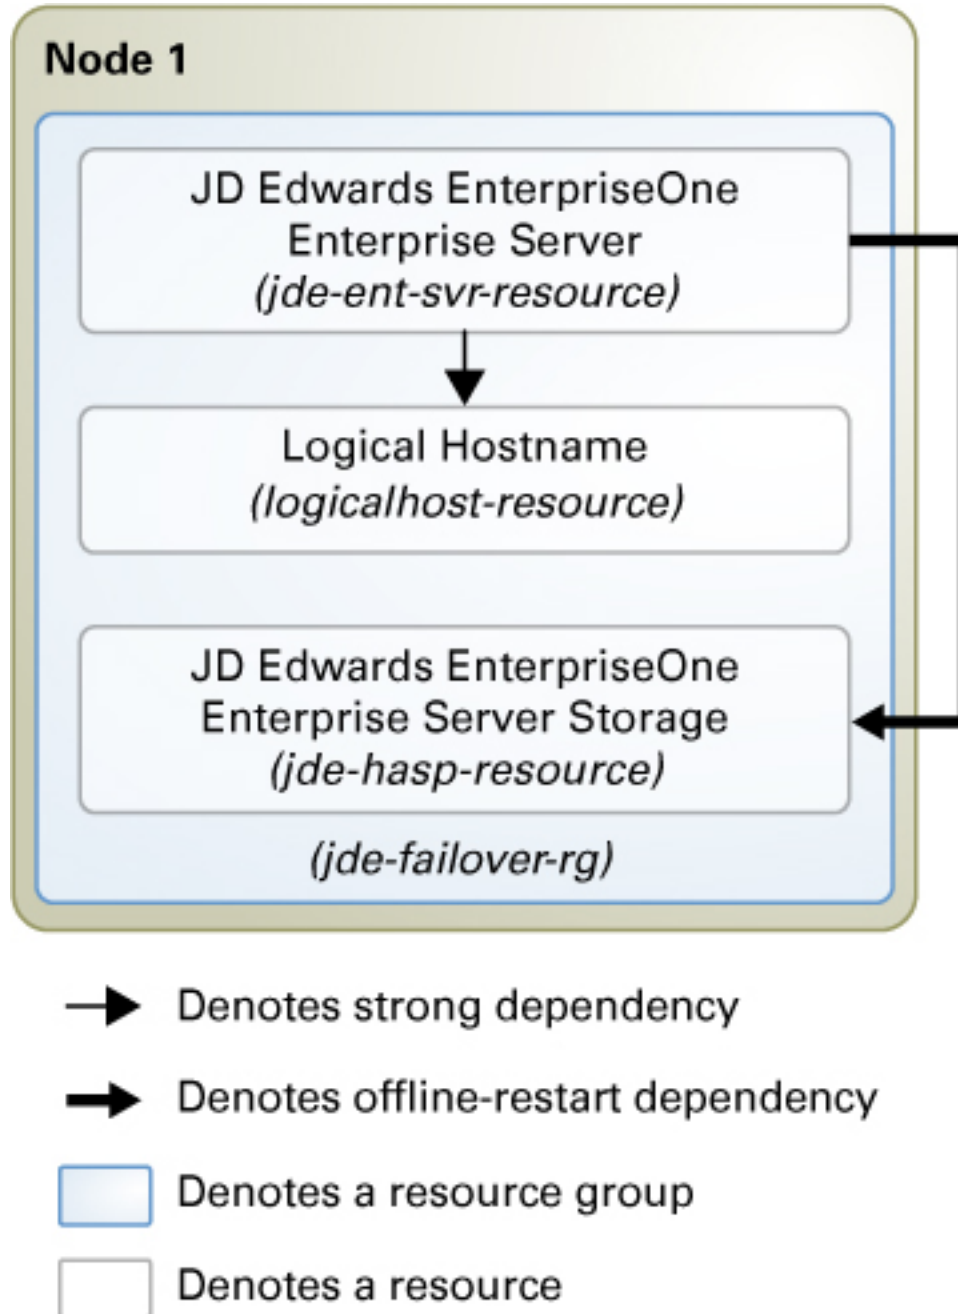
HA for JD Edwards EnterpriseOne Enterprise Server Configurations

Use the data service configurations in this section to plan the installation and configuration of the HA for JD Edwards EnterpriseOne Enterprise Server data service.

Failover Configuration

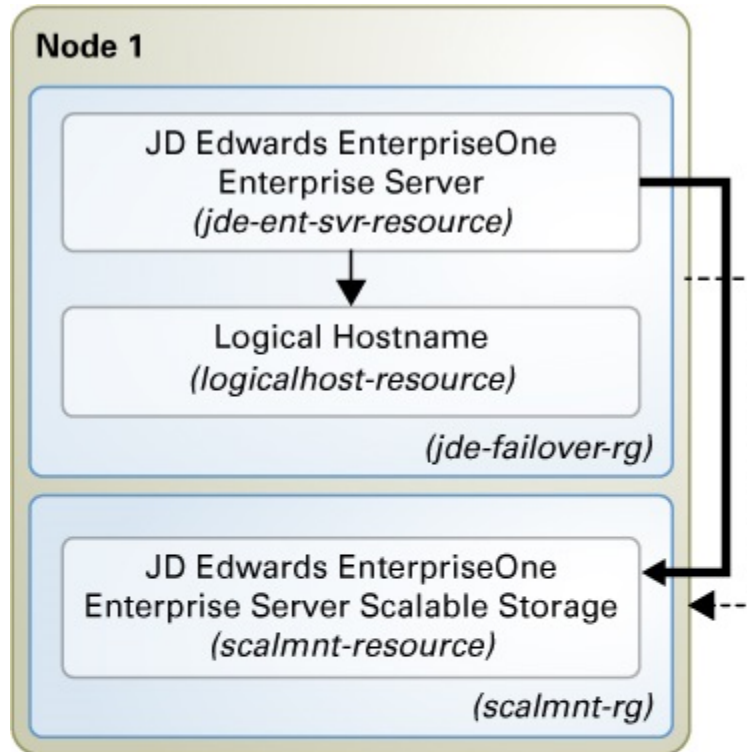
When the HA for JD Edwards EnterpriseOne Enterprise Server software is installed in a traditional file system, a failover deployment requires a configuration where one failover resource group contains the HA for JD Edwards EnterpriseOne Enterprise Server resource, the logical hostname resource, and the failover storage resource. This type of configuration is shown in [Figure 1, “HA for JD Edwards EnterpriseOne Enterprise Server Configured for Failover with Traditional File Storage,”](#) on page 23.

FIGURE 1 HA for JD Edwards EnterpriseOne Enterprise Server Configured for Failover with Traditional File Storage



When the HA for JD Edwards EnterpriseOne Enterprise Server software is installed on NAS, a failover resource group is configured with the logical hostname resource and HA for JD Edwards EnterpriseOne Enterprise Server resource. A scalable resource group is configured with the NAS storage resource. This type of configuration is shown in [Figure 2, “JD Edwards EnterpriseOne Enterprise Server Configured for Failover with NAS,”](#) on page 25.

FIGURE 2 JD Edwards EnterpriseOne Enterprise Server Configured for Failover with NAS



- Denotes strong dependency
- ➔ Denotes offline-restart dependency
- ▶ Denotes strong positive affinity
- Denotes a resource group
- Denotes a resource

