

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

SPARC Platform Edition

ORACLE®

Part No: E40215
July 2014, E40215-01

Copyright © 2000, 2014, 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 END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

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

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

Copyright © 2000, 2014, 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 disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

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

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

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

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

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

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

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

Contents

- Using This Documentation** 7

- 1 Getting Started with JD Edwards EnterpriseOne Enterprise Server** 9
 - Overview 9
 - Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package 10
 - ▼ How to Install the HA for JD Edwards EnterpriseOne Enterprise Server Package 11

- 2 Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server** 13
 - Overview 13
 - Planning the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration 14
 - Configuration Restrictions for JD Edwards EnterpriseOne Enterprise Server 14
 - Configuration Requirements for JD Edwards EnterpriseOne Enterprise Server 15
 - JD Edwards EnterpriseOne Enterprise Server Data Service Configurations 17
 - Preparing the Nodes and Disks 25
 - ▼ How to Prepare the Nodes for a Failover Configuration 25
 - ▼ How to Prepare the Nodes for a Multi-Instance Configuration 26
 - ▼ How to Prepare the Nodes for a Multi-Master Configuration 27
 - Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server 28
 - ▼ How to Install JD Edwards Enterprise Server Software 28
 - Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration 29
 - ▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Failover Services 30
 - ▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Multi-Master and Multi-Instance Services 31
 - Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package 32

▼ How to Install the HA for JD Edwards EnterpriseOne Enterprise Server Package	32
Registering and Configuring the JD Edwards EnterpriseOne Enterprise Server	33
▼ How to Register and Configure JD Edwards EnterpriseOne Enterprise Server for Failover	33
▼ How to Register and Configure JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Data Service	35
▼ How to Register and Configure JD Edwards EnterpriseOne Enterprise Server as a Multi-Instance Data Service	36
▼ How to Remove a JD Edwards EnterpriseOne Enterprise Server Resource From a Failover Resource Group	37
Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration	37
▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Failover Configuration	38
▼ How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Multi-Master or Multi-Instance Configuration	38
Tuning the JD Edwards EnterpriseOne Enterprise Server Fault Monitors	39
Resource Properties	40
Probing Algorithm and Functionality	40
Operation of the JD Edwards EnterpriseOne Enterprise Server Probe	40
Debugging the JD Edwards EnterpriseOne Enterprise Server	41
▼ How to Activate Debugging for JD Edwards EnterpriseOne Enterprise Server	41
▼ How to Deactivate Debugging for JD Edwards Enterprise Server	42
A Oracle HA for JD Edwards EnterpriseOne Enterprise Server Extension Properties	43
ORCL.JDE_Enterprise_Server Extension Properties	43
Index	47

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

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E39579>.

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.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

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

Getting Started with JD Edwards EnterpriseOne Enterprise Server

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

This chapter contains the following sections.

- [“Overview” on page 9](#)
- [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package” on page 10](#)

Overview

The 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 provides fault monitoring and automatic failover for the Enterprise Server to eliminate single points of failure in a JD Edwards EnterpriseOne system. [Table 1-1](#) 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 1-1 JD Edwards EnterpriseOne System Components and Their Data Services

JD Edwards EnterpriseOne System Component	Data Service
Database server	The data service for the database that you are using:

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

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).

For conceptual information about failover and multi-master services, see the [“Oracle Solaris Cluster 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 Software Installation Guide”](#) for details on initial installation of clusters and data service software. You register the 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.

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 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 11.2”](#).

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 11.2”](#).

3. **Install the HA for JD Edwards EnterpriseOne 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 System Administration Guide”](#).

Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server

This chapter provides an overview of the JD Edwards EnterpriseOne Enterprise Server and explains how to install and configure it.

This chapter contains the following sections:

- [“Overview” on page 13](#)
- [“Planning the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 14](#)
- [“Preparing the Nodes and Disks” on page 25](#)
- [“Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server” on page 28](#)
- [“Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 29](#)
- [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package” on page 10](#)
- [“Registering and Configuring the JD Edwards EnterpriseOne Enterprise Server” on page 33](#)
- [“Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 37](#)
- [“Tuning the JD Edwards EnterpriseOne Enterprise Server Fault Monitors” on page 39](#)
- [“Debugging the JD Edwards EnterpriseOne Enterprise Server” on page 41](#)

Overview

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

[Table 2-1](#) 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 2-1 Tasks for Installing and Configuring the HA for JD Edwards EnterpriseOne Enterprise Server

Task	Instructions
1. Plan the installation	“Planning the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 14
2. Install and configure the JD Edwards EnterpriseOne Enterprise Server	“Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server” on page 28
3. Verify the JD Edwards Enterprise One Enterprise Server installation and configuration	“Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 29
4. Register and configure JD Edwards EnterpriseOne Enterprise Server resources	“Registering and Configuring the JD Edwards EnterpriseOne Enterprise Server” on page 33
5. Verify the JD Edwards Enterprise One Enterprise Server installation and configuration	“Verifying the HA for JD Edwards EnterpriseOne Enterprise Server Installation and Configuration” on page 37
6. Tune the JD Edwards EnterpriseOne Enterprise Server fault monitor	“Tuning the JD Edwards EnterpriseOne Enterprise Server Fault Monitors” on page 39
7. Debug the JD Edwards EnterpriseOne Enterprise Server	“Debugging the JD Edwards EnterpriseOne Enterprise Server” on page 41

Planning the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration

This section contains the information you need to plan your JD Edwards EnterpriseOne Enterprise Server installation and configuration.