Multi-Instance Configuration

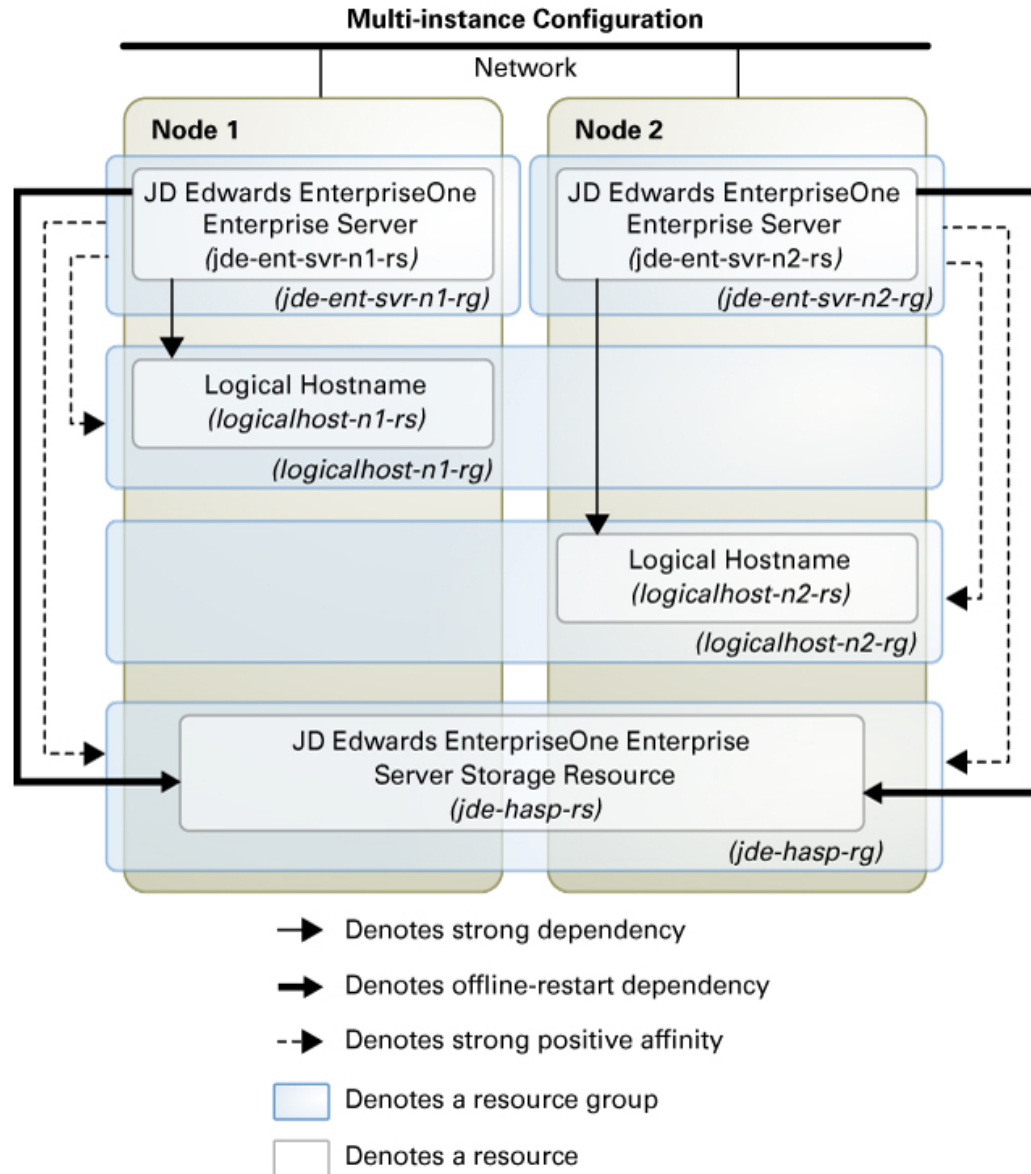
Multi-instance configuration is an application deployment topology where multiple instances of the same application provide an aggregation of services. This topology can be achieved independently from using a data service because you can manually start and stop the instances on the cluster nodes. When you require high availability of such instances, you can enable a data service for the instances by creating multiple single-node resource groups or a few multi-master resource groups.

[Figure 3, “HA for JD Edwards EnterpriseOne Enterprise Server Configured as a Multi-Instance Application With Single-Node Resource Groups,” on page 27](#) illustrates a multi-instance configuration using single-node resource groups. A single-node resource group is created for each of the HA for JD Edwards EnterpriseOne Enterprise Server resources. Each resource group has a strong positive affinity on a storage resource group whose primary node contains the HA for JD Edwards EnterpriseOne Enterprise Server resource group.

Each JD Edwards EnterpriseOne Enterprise Server instance runs on each node of the cluster and is managed by a single-node resource group. The suitability is entirely based on the need and environment. As an example use case, a multi-instance configuration can be used when initially there is only one JD Edwards EnterpriseOne Enterprise Server instance in the failover configuration on Node 1 and Node 2, but later there is a need to create another JD Edwards EnterpriseOne Enterprise Server instance.

An easier reconfiguration path would be to reduce the node list of the first Enterprise Server instance resource to Node 1, then create the second server instance in another single-node resource group using Node 2 in the `NodeList` property. While doing so, it is also important to ensure that rest of the resources follow the multi-instance configuration described in the following diagram.

FIGURE 3 HA for JD Edwards EnterpriseOne Enterprise Server Configured as a Multi-Instance Application With Single-Node Resource Groups



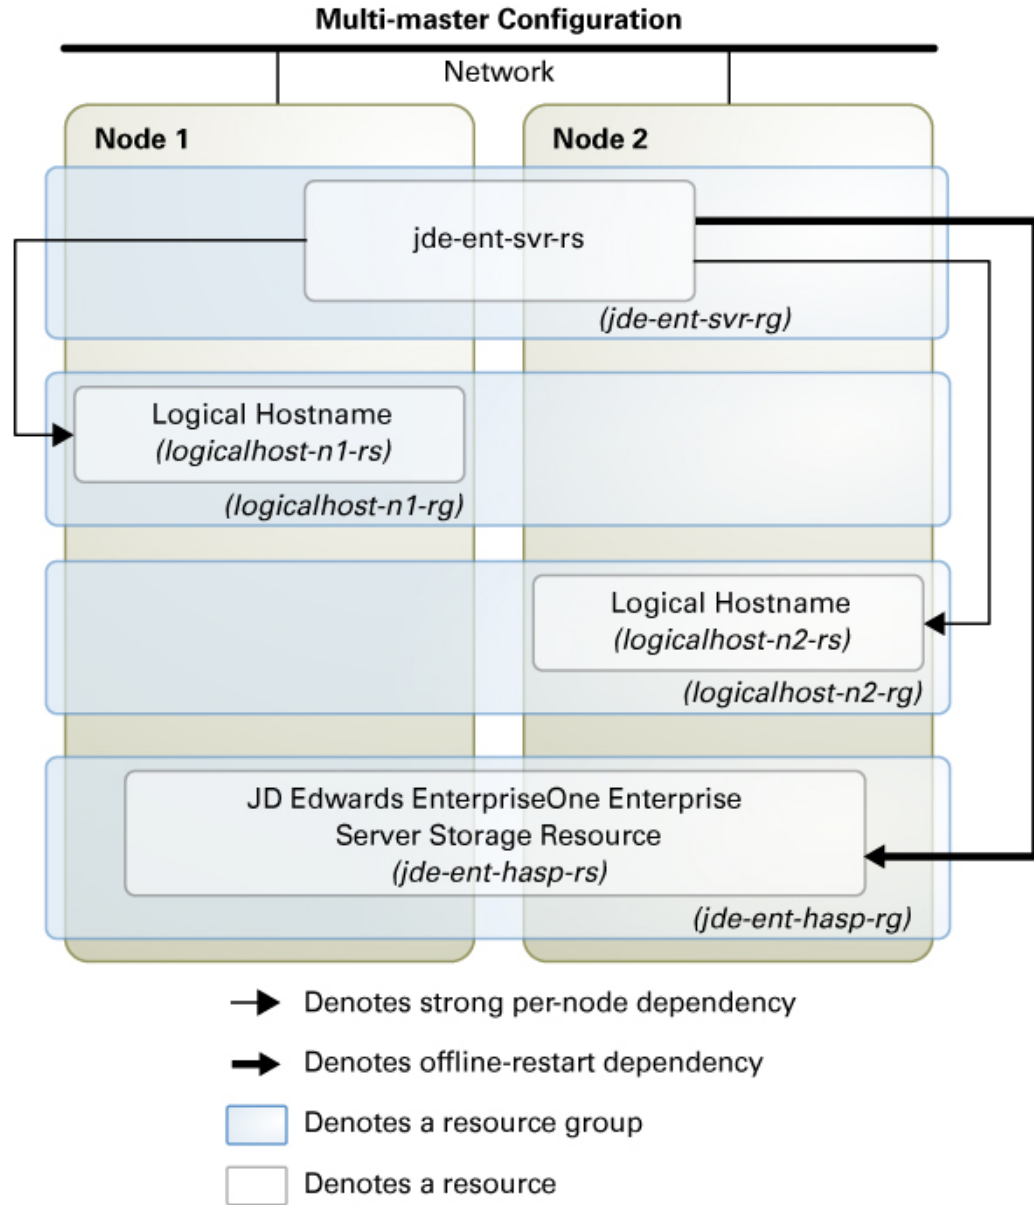
Multi-Master Configuration

The multi-master configuration is similar to a multi-instance configuration, except that all the instances of JD Edwards EnterpriseOne Enterprise Server are managed by a single resource. In a multi-master data service configuration, the JD Edwards EnterpriseOne Enterprise Server resource is created in a scalable resource group. The server resource is online on multiple nodes at the same time.

[Figure 4, “HA for JD Edwards EnterpriseOne Enterprise Server in a Multi-Master Configuration Using a Scalable Resource Group,” on page 29](#) illustrates a multi-master configuration using a scalable resource group. A scalable resource group is created for managing multiple instances of the JD Edwards EnterpriseOne Enterprise Server in a single resource. The resource group has a strong positive affinity on a storage resource group whose primary node contains the JD Edwards EnterpriseOne Enterprise Server resource group.

The suitability is entirely based on the need and environment. If there are multiple JD Edwards EnterpriseOne Enterprise Server instances configured from the beginning, a multi-master configuration can be used to make each instance of Enterprise Server run on each node of the cluster and be managed by a single resource.

FIGURE 4 HA for JD Edwards EnterpriseOne Enterprise Server in a Multi-Master Configuration Using a Scalable Resource Group



Preparing the Nodes and Disks

This section contains the following procedures to prepare the nodes and disks:

- [“How to Prepare the Nodes for a Failover Configuration” on page 30](#)
- [“How to Prepare the Nodes for a Multi-Instance Configuration” on page 32](#)
- [“How to Prepare the Nodes for a Multi-Master Configuration” on page 33](#)

Note - The following steps are based on the configurations that are described in [“Standard Data Service Configurations” in Oracle Solaris Cluster Data Service for Oracle WebLogic Server Guide](#).

▼ How to Prepare the Nodes for a Failover Configuration

Use this procedure to prepare for the installation and configuration of HA for JD Edwards EnterpriseOne Enterprise Server in a failover configuration.

Before You Begin Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

1. **Create a failover resource group to hold the logical hostname and the HA for JD Edwards EnterpriseOne Enterprise Server resources.**

```
# clresourcegroup create [-n nodelist] jde-failover-rg
```

2. **Add the logical hostname resource.**

The logical hostname `logicalhost-1` is where JD Edwards EnterpriseOne Enterprise Server will listen.

```
# clreslogicalhostname create -g jde-failover-rg -h logicalhost-1 logicalhost-resource
```

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.

3. **Create a highly-available local file system managed by the HAStoragePlus resource.**

```
# clresource create -g jde-failover-rg -t SUNW.HAStoragePlus \
```

```
-p FileSystemMountPoints=jde-mount-points jde-hasp-resource
```

For more information, see [“How to Set Up the HAStoragePlus Resource Type for New Resources”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide*.

4. Bring the resource group online.

```
# clresourcegroup online -eM jde-failover-rg
```

5. If you are running HA for JD Edwards EnterpriseOne Enterprise Server on NAS storage, create a file system that is managed by a SUNW.ScalMountPoint resource.

a. Create a scalable resource group.

```
# clresourcegroup create -p RG_mode=Scalable \
-p Desired_primaries=number-active-primaries \
-p Maximum_primaries=maximum-number-active-primaries scalmnt-rg
```

See the [clresourcegroup\(1CL\)](#) man page for information about clresourcegroup properties.

b. Register the SUNW.ScalMountPoint resource type.

```
# clresourcetype register SUNW.ScalMountPoint
```

c. Create a SUNW.ScalMountPoint resource.

Specify to the TargetFileSystem property the NAS NFS file system and path.

```
# clresource create -g scalmnt-rg \
-t SUNW.ScalMountPoint \
-p TargetFileSystem=nas-device:path \
-p FileSystemType=nas \
-p MountPointDir=file-system-mount-point scalmnt-resource
```

See the [SUNW.ScalMountPoint\(5\)](#) man page for information about SUNW.ScalMountPoint resource type properties.

d. Bring the resource group online.

```
# clresourcegroup online -eM scalmnt-rg
```

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

Use this procedure to prepare for the installation and configuration of JD Edwards EnterpriseOne Enterprise Server in a multi-instance configuration that uses single-node resource groups.

Before You Begin Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

1. Create failover resource groups for the logical hostnames for the HA for JD Edwards EnterpriseOne Enterprise Server resources.

```
# clresourcegroup create -p Nodelist=node1,node2 logicalhost-n1-rg
# clresourcegroup create -p Nodelist=node2,node1 logicalhost-n2-rg
```

2. Create logical hostname resources in the resource groups.

```
# clreslogicalhostname create -g logicalhost-n1-rg -h logicalhost-n1 logicalhost-n1-rs
# clreslogicalhostname create -g logicalhost-n2-rg -h logicalhost-n2 logicalhost-n2-rs
```

3. Bring the logical hostname resource groups online.

```
# clresourcegroup online -eM logicalhost-n1-rg
# clresourcegroup online -eM logicalhost-n2-rg
```

4. Set the Failback property for each logical hostname resource group.

```
# clresourcegroup set -p Failback=TRUE logicalhost-n1-rg
# clresourcegroup set -p Failback=TRUE logicalhost-n2-rg
```

5. Configure a cluster file system to be managed by an HAStoragePlus resource in a scalable resource group.

For more information, see [“How to Set Up the HAStoragePlus Resource Type for New Resources”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide*.

a. Create a scalable resource group to contain the HAStoragePlus resource.

The optional `-n nodelist` specifies where the JD Edwards EnterpriseOne Enterprise Server instances can be brought online. If you use this option, observe the following requirements:

- 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.

- The application resources that depend on the file system must have a strong resource dependency set to the `HASStoragePlus` resource.

```
# clresourcegroup create -p Maximum primaries=maximum-number-active-primaries \  
-p Desired primaries=number-active-primaries \  
[-n nodelist] jde-ent-hasp-rg
```

See the `clresourcegroup(1CL)` man page for information about `clresourcegroup` options and properties.

b. Register the `SUNW.HASStoragePlus` resource type.

```
# clresourcetype register SUNW.HASStoragePlus
```

For more information about the `SUNW.HASStoragePlus` resource type, see the `SUNW.HASStoragePlus(5)` man page.

c. Create an `HASStoragePlus` resource.

Specify the value of the `FileSystemMountPoints` property as follows:

- If you installed the JD Edwards EnterpriseOne Enterprise Server instances on separate directories under the same mount point, specify only that one mount point.
- If you installed each JD Edwards EnterpriseOne Enterprise Server instance on a separate mount point, specify a comma-separated list of all the mount points.

```
# clresource create -g jde-ent-hasp-rg \  
-t SUNW.HASStoragePlus \  
-p FileSystemMountPoints=mount-point \  
-p AffinityOn=false jde-hasp-rs
```

See the `SUNW.HASStoragePlus(5)` man page for more information about `SUNW.HASStoragePlus` resource type properties.

6. Bring online the resource group.

```
clresourcegroup online -eM jde-ent-hasp-rg
```

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

Use this procedure to prepare for the installation and configuration of HA for JD Edwards EnterpriseOne Enterprise Server in a multi-master configuration.

This procedure assume two nodes in the cluster and implements the configuration shown in [Figure 4, “HA for JD Edwards EnterpriseOne Enterprise Server in a Multi-Master Configuration Using a Scalable Resource Group,”](#) on page 29.

Before You Begin Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

1. Create failover resource groups for the logical hostname for the HA for JD Edwards EnterpriseOne Enterprise Server resources.

```
# clresourcegroup create -p Nodelist=node1,node2 logicalhost-n1-rg
# clresourcegroup create -p Nodelist=node2,node1 logicalhost-n2-rg
```

2. Create logical hostname resources in the resource groups.

```
# clreslogicalhostname create -g logicalhost-n1-rg -h logicalhost-n1 logicalhost-n1-rs
# clreslogicalhostname create -g logicalhost-n2-rg -h logicalhost-n2 logicalhost-n2-rs
```

3. Bring the logical hostname resource groups online.

```
# clresourcegroup online -eM logicalhost-n1-rg
# clresourcegroup online -eM logicalhost-n2-rg
```

4. Set the Failback property for each logical hostname resource group.

```
# clresourcegroup set -p Failback=TRUE logicalhost-n1-rg
# clresourcegroup set -p Failback=TRUE logicalhost-n2-rg
```

5. Configure a cluster file system to be managed by an HAStoragePlus resource in a scalable resource group.

For more information, see [“How to Set Up the HAStoragePlus Resource Type for New Resources”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide*.

a. Create a scalable resource group.

The optional `-n nodelist` specifies where the JD Edwards EnterpriseOne Enterprise Server instances can be brought online.

```
# clresourcegroup create -p Maximum primaries=maximum-number-active-primaries \
-p Desired primaries=number-active-primaries \
[-n nodelist] jde-ent-hasp-rg
```

See the `clresourcegroup(1CL)` man page for information about `clresourcegroup` options and properties.

b. Register the SUNW.HAStoragePlus resource type.

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

c. Create an HAStoragePlus resource.

Specify the value of the `FileSystemMountPoints` property as follows:

- If you installed the JD Edwards EnterpriseOne Enterprise Server instances on separate directories under the same mount point, specify only that one mount point.
- If you installed each JD Edwards EnterpriseOne Enterprise Server instance on a separate mount point, specify a comma-separated list of all the mount points.

```
# clresource create -g jde-ent-hasp-rg \  
-t SUNW.HAStoragePlus \  
-p FileSystemMountPoints=mount-point \  
-p AffinityOn=false jde-ent-hasp-rs
```

See the [SUNW.HAStoragePlus\(5\)](#) man page for more information about SUNW.HAStoragePlus resource type properties.

d. Bring online the resource group.

```
clresourcegroup online -eM jde-ent-hasp-rg
```

Installing and Configuring JD Edwards EnterpriseOne Enterprise Server Software

This section contains the procedure to install and configure a JD Edwards EnterpriseOne Enterprise Server as a cluster resource.