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

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

- Configure the database resource to be in a different resource group from the JD Edwards EnterpriseOne Enterprise Server resource, if possible.

However, if you are using Oracle as the database, and if you must configure Oracle 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 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.

- Each JD Edwards EnterpriseOne Enterprise Server instance should be installed as a different user. If the installation is on a local file system, you are allowed to use the same user ID.



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

For restrictions that apply to all data services, see the “[Oracle Solaris Cluster 4.2 Release Notes](#)”.

Configuration Requirements for JD Edwards EnterpriseOne Enterprise Server

Use 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 version 9.1.x is published at [JD Edwards EnterpriseOne Tools Documentation Library \(http://docs.oracle.com/cd/E17984_01/index.htm\)](http://docs.oracle.com/cd/E17984_01/index.htm).

For requirements that apply to all data services, see [Chapter 1, “Planning for Oracle Solaris Cluster Data Services,”](#) in “[Oracle Solaris Cluster 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 (for example, jde910) on that node. 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 jde910 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.

- **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 JD Edwards Enterprise Server Software” on page 28](#).
- **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:
 1. 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:

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

The *LH-rs-hostname* is the hostname of the per-node logical host resource.

2. The `listenOnSpecificIP` setting in the `JDE.INI` file should be set to 1:

`listenOnSpecificIP=1`

3. The `SecurityServer` attribute under the `[SECURITY]` section should 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:

`SecurityServer=192.1.2.36`

The IP address of the per-node logical host resource hostname configured for the JD Edwards EnterpriseOne Enterprise Server is 192.1.2.36.

4. The `PrimaryNode` attribute under the `[CLUSTER]` section should 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:

`PrimaryNode=192.1.2.36`

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.

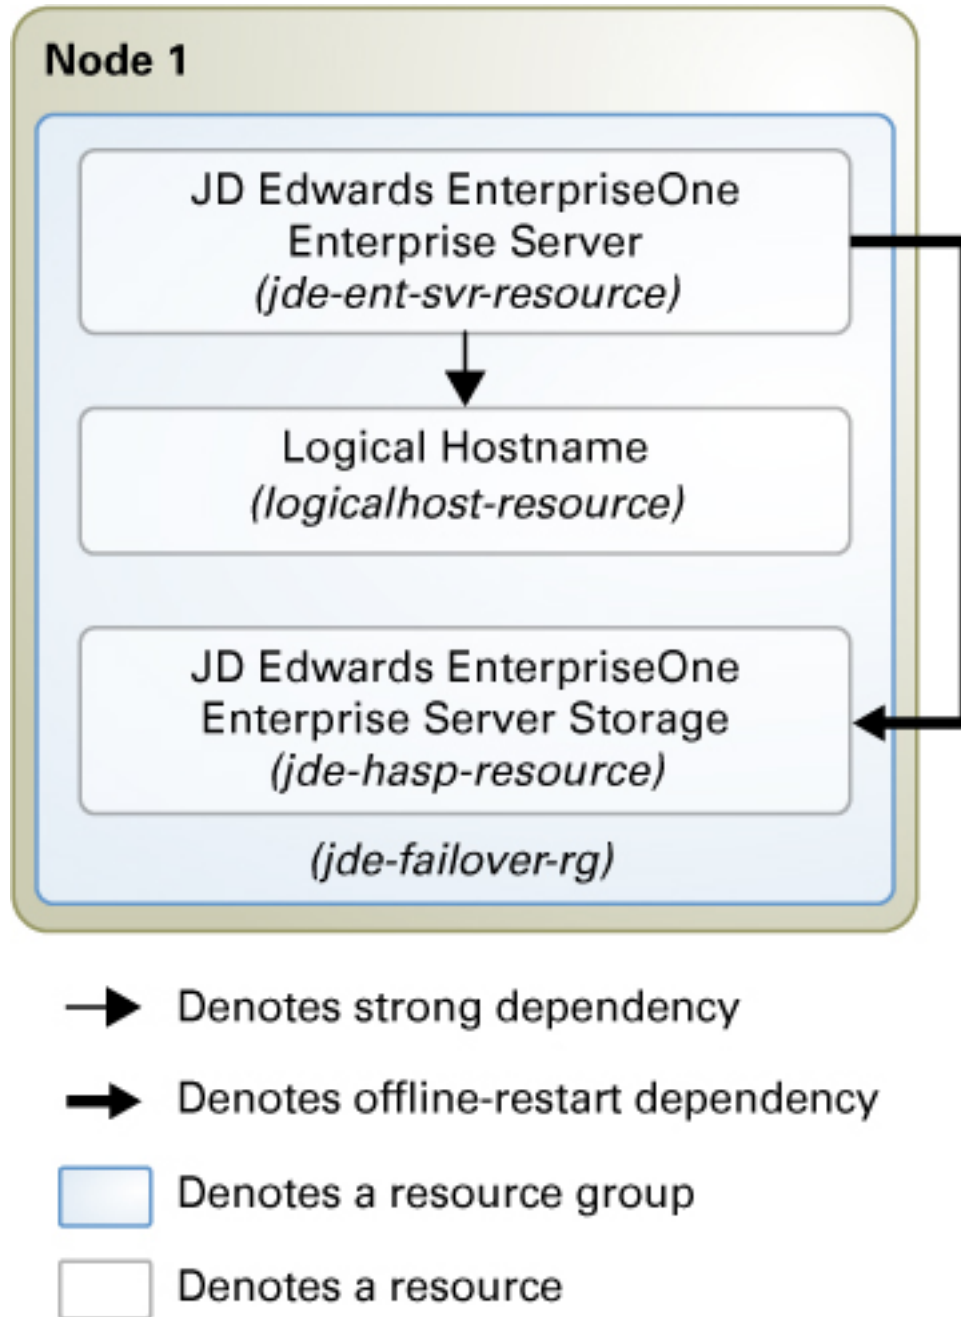
JD Edwards EnterpriseOne Enterprise Server Data Service Configurations