▼ How to Install and Configure JD Edwards EnterpriseOne Enterprise Server Software

1. On the cluster member where the HA for JD Edwards EnterpriseOne Enterprise Server resource group is online, assume the `root` role.
2. Follow the instructions in the appropriate JD Edwards EnterpriseOne documentation for your version of the product.

See the [JD Edwards EnterpriseOne Tools Documentation Library](#).

Ensure that you observe the configuration requirements in “[Configuration Requirements for HA for JD Edwards EnterpriseOne Enterprise Server](#)” on page 19.

- 3. When you provide details about JD Edwards EnterpriseOne Enterprise Server on Enterprise Server Revisions of the Installation Plan using the Solution Explorer, type the hostname of the logical host resource in the Machine Name field.**

For example, if the hostname of the logical host resource is `logicalhost-1`, type `LOGICALHOST-1` in the Machine Name field.

Note - If you are using HA for Oracle WebLogic Server to configure a JD Edwards EnterpriseOne HTML Web Server, the Machine Name field in the HTML Web Server Revisions of the Installation Plan must contain the hostname of the logical host resource that is configured for the WebLogic Server resource.

- 4. For multi-master configurations, ensure that each JD Edwards EnterpriseOne Enterprise Server is configured to listen on each logical hostname or its IP address.**

To create multiple JD Edwards EnterpriseOne Enterprise Servers, follow the steps in the *JD Edwards EnterpriseOne Tools Server and Workstation Administration Guide* (http://docs.oracle.com/cd/E17984_01/doc.898/e14718/unix_and_linux_servers.htm#g8d6ab57f7cedeaac_ef90c_10a77c8e3f7__6057).

- 5. After the installation is complete, on each JD Edwards EnterpriseOne Enterprise Server you configured, set environment variables for the profile of the Enterprise Server user.**

In addition to the required environment variables that are explained in the *JD Edwards EnterpriseOne Installation Guide*, you must set up the following variables after you configure the JD Edwards EnterpriseOne Enterprise Server:

- `SC_LHOSTNAME`
- `LD_PRELOAD_32`
- `LD_PRELOAD_64`

Set `SC_LHOSTNAME` to the logical hostname where the JD Edwards EnterpriseOne Enterprise Server can be reached from the web tier. For more details, refer to the `libscho.st.so.1(1)` man page.

Set these environment variables for the profile of the user that operates the JD Edwards EnterpriseOne Enterprise Server. Ensure that the login for the user is noninteractive. If you invoke the root user, you must see these variables displayed. In the following example, the user name is `jde910`:

```
# su - jde910 -c "env"
```

Symbolic links are created for the `libscho.st.so.1` file to `/usr/lib/secure/`.

```
# ln -s /usr/cluster/lib/libscho.st.so.1 /usr/lib/secure/libscho.st.so.1
# ln -s /usr/cluster/lib/64/libscho.st.so.1 /usr/lib/secure/64/libscho.st.so.1
```

The following is part of a sample `.profile` file that contains the environment variables for the example `jde910` user:

```
LD_PRELOAD_32=/usr/lib/secure/libscho.st.so.1
LD_PRELOAD_64=/usr/lib/secure/64/libscho.st.so.1
SC_LHOSTNAME=logicalhost-1
export LD_PRELOAD_32 LD_PRELOAD_64 SC_LHOSTNAME
```

Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration

This section contains the following procedures to verify the JD Edwards EnterpriseOne Enterprise Server installation and configuration:

- [“How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Failover Services” on page 37](#)
- [“How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Multi-Master and Multi-Instance Services” on page 39](#)

▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Failover Services

Use this procedure to verify the JD Edwards EnterpriseOne Enterprise Server installation and configuration. This procedure does not verify that your application is highly available because you have not yet installed the data service.

Note - Use the following steps for the failover configuration that is described in [“Standard Data Service Configurations” in *Oracle Solaris Cluster Data Service for Oracle WebLogic Server Guide*](#). If you are planning a different configuration, you must configure the JD Edwards EnterpriseOne Enterprise Server accordingly.

1. **Verify that the *logicalhost-1* configured in *jde-failover-rg* is online.**
2. **If you are using a database, verify that the database is running.**
3. **Go to the directory that contains the JD Edwards EnterpriseOne Enterprise Server installation.**
4. **Assume the *root* role, log in to the node that currently hosts the *jde-failover-rg* resource group, and start the JD Edwards EnterpriseOne Enterprise Server by running the *RunOneWorld.sh* script.**

In the following command, *jde-user* represents the installation username of the JD Edwards EnterpriseOne Enterprise Server.

```
# su - jde-user -c "/path-to-jde-binaries/RunOneWorld.sh"
```

5. **Verify the status of the JD Edwards EnterpriseOne Enterprise Server.**

```
# su - jde-user -c "/path-to-jde-binaries/jdejobs"  
# su - jde-user -c "/path-to-jde-binaries/porttest user password environment"
```

6. **Stop the JD Edwards EnterpriseOne Enterprise Server.**

```
# su - jde-user -c "/path-to-jde-binaries/EndOneWorld.sh"
```

7. **Switch the *jde-failover-rg* resource group to another cluster member.**

```
# clresourcegroup switch -n node jde-failover-rg
```

8. **Repeat all steps until you have tested all the potential nodes where the JD Edwards EnterpriseOne Enterprise Server can run.**

If the clients can successfully connect to the JD Edwards EnterpriseOne Enterprise Server on each of the potential master nodes, you have successfully configured the JD Edwards EnterpriseOne Enterprise Server to work with HA for JD Edwards EnterpriseOne Enterprise Server.

9. **Stop the JD Edwards EnterpriseOne Enterprise Server instance.**

```
# su - jde-user -c "/path-to-jde-binaries/EndOneWorld.sh"
```

Next Steps Go to [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package”](#) on page 40.

▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Multi-Master and Multi-Instance Services

Use this procedure to verify the JD Edwards EnterpriseOne Enterprise Server installation and configuration for both multi-master and multi-instance services. This procedure does not verify that your application is highly available because you have not installed the data service yet.

Note - Use the following steps for the multi-master and multi-instance configuration that is described in [“Standard Data Service Configurations” in Oracle Solaris Cluster Data Service for Oracle WebLogic Server Guide](#). If you are planning a different configuration, you must configure the JD Edwards EnterpriseOne Enterprise Server accordingly.

1. **Verify that the *logicalhost-n1-rs* and *logicalhost-n2-rs* configured in *logicalhost-n1-rg* and *logicalhost-n2-rg* are online on *node1* and *node2* respectively.**

2. **If you are using a database, verify that the database is running.**

3. **Go to the directory that contains the JD Edwards EnterpriseOne Enterprise Server START script on *node1*.**

This directory contains the JD Edwards EnterpriseOne Enterprise Server *instance 1* installation.

4. **Assume the root role, log in to *node1*, and start the server.**

In the following command, *jde-user* represents the installation username of JD Edwards EnterpriseOne Enterprise Server *instance 1*.

```
# su - jde-user -c "/path-to-jde-binaries/RunOneWorld.sh"
```

5. **Verify the status of JD Edwards EnterpriseOne Enterprise Server *instance 1*.**

```
# su - jde-user -c "/path-to-jde-binaries/jdejobs"  
# su - jde-user -c "/path-to-jde-binaries/porttest user password environment"
```

6. **Assume the root role, log in to *node2*, and start the server.**

In the following command, *jde-user2* represents the installation username of JD Edwards EnterpriseOne Enterprise Server *instance 2*.

```
# su - jde-user2 -c "/path-to-jde-binaries/RunOneWorld.sh"
```

7. **Verify the status of JD Edwards EnterpriseOne Enterprise Server *instance 2*.**

```
# su - jde-user2 -c "/path-to-jde-binaries/jdejobs"  
# su - jde-user2 -c "/path-to-jde-binaries/porttest user password environment"
```

If the clients can successfully connect to the JD Edwards EnterpriseOne Enterprise Servers on each of the potential master nodes, you have successfully configured the JD Edwards EnterpriseOne Enterprise Servers to work with the HA for JD Edwards EnterpriseOne Enterprise Server in a multi-master or multi-instance configuration.

8. Stop the JD Edwards EnterpriseOne Enterprise Server instances.

```
# su - jde-user -c "/path-to-jde-binaries/EndOneWorld.sh"  
# su - jde-user2 -c "/path-to-jde-binaries/EndOneWorld.sh"
```

Next Steps Go to [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package”](#) on page 40.

Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package

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

▼ How to Install the HA for JD Edwards EnterpriseOne Enterprise Server Package

Perform this procedure on each cluster node where you want the HA for JD Edwards EnterpriseOne Enterprise Server software to run.

1. **On the cluster node where you are installing the data service package, assume the root role.**
2. **Ensure that the data service package is available from the configured publisher and that the `solaris` and `ha-cluster` publishers are valid.**

```
# pkg list -a ha-cluster/data-service/jd-edwards-enterpriseone  
# pkg publisher  
PUBLISHER                                TYPE      STATUS   P  LOCATION  
solaris                                  origin   online   F  solaris-repository
```



```
ha-cluster                                origin  online  F  ha-cluster-repository
```

For information about setting the `solaris` publisher, see [“Adding, Modifying, or Removing Package Publishers”](#) in *Adding and Updating Software in Oracle Solaris 11.3*.

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,”](#) in *Adding and Updating Software in Oracle Solaris 11.3*.

3. Install the HA for JD Edwards EnterpriseOne Enterprise Server software package.

```
# pkg install ha-cluster/data-service/jd-edwards-enterpriseone
```

4. Verify that the package installed successfully.

```
# pkg info ha-cluster/data-service/jd-edwards-enterpriseone
```

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 your software, see [Chapter 11, “Updating Your Software”](#) in *Oracle Solaris Cluster 4.3 System Administration Guide*.

Registering and Configuring HA for JD Edwards EnterpriseOne Enterprise Server

This section contains the following procedures to configure HA for JD Edwards EnterpriseOne Enterprise Server:

- [“How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server for Failover”](#) on page 42
- [“How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Data Service”](#) on page 45

- [“How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Instance Data Service” on page 48](#)

▼ How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server for Failover

Use this procedure to configure HA for JD Edwards EnterpriseOne Enterprise Server as a failover data service. The following steps assume that you are using the failover configuration as described in [“Standard Data Service Configurations” in Oracle Solaris Cluster Data Service for Oracle WebLogic Server Guide](#) and have configured the JD Edwards EnterpriseOne Enterprise Server as described in [“Installing and Configuring JD Edwards EnterpriseOne Enterprise Server Software” on page 35](#).

Before You Begin Ensure that the data service package is installed. For more information, see [“How to Install and Configure JD Edwards EnterpriseOne Enterprise Server Software” on page 35](#).

1. **On the cluster node that hosts the JD Edwards EnterpriseOne Enterprise Server, assume the root role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.**

2. **Register the `ORCL.JDE_Enterprise_Server` resource type.**

```
# clresourcetype register ORCL.JDE_Enterprise_Server
```

3. **Create an HA for JD Edwards EnterpriseOne Enterprise Server resource in the failover resource group.**

Include the following settings:

- Set the `JDE_USER` and `JDE_HOME` properties with the JD Edwards EnterpriseOne Enterprise Server installation username and installation directory, respectively.
- Set the `Resource_dependencies_offline_restart` property with the name of the `SUNW.ScalMountPoint` or `SUNW.HASStoragePlus` storage resource you are using.

```
# clresource create -g jde-failover-rg -d \  
-t ORCL.JDE_Enterprise_Server \  
-p JDE_USER=jde-user \  
-p JDE_HOME=jde-home-directory \  
-p Resource_dependencies=logicalhost-resource \  
-p Resource_dependencies_offline_restart=jde-hasp-resource  
jde-ent-svr-resource
```

For more information, see [Appendix A, “HA for JD Edwards EnterpriseOne Enterprise Server Extension Properties”](#) and the `r_properties(5)` man page.

4. **If the database is deployed on the same global cluster, configure a strong dependency to the resources for the database instance and database listener.**

```
# clresource set -p Resource_dependencies+=db-instance-resource,db-listener-resource \
jde-ent-svr-resource
```

5. **Enable the HA for JD Edwards EnterpriseOne Enterprise Server resource.**

If multiple instances were created, repeat this step for each JD Edwards EnterpriseOne Enterprise Server instance.

```
# clresource status
# clresource enable jde-ent-svr-resource
```

Example 1 Configuring HA for JD Edwards EnterpriseOne Enterprise Server as a Failover Service on a Two-Node Cluster With Traditional File Storage

This example uses the following sample component names:

- Node names: node1, node2
- Logical hostname: logicalhost-n1
- Resource group: jde-failover-rg (for all resources)
- Resources:
 - logicalhost-resource (logical host resource)
 - jde-hasp-resource (HASStoragePlus resource)
 - jde-ent-svr-resource (JD Edwards EnterpriseOne Enterprise Server resource)

Create the resource group to contain all resources

```
# clresourcegroup create jde-failover-rg
```

Add the logical hostname resource to the resource group

```
# clreslogicalhostname create -g jde-failover-rg -h logicalhost-1
```

Add the HASStoragePlus resource to the resource group

```
# clresource create -g jde-failover-rg -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=/global/jde jde-hasp-resource
```

(Bring the failover resource group online.)

```
# clresourcegroup online -eM jde-failover-rg
```

Register the resource type for HA for JD Edwards EnterpriseOne Enterprise Server

```
# clresourcetype register ORCL.JDE_Enterprise_Server
```

```

    Add a JD Edwards EnterpriseOne Enterprise Server resource instance
# clresource create -g jde-failover-rg -d \
-t ORCL.JDE_Enterprise_Server \
-p JDE_USER=jde910 \
-p JDE_HOME=/global/jde/jdedwardsppack \
-p Resource_dependencies=logicalhost-resource,db-instance-resource \
-p Resource_dependencies_offline_restart=jde-hasp-resource jde-ent-svr-resource

    Enable the JD Edwards EnterpriseOne Enterprise Server resource instance
# clresource enable jde-ent-svr-resource

```

Example 2 Registering HA for JD Edwards EnterpriseOne Enterprise Server as a Failover Service on a Two-Node Cluster With NAS

This example uses the following sample component names:

- Node names: node1, node2
- Logical hostname: logicalhost-1
- Resource groups:
 - jde-failover-rg (for all resources)
 - scalmnt-rg (scalable mount-point resource group)
- Resources:
 - logicalhost-resource (logical host resource)
 - scalmnt-resource (scalable mount-point resource)
 - jde-ent-svr-resource (JD Edwards EnterpriseOne Enterprise Server resource)

```

    Create the resource group to contain all resources
# clresourcegroup create jde-failover-rg

    Add the logical hostname resource to the resource group
# clreslogicalhostname create -g jde-failover-rg -h logicalhost-1

Create a scalable mount-point resource group
# clresourcegroup create -p RG_mode=Scalable \
-p Desired primaries=2 \
-p Maximum primaries=2 scalmnt-rg

    Register the resource type for the scalable mount-point resource
# clresourcetype register SUNW.ScalMountPoint

    Add a scalable mount-point resource to scalmnt-rg
# clresource create -g scalmnt-rg -t SUNW.ScalMountPoint \
-p TargetFileSystem=jde4nas:/export/home/jde \
-p FileSystemType=nas \

```

```
-p MountPointDir=/global/jde scalmnt-resource

    Bring the scalable mount-point resource group online
# clresourcegroup online -eM scalmnt-rg

    Set strong positive RG_affinities on the storage resource group
# clresourcegroup set -p RG_affinities=++scalmnt-rg jde-failover-rg

    Bring the failover resource group online
# clresourcegroup online -eM jde-failover-rg

Register the resource type for HA for JD Edwards EnterpriseOne Enterprise Server
# clresourcetype register ORCL.JDE_Enterprise_Server

    Add an HA for JD Edwards EnterpriseOne Enterprise Server resource instance
# clresource create -g jde-failover-rg -d \
-t ORCL.JDE_Enterprise_Server \
-p JDE_USER=jde910 \
-p JDE_HOME=/global/jde/jdedwardsppack \
-p Resource_dependencies=logicalhost-resource,db-instance-resource \
-p Resource_dependencies_offline_restart=scalmnt-resource jde-ent-svr-resource

Enable the HA for JD Edwards EnterpriseOne Enterprise Server resource instance
# clresource enable jde-ent-svr-resource
```

▼ How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Data Service

Before You Begin Use this procedure to configure HA for JD Edwards EnterpriseOne Enterprise Server in a multi-master configuration. The following steps assume that you are using the configuration that is described in [“Multi-Master Configuration” on page 28](#) and have configured the server as described in [“Installing and Configuring JD Edwards EnterpriseOne Enterprise Server Software” on page 35](#). If you are planning a different configuration, you must configure the JD Edwards EnterpriseOne Enterprise Server resources accordingly.