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

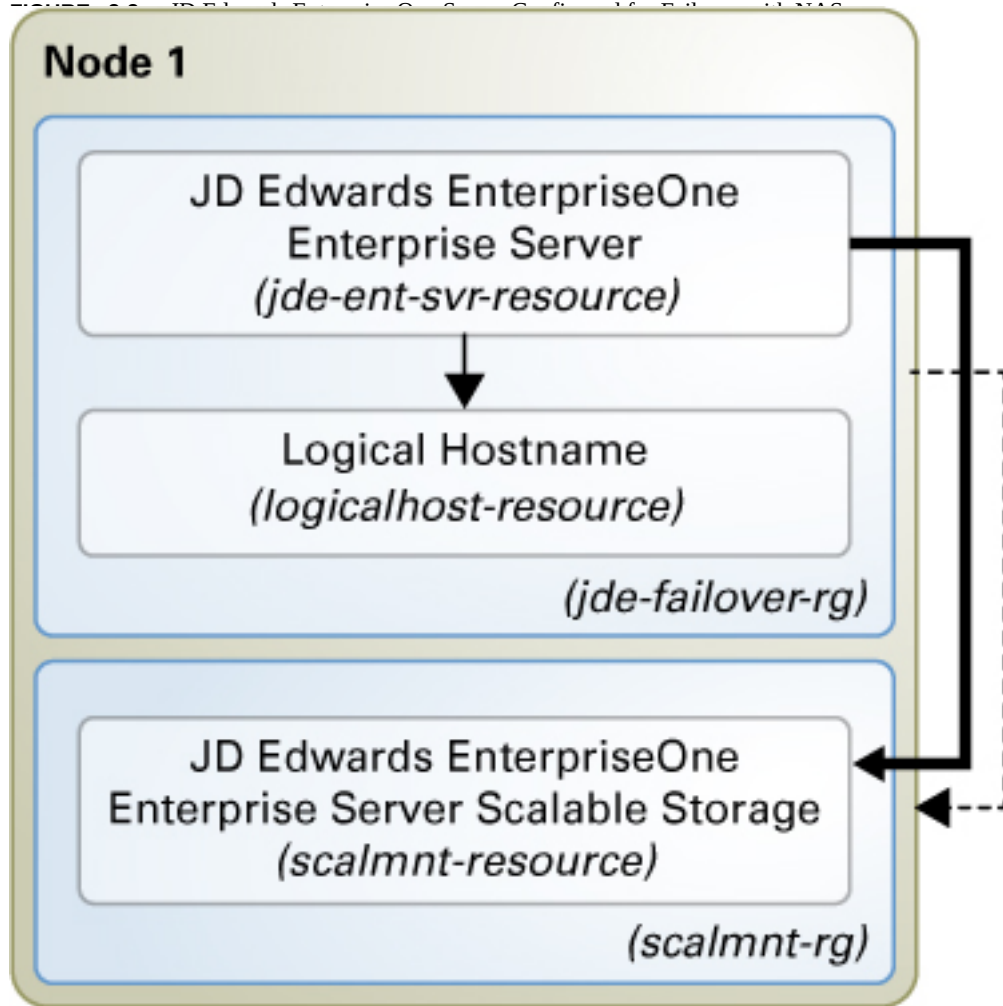
Failover Configuration

When the JD Edwards software is installed in a traditional file system, a failover deployment requires a configuration where one failover resource group contains the JD Edwards EnterpriseOne Enterprise Server resource, the logical hostname resource, and the failover storage resource. This type of configuration is shown in [Figure 2-1](#).

FIGURE 2-1 JD Edwards EnterpriseOne Enterprise Server Configured for Failover with Traditional File Storage



When the JD Edwards software is installed on NAS, a failover resource group is configured with the logical hostname resource and 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 2-2](#).