1. **Add the HA for JD Edwards EnterpriseOne Enterprise Server resource type.**
2. **Create a scalable resource group for the JD Edwards EnterpriseOne Enterprise Server.**

```
# clresourcetype register ORCL.JDE_Enterprise_Server

# clresourcegroup create -p Maximum primaries=m \
```

-p Desired_primaries=n *jde-ent-svr-rg*

-p Maximum_primaries=m

Specifies the maximum number of active primary nodes permitted for this resource group.
The default is one.

-p Desired_primaries=n

Specifies the desired number of active primary nodes permitted for this resource group.
The default is one.

jde-ent-svr-rg

Specifies the resource group name.

3. Bring the resource group online.

```
# clresourcegroup online -eM jde-ent-svr-rg
```

4. Create a multi-master resource to start the JD Edwards EnterpriseOne Enterprise Server instances simultaneously.

The following example starts the JD Edwards EnterpriseOne Enterprise Server instances on *node1* and *node2*.

Set the `JDE_USER` and `JDE_HOME` properties with the JD Edwards EnterpriseOne Enterprise Server installation username and installation directory, respectively, for each node.

Set the `Resource_dependencies_offline_restart` property with the name of the SUNW.ScalMountPoint or SUNW.HASStoragePlus storage resource you are using.

```
# clresource create -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-rg \  
-p JDE_USER{node1}=jde-user \  
-p JDE_USER{node2}=jde-user2 \  
-p JDE_HOME{node1}=jde-home-directory \  
-p JDE_HOME{node2}=jde-home-directory \  
-p Resource_dependencies=logicalhost-n1-rs@node1,logicalhost-n2-rs@node2,db-instance-resource \  
-p Resource_dependencies_offline_restart=jde-ent-hasp-rs \  
jde-ent-svr-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.

Example 3 Configuring HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Service

This example uses the following sample component names:

- Node names: node1, node2
- Logical hostnames: logicalhost-n1, logicalhost-n2
- Resource groups:
 - logicalhost-n1-rg (for the logicalhost-n1-rs resource)
 - logicalhost-n2-rg (for the logicalhost-n2-rs resource)
 - jde-ent-hasp-rg (for the scalable HAStoragePlus resource)
 - jde-ent-svr-rg (for the JD Edwards EnterpriseOne Enterprise Server resource)
- Resources:
 - logicalhost-n1-rs logicalhost-n2-rs (logical hostname resources)
 - jde-ent-hasp-rs (scalable HAStoragePlus resource)
 - jde-ent-svr-rs (multi-master JD Edwards EnterpriseOne Enterprise Server resource)
- JD Edwards EnterpriseOne Enterprise Server instance 1 installation directory: /global/jde/u01 Username: jde910
- JD Edwards EnterpriseOne Enterprise Server instance 2 installation directory: /global/jde/u02 Username: jde2910

Create failover resource groups to contain logical hostnames for JD Edwards EnterpriseOne Enterprise Server instances

```
# clresourcegroup create -p Nodelist=node1,node2 logicalhost-n1-rg
# clresourcegroup create -p Nodelist=node2,node1 logicalhost-n2-rg
```

Add the logical hostname resources to the failover resource groups.

```
# clreslogicalhostname create -g logicalhost-n1-rg -h logicalhost-n1 logicalhost-n1-rs
# clreslogicalhostname create -g logicalhost-n2-rg -h logicalhost-n2 logicalhost-n2-rs
```

Bring the logical hostname resource groups online

```
# clresourcegroup online -eM logicalhost-n1-rg
# clresourcegroup online -eM logicalhost-n2-rg
```

Set the Failback property for each logical hostname resource group

```
# clresourcegroup set -p Failback=TRUE logicalhost-n1-rg
# clresourcegroup set -p Failback=TRUE logicalhost-n2-rg
```

Add a scalable resource group for HAStoragePlus resource

```
# clresourcegroup create -p Maximum primaries=2 \
-p Desired primaries=2 \
jde-ent-hasp-rg
```

Register the resource type for HAStoragePlus resource

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

Add an HAStoragePlus resource to the jde-ent-hasp-rg resource group

```
# clresource create -g jde-ent-hasp-rg -t SUNW.HAStoragePlus \
```

```
-p FileSystemMountPoints=/global/jde affinityon=false jde-ent-hasp-rs

    Bring the HASToragePlus resource group online
# clresourcegroup online -eM jde-ent-hasp-rg

    Create a scalable resource group for JD Edwards EnterpriseOne Enterprise Server resource
# clresourcegroup create -p Maximum primaries=2 -p Desired primaries=2 jde-ent-svr-rg

    Add a multi-master resource for JD Edwards EnterpriseOne Enterprise Server instances
# clresource create -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-rg \
-p JDE_USER{node1}=jde910 \
-p JDE_USER{node2}=jde2910 \
-p JDE_HOME{node1}=/global/jde/u01/jdedwardsppack \
-p JDE_HOME{node2}=/global/jde/u02/jdedwardsppack \
-p Resource_dependencies=logicalhost-n1-rs@node1,logicalhost-n2-rs@node2,db-instance-
resource \
-p Resource_dependencies_offline_restart=jde-ent-hasp-rs jde-ent-svr-rs

    Bring the resource group online
# clresourcegroup online -eM jde-ent-svr-rg
```

▼ How to Register and Configure HA for JD Edwards EnterpriseOne Enterprise Server as a Multi-Instance Data Service

Before You Begin Use this procedure to configure an HA for JD Edwards EnterpriseOne Enterprise Server multi-instance configuration that uses single-node resource groups. A single node resource group is created for each of the server resources managing a server instance. Each resource group has a strong positive affinity on a storage resource group.

This procedure assumes that you are using a global cluster. If you are performing this configuration in a zone cluster, add the `-Z zone-name` option to the command.

1. **Create a single-node resource group for a JD Edwards EnterpriseOne Enterprise Server instance 1, with strong positive affinities for the storage resource group.**

```
# clresourcegroup create -p Nodelist=node1 \
-p RG_affinities=++jde-hasp-rg jde-ent-svr-n1-rg
```

2. **Bring the resource group online.**

```
# clresourcegroup online -eM jde-ent-svr-n1-rg
```


3. Create a single-node resource group for a JD Edwards EnterpriseOne Enterprise Server instance 2, with strong positive affinities for the storage resource group.

```
# clresourcegroup create -p Nodelist=node2 \  
-p RG_affinities=++jde-hasp-rg jde-ent-svr-n2-rg
```

4. Bring the second resource group online.

```
# clresourcegroup online -eM jde-ent-svr-n2-rg
```

5. Create the HA for JD Edwards EnterpriseOne Enterprise Server resource and set its properties.

Set the JDE_USER and JDE_HOME properties with the JD Edwards EnterpriseOne Enterprise Server installation username and installation directory, respectively, for instance 1.

```
# clresource create -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-n1-rg \  
-p JDE_USER=jde-user \  
-p JDE_HOME=jde-installation-dir \  
-p Resource_dependencies=logicalhost-n1-rs \  
-p Resource_dependencies_offline_restart=jde-hasp-rs \  
jde-ent-svr-n1-rs
```

The resource is created in the enabled state.

6. Create the second HA for JD Edwards EnterpriseOne Enterprise Server resource for instance 2 and set its properties.

Set the JDE_USER and JDE_HOME properties with the JD Edwards EnterpriseOne Enterprise Server installation username and installation directory, respectively, for instance 2.

```
# clresource create -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-n2-rg \  
-p JDE_USER=jde-user2 \  
-p JDE_HOME=jde-installation-dir \  
-p Resource_dependencies=logicalhost-n2-rs \  
-p Resource_dependencies_offline_restart=jde-hasp-rs \  
jde-ent-svr-n2-rs
```

The resource is created in the enabled state.

Example 4 Configuring HA for JD Edwards EnterpriseOne Enterprise Server in a Multi-Instance Service

This example uses the following sample component names:

- Node names: node1, node2
- Logical hostnames: logicalhost-n1, logicalhost-n2
- Resource groups:
 - logicalhost-n1-rg (for the logicalhost-n1-rs resource)

- logicalhost-n2-rg (for the logicalhost-n2-rs resource)
- jde-hasp-rg (for the scalable HAStoragePlus resource)
- jde-ent-svr-n1-rg, jde-ent-svr-n2-rg (for the JD Edwards EnterpriseOne Enterprise Server resources)
- Resources:
 - logicalhost-n1-rs, logicalhost-n2-rs (logical hostname resources)
 - jde-hasp-rs (scalable HAStoragePlus resource)
 - jde-ent-svr-n1-rs, jde-ent-svr-n2-rs (multi-master JD Edwards EnterpriseOne Enterprise Server resources)
- JD Edwards EnterpriseOne Enterprise Server instance 1 installation directory: /global/jde/u01 Username: jde910
- JD Edwards EnterpriseOne Enterprise Server instance 2 installation directory: /global/jde/u02 Username: jde2910

Create failover resource groups to contain logical hostnames for JD Edwards EnterpriseOne Enterprise Server instances