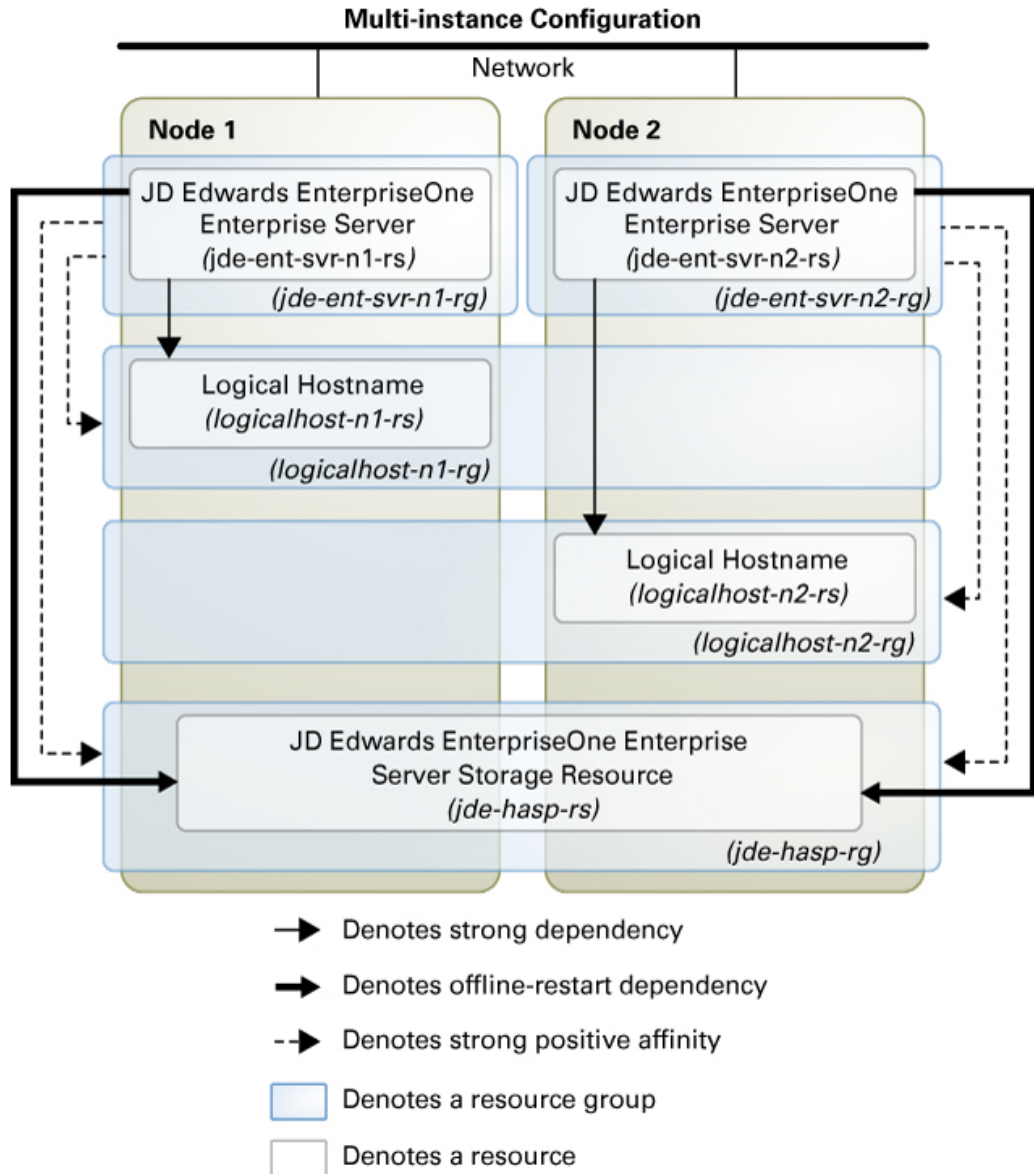
- Denotes strong dependency
- ➔ Denotes offline-restart dependency
- ➔ Denotes strong positive affinity
- Denotes a resource group
- Denotes a resource

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 2-3](#) illustrates a multi-instance configuration using single-node resource groups. A single-node resource group is created for each of the JD Edwards EnterpriseOne Enterprise Server resources. Each resource group has a strong positive affinity on a storage resource group whose primary node is the node containing the JD Edwards EnterpriseOne Enterprise Server resource group.

FIGURE 2-3 JD Edwards Enterprise One Server Configured as a Multi-Instance Application With Single-Node Resource Groups

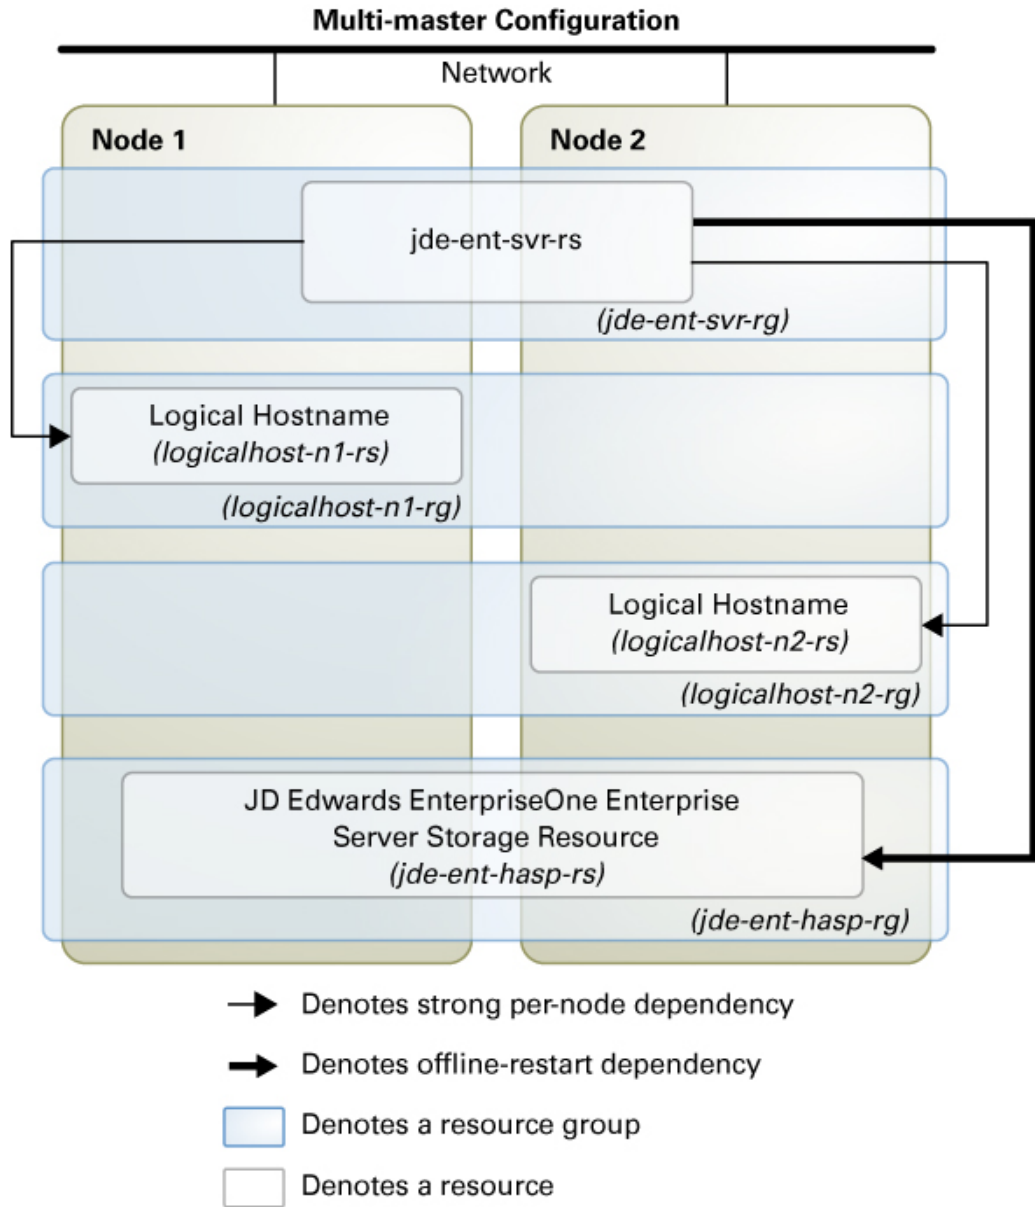


Multi-Master Configuration

The multiple-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 2-4](#) 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 is the node containing the JD Edwards EnterpriseOne Enterprise Server resource group.

FIGURE 2-4 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 25](#)
- [“How to Prepare the Nodes for a Multi-Instance Configuration” on page 26](#)
- [“How to Prepare the Nodes for a Multi-Master Configuration” on page 27](#)

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 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 JD Edwards EnterpriseOne Enterprise Server resources.**

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

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

The `logicalhost-resource` is the host name where JD Edwards EnterpriseOne Enterprise Server will be listening.

```
# clreslogicalhostname create -g jde-ent-failover-rg 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-ent-failover-rg -t SUNW.HAStoragePlus \
-p FilesystemMountPoints=jde-mount-points jde-ent-hasp-rs
```

For more information, see [Chapter 5, “Creating a Cluster File System,”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

4. **Bring the resource group online.**

```
# clresourcegroup online -M jde-ent-failover-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 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 -emM logicalhost-n1-rg
# clresourcegroup online -emM 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. **Create a cluster file system managed by the HASStoragePlus resource for each JD Edwards EnterpriseOne Enterprise Server instance.**

Create storage resource groups for each JD Edwards EnterpriseOne Enterprise Server instance.