```
# clresourcegroup create -p Nodelist=node1,node2 logicalhost-n1-rg
# clresourcegroup create -p Nodelist=node2,node1 logicalhost-n2-rg
```

Add the logical hostname resources to the failover resource groups

```
# clreslogicalhostname create -g logicalhost-n1-rg -h logicalhost-n1 logicalhost-n1-rs
# clreslogicalhostname create -g logicalhost-n2-rg -h logicalhost-n2 logicalhost-n2-rs
```

Bring the logical hostname resource groups online

```
# clresourcegroup online -eM logicalhost-n1-rg
# clresourcegroup online -eM logicalhost-n2-rg
```

Set the Failback property for each logical hostname resource group

```
# clresourcegroup set -p Failback=TRUE logicalhost-n1-rg
# clresourcegroup set -p Failback=TRUE logicalhost-n2-rg
```

Add a scalable resource group for the HAStoragePlus resource

```
# clresourcegroup create -p Maximum primaries=2 \
-p Desired primaries=2 \
jde-hasp-rg
```

Register the resource type for HAStoragePlus resource

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

Add an HAStoragePlus resource to the jde-hasp-rg resource group

```
# clresource create -g jde-hasp-rg -t SUNW.HAStoragePlus \
-p FileSystemMountPoints=/global/jde affinityon=false jde-hasp-rs
```

Bring the HAStoragePlus resource group online

```

# clresourcegroup online -eM jde-hasp-rg

    Create single-node resource groups for JD Edwards EnterpriseOne Enterprise Server
    instances 1 and 2
# clresourcegroup create -p Nodelist=node1 \
-p RG_affinities=++jde-hasp-rg jde-ent-svr-n1-rg

# clresourcegroup create -p Nodelist=node2 \
-p RG_affinities=++jde-hasp-rg jde-ent-svr-n2-rg

    Bring the resource groups online
# clresourcegroup online -eM jde-ent-svr-n1-rg
# clresourcegroup online -eM jde-ent-svr-n2-rg

    Add resources for JD Edwards EnterpriseOne Enterprise Server instance 1 and 2
# clresource create -d -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-n1-rg \
-p JDE_USER=jde910 \
-p JDE_HOME=/global/jde/u01/jdedwardsppack \
-p Resource_dependencies=logicalhost-n1-rs,db-instance-resource \
-p Resource_dependencies_offline_restart=jde-hasp-rs \
jde-ent-svr-n1-rs

# clresource create -d -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-n2-rg \
-p JDE_USER=jde2910 \
-p JDE_HOME=/global/jde/u02/jdedwardsppack \
-p Resource_dependencies=logicalhost-n2-rs,db-instance-resource \
-p Resource_dependencies_offline_restart=jde-hasp-rs \
jde-ent-svr-n2-rs

    Bring the resources online
# clresource enable jde-ent-svr-n1-rs
# clresource enable jde-ent-svr-n2-rs

```

See Also For information about administering resources and resource groups, see [Chapter 2, “Administering Data Service Resources”](#) in *Oracle Solaris Cluster 4.3 Data Services Planning and Administration Guide*.

Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration

This section contains the following procedures to verify that you installed and configured the HA for JD Edwards EnterpriseOne Enterprise Server resource correctly:

- [“How to Verify the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Failover Configuration” on page 52](#)
- [“How to Verify the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Multi-Master or Multi-Instance Configuration” on page 53](#)

▼ How to Verify the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Failover Configuration

Use this procedure to verify that you installed and configured the JD Edwards EnterpriseOne Enterprise Server correctly.

1. **On the cluster node that hosts the resource group that contains the HA for JD Edwards EnterpriseOne Enterprise Server resource, assume the `root` role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.**

2. **Create a failover resource group.**

```
# clresourcegroup create jde-failover-rg
```

3. **Switch the resource group to another cluster member.**

```
# clresourcegroup switch -n node jde-failover-rg
```

4. **Verify the status of the JD Edwards EnterpriseOne Enterprise Server.**

```
# su - jde-user -c "/path-to-jde-binaries/jdejobs"  
# su - jde-user -c "/path-to-jde-binaries/porttest user password environment"
```

5. **Verify the status of the HA for JD Edwards EnterpriseOne Enterprise Server resource.**

```
# clresource status jde-ent-svr-resource
```

6. **Repeat all steps until you have tested all the potential nodes where the JD Edwards EnterpriseOne Enterprise Server can run.**

After the JD Edwards EnterpriseOne Enterprise Servers are in production under Oracle Solaris Cluster control, do not start or stop the servers manually or through the Server Management Console. If you do need to perform maintenance on the servers, you must first disable their associated resources.

▼ How to Verify the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Multi-Master or Multi-Instance Configuration

This procedure verifies that you installed and configured HA for JD Edwards EnterpriseOne Enterprise Server correctly for a multi-master or multi-instance configuration.

1. **Ensure that you can successfully connect to the JD Edwards EnterpriseOne Enterprise Servers from a client and that all the servers are running.**

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

The Fault Monitor must detect the killed process and restart the JD Edwards EnterpriseOne Enterprise Server.

3. **Verify the status of the JD Edwards EnterpriseOne Enterprise Server.**

Issue these commands for each JD Edwards EnterpriseOne Enterprise Server instance, on each cluster node where the instance can come online..

```
# su - jde-user -c "/path-to-jde-binaries/jdejobs"  
# su - jde-user -c "/path-to-jde-binaries/porttest user password environment"
```

4. **If a JD Edwards HTML Web Server is configured for the corresponding JD Edwards EnterpriseOne Enterprise Server on Oracle WebLogic Server, deploy a sample application that consists of the WebLogic server instances configured for the JD Edwards HTML Web Server.**

5. **Use a proxy web server to access the sample application.**

6. **Shut down one of the nodes and access the application using the proxy web server address.**

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 HA for JD Edwards EnterpriseOne Enterprise Server in a multi-master or multi-instance configuration.

Tip - After the JD Edwards EnterpriseOne Enterprise Server is in production under Oracle Solaris Cluster control, do not start or stop the servers manually or through the Server Manager Management Console. If you do need to perform maintenance on the servers, you must first disable their associated resources.

Tuning the HA for JD Edwards EnterpriseOne Enterprise Server Fault Monitors

This section provides the following information about the HA for JD Edwards EnterpriseOne Enterprise Server fault monitor's probing algorithm or functionality, and the conditions, messages, and recovery actions associated with unsuccessful probing:

- [“Resource Properties” on page 54](#)
- [“Probing Algorithm and Functionality” on page 54](#)
- [“Operation of the HA for JD Edwards EnterpriseOne Enterprise Server Probe” on page 55](#)

For conceptual information about fault monitors, see the [Oracle Solaris Cluster 4.3 Concepts Guide](#).

Resource Properties

The HA for JD Edwards EnterpriseOne Enterprise Server fault monitor uses the resource properties that are specified in the resource type `ORCL.JDE_Enterprise_Server`. See the [r_properties\(5\)](#) man page for a list of general resource properties used. See [Appendix A, “HA for JD Edwards EnterpriseOne Enterprise Server Extension Properties”](#) for a list of resource properties for this resource type.

Probing Algorithm and Functionality

The JD Edwards EnterpriseOne Enterprise Server is controlled by extension properties that control the probing frequency. The default values of these properties determine the preset behavior of the fault monitor and are suitable for most Oracle Solaris Cluster installations.

You can modify this preset behavior by modifying the following settings:

- The interval between fault monitor probes (`Thorough_probe_interval`).
- The timeout for fault monitor probes (`Probe_timeout`).
- The number of times the fault monitor attempts to restart the resource (`Retry_count`). The HA for JD Edwards EnterpriseOne Enterprise Server fault monitor checks the server status within an infinite loop. During each cycle, the fault monitor checks the server state and reports failure or success.

- If the fault monitor is successful, it returns to its infinite loop and continues the next cycle of probing and sleeping.
- If the fault monitor reports a failure, a request is made to the cluster to restart the resource. If the fault monitor reports another failure, another request is made to the cluster to restart the resource. This behavior continues whenever the fault monitor reports a failure. If successive restarts exceed the `Retry_count` within the `Thorough_probe_interval`, a request is made to fail over the resource group onto a different node.

Operation of the HA for JD Edwards EnterpriseOne Enterprise Server Probe

The following list explains how the HA for JD Edwards EnterpriseOne Enterprise Server probe operates:

- If the `control_app_server` script for the resource is still running with the start option, the probe returns `100`. This implements *wait for online* during start. Otherwise, the probe continues.
- If the output from `jdejobs` returns no running processes, the probe returns `100` to indicate a failed start. Otherwise, the probe continues.
- If the return code from `porttest` returns a non-zero value, the status of the resource is changed to `Degraded` with the status message *porttest failed* and the probe returns `0`.
- On a subsequent probe, if the `porttest` returns a value `0`, the probe returns `0` and the status messages of the resource are changed to *Online* and *Service is Online* respectively.
- If the HA for JD Edwards EnterpriseOne Enterprise Server resource is repeatedly restarted and subsequently exhausts the `Retry_count` within the `Retry_interval`, and if `Failover_enabled` is set to `TRUE`, a failover to another node is initiated for the resource group.

Debugging HA for JD Edwards EnterpriseOne Enterprise Server

HA for JD Edwards EnterpriseOne Enterprise Server has an extension property named `Debug_level`. This extension property enables you to activate and deactivate debugging for HA for JD Edwards EnterpriseOne Enterprise Server resources.

This section contains the following procedures:

- [“How to Activate Debugging for HA for JD Edwards EnterpriseOne Enterprise Server” on page 56](#)
- [“How to Deactivate Debugging for HA for JD Edwards EnterpriseOne Enterprise Server” on page 56](#)

▼ How to Activate Debugging for HA for JD Edwards EnterpriseOne Enterprise Server

1. **Edit the `/etc/syslog.conf` file on the appropriate node to change `daemon.notice` to `daemon.debug`.**
2. **Add `/var/cluster/logs/DS/ORCL.JDE_Enterprise_Server/message_log_$RESOURCE` to the `/etc/syslog.conf` file.**
3. **Confirm that debugging for HA for JD Edwards EnterpriseOne Enterprise Server is active.**

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit /var/cluster/logs/DS
/ORCL.JDE_Enterprise_Server/message_log_$RESOURCE
*.alert;kern.err;daemon.err operator
```