```
# clresource create -g jde-ent-hasp-rg-1 -l -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=jde-mount-points-1 jde-ent-hasp-rs-1
# clresource create -g jde-ent-hasp-rg-2 -l -t SUNW.HASStoragePlus \
-p FilesystemMountPoints=jde-mount-points-1 jde-ent-hasp-rs-2
```

For more information, see [Chapter 5, “Creating a Cluster File System,”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

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

Use this procedure to prepare for the installation and configuration of 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 2-4](#).

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 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 -emM logicalhost-n1-rg
# clresourcegroup online -emM 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. Create a cluster file system managed by the HAStoragePlus resource for each JD Edwards EnterpriseOne Enterprise Server instance.**

Create storage resource groups for each JD Edwards EnterpriseOne Enterprise Server instance.

```
# clresource create -g jde-ent-hasp-rg-1 -l -t SUNW.HAStoragePlus \
-p FilesystemMountPoints=jde-mount-points-1 jde-ent-hasp-rs-1
# clresource create -g jde-ent-hasp-rg-2 -l -t SUNW.HAStoragePlus \
-p FilesystemMountPoints=jde-mount-points-1 jde-ent-hasp-rs-2
```

For more information, see [Chapter 5, “Creating a Cluster File System,”](#) in [“Oracle Solaris Cluster Software Installation Guide”](#).

Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server

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

▼ How to Install JD Edwards Enterprise Server Software

1. **On the cluster member where the *jde-ent-rg* resource group is online, assume the root role.**
2. **Follow the instructions in the appropriate JD Edwards 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 JD Edwards EnterpriseOne Enterprise Server](#)” on page 15.

3. **When you provide details about JD Edwards EnterpriseOne Enterprise Server on Enterprise Server Revisions of the Installation Plan using the Solution Explorer, enter 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`, enter `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 should contain the hostname of the logical host resource 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#g8d6ab57f7cedaac_ef90c_10a77c8e3f7__6057).**
5. **After the installation is complete, you must 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 [libschohost.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:

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

Symlinks are created for the `libschohost.so.1` to `/usr/lib/secure/`.

```
# ln -s /usr/cluster/lib/libschohost.so.1 /usr/lib/secure/libschohost.so.1
```

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

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

```
LD_PRELOAD_32=/usr/lib/secure/libschohost.so.1
LD_PRELOAD_64=/usr/lib/secure/64/libschohost.so.1
SC_LHOSTNAME=pole-2
export LD_PRELOAD_32 LD_PRELOAD_64 SC_LHOSTNAME
```

Verifying the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration

This section contains the procedure 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 30](#)
- [“How to Verify the JD Edwards EnterpriseOne Enterprise Server Installation and Configuration for Multi-Master and Multi-Instance Services” on page 31](#)

▼ 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-ent-failover-rg* is online on *node1*.**
2. **If you are using a database, verify that the database is running.**
3. **Go to the directory that contains the *RunOneWorld.sh* script that you configured in [“Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server”](#) on page 28.**
4. **Assume the *root* role and log into the node that currently hosts the *jde-ent-failover-rg* resource group and start the JD Edwards EnterpriseOne Enterprise Server by running the *RunOneWorld.sh* script.**

```
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/RunOneWorld.sh"
```

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

```
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/jdejobs"
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/porttest JDE JDE DV910"
```

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

```
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/EndOneWorld.sh"
```

7. **Switch the JD Edwards resource group to another cluster member.**

```
# clresourcegroup switch -n node jde-ent-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 the HA for Enterprise Server.

9. Stop the JD Edwards EnterpriseOne Enterprise Server and proceed to [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package”](#) on page 10.

▼ 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*.
4. Assume the *root* role and log into the node that currently hosts the *jde-ent-rg* resource group and start the *RunOneWorld.sh* script.

```
# su - jde_User -c "/u01/jdedwardspack/e910/system/bin32/RunOneWorld.sh"
```

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

```
# su - jde_User -c "/u01/jdedwardspack/e910/system/bin32/jdejobs"
# su - jde_User -c "/u01/jdedwardspack/e910/system/bin32/porttest JDE JDE DV910"
```

6. Start the server by going to the directory containing the JD Edwards EnterpriseOne Enterprise Server installation on *node2*.

```
# su - jde_User_2 -c "/u02/jdedwardspack/e910/system/bin32/RunOneWorld.sh"
```

7. Verify the status of the JD Edwards EnterpriseOne Enterprise Server.

```
# su - jde_User -c "/u01/jdedwardspack/e910/system/bin32/jdejobs"
# su - jde_User -c "/u01/jdedwardspack/e910/system/bin32/porttest JDE JDE DV910"
```

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 configuration.

8. **Stop the JD Edwards EnterpriseOne Enterprise Server and proceed to [“Installing the HA for JD Edwards EnterpriseOne Enterprise Server Package” on page 10.](#)**

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 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 11.2”](#).

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 11.2”](#).

3. Install the HA for JD Edwards EnterpriseOne 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 System Administration Guide”](#).

Registering and Configuring the JD Edwards EnterpriseOne Enterprise Server

This section contains the procedures to configure or unconfigure the JD Edwards EnterpriseOne Enterprise Server.

- [“How to Register and Configure JD Edwards EnterpriseOne Enterprise Server for Failover”](#) on page 33
- [“How to Register and Configure JD Edwards EnterpriseOne Enterprise Server as a Multi-Master Data Service”](#) on page 35
- [“How to Register and Configure JD Edwards EnterpriseOne Enterprise Server as a Multi-Instance Data Service”](#) on page 36
- [“How to Remove a JD Edwards EnterpriseOne Enterprise Server Resource From a Failover Resource Group”](#) on page 37

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

Use this procedure to configure Oracle Solaris Cluster 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 the JD Edwards EnterpriseOne Enterprise Server” on page 28.

Before You Begin Install the data service package during your initial Oracle Solaris Cluster installation. If you did not install the JD Edwards EnterpriseOne Enterprise Server package as part of your initial Oracle Solaris installation, go to [Broken Link \(Target ID: GCYUY-1\)](#).

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 a JD Edwards EnterpriseOne Enterprise Server resource in the failover resource group.**

```
# clresource create -g jde-ent-failover-rg -d \  
-t ORCL.JDE_Enterprise_Server \  
-p JDE_USER=jde-username \  
-p JDE_HOME=jde-home-directory \  
-p resource_dependencies=logicalhost-resource \  
-p resource_dependencies_offline_restart=jde-ent-hasp-rs jde-ent-svr-rs
```

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.**

Perform this step if the database tier is deployed in a different zone cluster of the same global cluster.

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

5. **Enable the 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-rs
```

▼ How to Register and Configure 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 23](#) and have configured the server as described in [“Installing and Configuring the JD Edwards EnterpriseOne Enterprise Server” on page 28](#). If you are planning a different configuration, you must configure the JD Edwards EnterpriseOne Enterprise Server resources accordingly.

1. Add the JD Edwards EnterpriseOne Enterprise Server resource type.

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

2. Create a scalable resource group for the JD Edwards EnterpriseOne 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 -M 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*.

```
# clresource create -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-rg \
-p JDE_USER{node1}=jde-username \
-p JDE_USER{node2}=jde-username \
-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-rs \  
-p resource_dependencies_offline_restart=jde-ent-hasp-rs-1,jde-ent-hasp-rs-2 \  
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.

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

Before You Begin Use this procedure to configure a 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 zone cluster. If you are performing this configuration in a global cluster, do not use the `-Z` option.

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

```
# clresourcegroup create -Z zone-name -p nodelist=node1 \  
-p RG_affinities=++jde-ent-hasp-rg1 jde-ent-svr-n1-rg
```

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

```
# clresourcegroup online -Z zone-name -M jde-ent-svr-n1-rg
```

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

```
# clresourcegroup create -Z zone-name -p nodelist=node2 \  
-p RG_affinities=++jde-ent-hasp-rg2 jde-ent-svr-n2-rg
```

4. **Bring the second resource group online.**

```
# clresourcegroup online -Z zone-name -M jde-ent-svr-n2-rg
```

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

```
# clresource create -Z zone-name -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-n1-rg \  

```

```
-p JDE_USER=username \  
-p JDE_HOME=/path-to-jde-installation-dir \  
-p resource_dependencies=logicalhost-n1-rs \  
-p resource_dependencies_offline_restart=jde-ent-hasp-rs1 \  
jde-ent-svr-n1-rs
```

The resource is created in the enabled state.

6. Create the second JD Edwards EnterpriseOne Enterprise Server resource and set its properties.

```
# clresource create -Z zone-name -t ORCL.JDE_Enterprise_Server -g jde-ent-svr-n2-rg \  
-p JDE_USER=username \  
-p JDE_HOME=/path-to-jde-installation-dir \  
-p resource_dependencies=logicalhost-n2-rs \  
-p resource_dependencies_offline_restart=jde-ent-hasp-rs2 \  
jde-ent-svr-n2-rs
```

The resource is created in the enabled state.

▼ How to Remove a JD Edwards EnterpriseOne Enterprise Server Resource From a Failover Resource Group

1. On the cluster node that hosts JD Edwards EnterpriseOne Enterprise Server resource, assume the root role that provides `solaris.cluster.modify` and `solaris.cluster.admin` RBAC authorizations.
2. Disable and remove the resource that is used by the JD Edwards EnterpriseOne Enterprise Server data service.

```
# clresourcegroup disable jde-ent-svr-rs  
# clresourcegroup delete jde-ent-svr-rs
```

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

This section contains the procedure to verify that you installed and configured the JD Edwards EnterpriseOne Enterprise Server resource correctly.

▼ How to Verify the 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 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-ent-failover-rg
```

3. **Switch the JD Edwards EnterpriseOne Enterprise Server resource group to another cluster member.**

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

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

```
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/jdejobs"  
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/porttest JDE JDE DV910"
```

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

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

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 JD Edwards EnterpriseOne Enterprise Server Installation and Configuration in a Multi-Master or Multi-Instance Configuration

This procedure verifies that you installed and configured Oracle Solaris Cluster HA for JD Edwards 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.**