4. **Restart the `syslogd` daemon.**
5. **Set the property `Debug_level` to level 2.**

```
# svcadm refresh svc:/system/system-log:default
```

```
# clresource set -p Debug_level=2 jde-ent-svr-resource
```

▼ How to Deactivate Debugging for HA for JD Edwards EnterpriseOne Enterprise Server

1. **Edit the `/etc/syslog.conf` file on the appropriate node to change `daemon.debug` to `daemon.notice`.**
2. **Determine whether debugging for HA for JD Edwards EnterpriseOne Enterprise Server is not active.**

```
# grep daemon /etc/syslog.conf
```



```
*.err;kern.debug;daemon.notice;mail.crit /var/cluster/logs/DS
/ORCL.JDE_Enterprise_Server/message_log_$RESOURCE
*.alert;kern.err;daemon.err operator
```

3. Restart the syslogd daemon.

```
# svcadm refresh svc:/system/system-log:default
```

4. Set the property Debug_level to level 0.

```
# clresource set -p Debug_level=0 jde-ent-svr-resource
```


HA for JD Edwards EnterpriseOne Enterprise Server Extension Properties

Extension properties for the HA for JD Edwards EnterpriseOne Enterprise Server resource type are described below.

For details about system-defined properties, see the [r_properties\(5\)](#) man page and the [rg_properties\(5\)](#) man page.

ORCL.JDE_Enterprise_Server Extension Properties

The `ORCL.JDE_Enterprise_Server` resource type represents the JD Edwards EnterpriseOne Enterprise Server in an Oracle Solaris Cluster configuration. The extension properties of this resource type are as follows:

`Monitor_retry_interval`

Indicates the time in minutes, over which the failures of the fault monitor are counted, and corresponds to the `-t` option passed to the `pmfadm` command. If the number of times the fault monitor fails exceeds the value of `Monitor_retry_count`, the fault monitor is not restarted by the process monitor facility.

Data Type: Integer

Default: 2

Tunable: Any time

`Probe_timeout`

This property is the timeout value (in seconds) used by the fault monitor to probe a JD Edwards EnterpriseOne Enterprise Server instance.

Data Type: Integer

Default: 30

Tunable: Any time

Child_mon_level

This property provides control over the processes that are monitored through the Process Monitor Facility (PMF). This property denotes the level to which the forked children processes are monitored. Omitting this property or setting this property to the default value is the same as omitting the -C option for `rpc.pmfd(1M)`: all children (and their descendents) are monitored.

Data Type: Integer

Default: -1

Tunable: At creation

Validate_command

This property validates the application.

Data Type: String

Default: ""

Tunable: When disabled

Network_aware

This property determines whether the application uses a network.

Data Type: Boolean

Default: TRUE

Tunable: At creation

Stop_signal

This property controls the stop signal sent to the application.

Data Type: Integer

Default: 15

Tunable: When disabled

Failover_enabled

This property determines whether to fail over when the `Retry_count` is exceeded during the `Retry_count`.

Data Type: Boolean

Default: TRUE

Tunable: When disabled

Log_level

This property determines the log level for event-based traces.

Data Type: Enum

Default: INFO

Tunable: Anytime

Debug_level

This property controls the debug level for the control script and its functions.

Data Type: Integer

Default: 0

Tunable: Anytime

JDE_USER

This property determines the username used to run the application.

Data Type: String

Default: None

Tunable: When disabled

JDE_HOME

This property defines the installation directory for the JD Edwards EnterpriseOne Enterprise Server binaries.

Data Type: String

Default: None

Tunable: When disabled

Index

C

- Child_mon_level extension property
 - ORCL.JDE_Enterprise_Server resource type, 60
- configuration modes, 21
 - failover, 22
 - multi-instance, 26
 - multi-master, 28
- configuration requirements
 - HA for JD Edwards EnterpriseOne Enterprise Server, 19
- configuration restrictions
 - JD Edwards EnterpriseOne Enterprise Server, 18
- configuring
 - JD Edwards EnterpriseOne Enterprise Server, 35, 42
- configuring as a multi-instance configuration
 - HA for JD Edwards EnterpriseOne Enterprise Server, 48
- configuring as a multi-master data service
 - HA for JD Edwards EnterpriseOne Enterprise Server, 45

D

- data services
 - JD Edwards EnterpriseOne system components, 16
- Debug_level extension property
 - ORCL.JDE_Enterprise_Server resource type, 61

E

- extension properties
 - ORCL.JDE_Enterprise_Server resource type, 59

F

- failover configuration, 22
- Failover_enabled extension property
 - ORCL.JDE_Enterprise_Server resource type, 60
- fault monitors
 - tuning, 54

H

- HA for JD Edwards EnterpriseOne Enterprise Server
 - configurations, 21
 - fault monitors, 54
 - installing, 40
 - overview, 16, 17
 - registering and configuring as a multi-instance data service, 48
 - registering and configuring as a multi-master data service, 45
 - software package, installing, 40
 - verifying installation, 52

I

- installing
 - HA for JD Edwards EnterpriseOne Enterprise Server, 40
 - JD Edwards EnterpriseOne Enterprise Server, 35

J

- JD Edwards EnterpriseOne Enterprise Server
 - configuring, 35
 - installing, 35

- planning the installation, 18
- registering and configuring, 42
- verifying installation, 37

JD Edwards EnterpriseOne Enterprise Server software configuration restrictions, 18

JD Edwards EnterpriseOne system components data services, 16

JDE_HOME extension property
ORCL.JDE_Enterprise_Server resource type, 61

JDE_USER extension property
ORCL.JDE_Enterprise_Server resource type, 61

L

Log_level extension property
ORCL.JDE_Enterprise_Server resource type, 61

M

Monitor_retry_interval extension property
ORCL.JDE_Enterprise_Server resource type, 59

multi-instance configuration, 26

multi-master configuration, 28

N

Network_aware extension property
ORCL.JDE_Enterprise_Server resource type, 60

O

Oracle Solaris Cluster software publisher, 40, 40

ORCL.JDE_Enterprise_Server resource type extension properties, 59

overview
HA for JD Edwards EnterpriseOne Enterprise Server, 16, 17

P

package, 40

planning the installation
JD Edwards EnterpriseOne Enterprise Server, 18

probe
HA for JD Edwards EnterpriseOne Enterprise Server, 55

Probe_timeout extension property
ORCL.JDE_Enterprise_Server resource type, 59

publisher
Oracle Solaris Cluster software, 40, 40

R

registering
JD Edwards EnterpriseOne Enterprise Server, 42

registering as a multi-instance configuration
HA for JD Edwards EnterpriseOne Enterprise Server, 48

registering as a multi-master data service
HA for JD Edwards EnterpriseOne Enterprise Server, 45

resource types
ORCL.JDE_Enterprise_Server extension properties, 59

restrictions
JD Edwards EnterpriseOne Enterprise Server software, 18

S

software package, 40

Stop_signal extension property
ORCL.JDE_Enterprise_Server resource type, 60

T

tuning
fault monitors, 54

V

Validate_command extension property

ORCL.JDE_Enterprise_Server resource type, 60
verifying
 HA for JD Edwards EnterpriseOne Enterprise
 Server, 52
 JD Edwards EnterpriseOne Enterprise Server
 installation, 37