```
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/jdejobs"  
# su - jde_User -c "/u01/jdedwardsppack/e910/system/bin32/portttest JDE JDE DV910"
```

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

Tip - After the JD Edwards EnterpriseOne Enterprise Server are 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 JD Edwards EnterpriseOne Enterprise Server Fault Monitors

This section describes the JD Edwards EnterpriseOne Enterprise Server fault monitor's probing algorithm or functionality, and states the conditions, messages, and recovery actions associated with unsuccessful probing:

- [“Resource Properties” on page 40](#)
- [“Probing Algorithm and Functionality” on page 40](#)

- [“Operation of the JD Edwards EnterpriseOne Enterprise Server Probe” on page 40](#)

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

Resource Properties

The 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 [“ORCL.JDE_Enterprise_Server Extension Properties” on page 43](#) 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 JD Edwards EnterpriseOne Enterprise Server Probe

The following list explains how the 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 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 the JD Edwards EnterpriseOne Enterprise Server

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

▼ How to Activate Debugging 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 the 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.**

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

5. **Set the property `debug_level` to level 2.**

```
# clresource set -p debug_level=2 jde-ent-svr-rs
```

▼ How to Deactivate Debugging for JD Edwards Enterprise Server

1. **Edit the `/etc/syslog.conf` file on the appropriate node to change `daemon.debug` to `daemon.notice`.**
2. **Determine if debugging for the 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
```



Oracle HA for JD Edwards EnterpriseOne Enterprise Server Extension Properties

Extension properties for the Oracle HA for JD Edwards EnterpriseOne Enterprise Server 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 `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, 44
- configuration modes, 17
 - failover, 17
 - multi-instance, 21
 - multi-master, 23
- configuration requirements
 - JD Edwards EnterpriseOne Enterprise Server, 15
- configuration restrictions
 - JD Edwards EnterpriseOne Enterprise Server, 14
- configuring
 - JD Edwards EnterpriseOne Enterprise Server, 28, 33
- configuring as a multi-instance configuration
 - HA for JD Edwards EnterpriseOne Enterprise Server, 36
- configuring as a multi-master data service
 - HA for JD Edwards EnterpriseOne Enterprise Server, 35

D

- data services
 - JD Edwards EnterpriseOne system components, 9
- Debug_Level extension property
 - ORCL.JDE_Enterprise_Server resource type, 45

E

- extension properties
 - ORCL.JDE_Enterprise_Server resource type, 43

F

- failover configuration, 17

- Failover_enabled extension property
 - ORCL.JDE_Enterprise_Server resource type, 44
- fault monitors
 - tuning, 39

H

- HA for JD Edwards
 - registering and configuring as a multi-instance data service
 - JD Edwards EnterpriseOne Enterprise Server, 36
 - registering and configuring as a multi-master data service
 - JD Edwards EnterpriseOne Enterprise Server, 35
- HA for JD Edwards EnterpriseOne Enterprise Server
 - installing, 10, 32
 - software package, installing, 10, 32

I

- installing
 - JD Edwards EnterpriseOne Enterprise Server, 10, 28, 32

J

- JD Edwards EnterpriseOne Enterprise Server
 - configurations, 17
 - configuring, 28
 - fault monitors, 39
 - installing, 28
 - overview, 9, 13
 - registering and configuring, 33

- verifying installation, 30, 38
- JD Edwards EnterpriseOne system components
 - data services, 9
- JDE_HOME extension property
 - ORCL.JDE_Enterprise_Server resource type, 45
- JDE_USER extension property
 - ORCL.JDE_Enterprise_Server resource type, 45
- L**
- Log_Level extension property
 - ORCL.JDE_Enterprise_Server resource type, 44
- M**
- Monitor_retry_interval extension property
 - ORCL.JDE_Enterprise_Server resource type, 43
- multi-instance configuration, 21
- multi-master configuration, 23
- N**
- Network_aware extension property
 - ORCL.JDE_Enterprise_Server resource type, 44
- O**
- Oracle Solaris Cluster software
 - publisher, 11, 11, 32, 32
- ORCL.JDE_Enterprise_Server resource type
 - extension properties, 43
- overview
 - JD Edwards EnterpriseOne Enterprise Server, 9, 13
- P**
- package, 10, 32
- planning the installation
 - JD Edwards EnterpriseOne Enterprise Server, 14
- probe
 - JD Edwards EnterpriseOne Enterprise Server, 40
- Probe_timeout extension property
 - ORCL.JDE_Enterprise_Server resource type, 43
- publisher
 - Oracle Solaris Cluster software, 11, 11, 32, 32
- R**
- registering
 - JD Edwards EnterpriseOne Enterprise Server, 33
- registering as a multi-instance configuration
 - HA for JD Edwards EnterpriseOne Enterprise Server, 36
- registering as a multi-master data service
 - HA for JD Edwards EnterpriseOne Enterprise Server, 35
- resource types
 - ORCL.JDE_Enterprise_Server
 - extension properties, 43
- S**
- software package, 10, 32
- Stop_signal extension property
 - ORCL.JDE_Enterprise_Server resource type, 44
- T**
- tuning
 - fault monitors, 39
- V**
- Validate_command extension property
 - ORCL.JDE_Enterprise_Server resource type, 44
- verifying
 - JD Edwards EnterpriseOne Enterprise Server, 38
 - JD Edwards EnterpriseOne Enterprise Server installation, 30