

JD Edwards EnterpriseOne High Availability Architecture

ORACLE TECHNICAL BRIEF | MAY 2017





Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



Table of Contents

Introduction	1
HAA Design for Oracle JD Edwards EnterpriseOne	2
Oracle RAC	3
Oracle Clusterware	4
Conventions	4
Installing and Configuring the Deployment Server	5
Installing and Configuring the Oracle Database on RAC	5
Installing and Configuring EnterpriseOne Using Clusterware 11g for an Active - Passive Cluster	5
Pre-configuration for Clusterware Install	6
Configure Disks for ASM	9
Install Clusterware	11
Installation of Oracle Clusterware Software	13
Configuring the Shared Disk Array with the ASMCA Utility	14
Installing the JD Edwards EnterpriseOne Platform Pack	15
Install the JDE Server Manager Agent	17
Installing the Clusterware Agent	18
Start the EnterpriseOne Application under Clusterware Control	19
Basic Commands to Control the Cluster and Resources	21
Setting Up a Batch Server	24
Setting Up the Web Server for Clustering	24
Complete JD Edwards EnterpriseOne High Availability Architecture System	24
Batch Behavior During Fail-over	25



High Availability (HA) Design for Oracle JD Edwards EnterpriseOne on the Oracle Public Cloud	32
Oracle RAC with Database Cloud Service (DBCS)	32
Oracle Traffic Director (OTD) and Oracle HTTP Server (OHS)	32
Virtual Hostnames Setup and Configuration	33
Aliasing for the Applications Server	33
Aliasing for the Batch Server	34
Modifying JD Edwards EnterpriseOne to Listen to the Logical Name	35
Disaster Recovery (DR) for Oracle JD Edwards EnterpriseOne on the Oracle Public Cloud	38
Active Data Guard	38
Setup and Configuration of RSYNC	38
HA and DR of Media Objects and Print Queue Directory	38
Setting up a Print Queue Directory on a DBCS Instance File Structure	39
Setting up the NFS Server on the DBCS Machine jdedbs92a	39
Setting up the NFS Client on Enterprise Server ent920	40
Appendix A: Script	42

Introduction

High Availability (HA) is the measurement of a system's ability to remain accessible in the event of a system component failure. Generally, HA is implemented by building in multiple levels of fault tolerance and/or load balancing capabilities into a system. On the other hand, disaster recovery is the process by which a system is restored to a previous acceptable state, after a natural or man-made disaster. While they both increase overall availability, a notable difference is that with HA there is, generally, no loss of service. HA refers to the retaining of the service and disaster recovery to the retaining of the data.

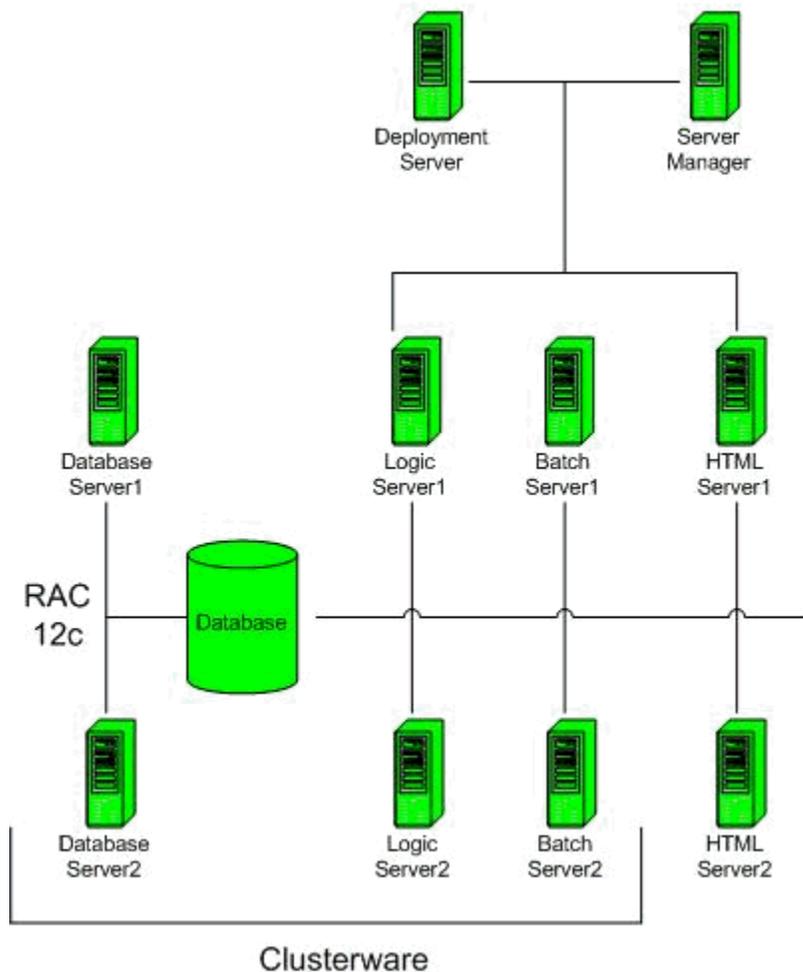
HIGH AVAILABILITY AND DISASTER RECOVERY COMPARISON

High Availability	Disaster Recovery
Addresses service availability, providing redundancy so that if one infrastructure component (network, servers, processes) becomes unavailable, overall service remains available.	Addresses service continuity, so that in case of disaster, service is maintained through a standby site.
A single system contains its own data (in the file system and database) and executables	Two independent environments, typically in separate and distinct facilities, each contain their own data (in the file system and database) and executables.
Data replication is unnecessary (although data should be backed up).	Data and configuration information are replicated between the production and standby sites.

High Availability Architecture (HAA) for Oracle JD Edwards EnterpriseOne is a combination of best practices and optimal architectures at the lowest cost and complexity. HAA is a state where the end user is unaware of hardware or software failures. Configurations within HAA for Oracle JD Edwards EnterpriseOne use Oracle Clusterware applied to both the database and middle-tier servers of the JD Edwards EnterpriseOne environment. Oracle Clusterware provides the ability for servers to communicate and operate as a collective unit. Servers with Oracle Clusterware installed and configured run additional processes that communicate with other Oracle Clusterware servers in order to appear as one system to both user and application. Oracle Clusterware provides the foundation for Oracle Real Application Clusters as well as management of defined resources for other applications. Defined resources can include the virtual IP addresses and application services.

To understand how this technology works with JD Edwards EnterpriseOne tests were performed to characterize how JD Edwards EnterpriseOne responds when a failure occurs. The tests evaluated the user experience during a Real Application Cluster (RAC) node, JD Edwards EnterpriseOne node, and Batch node failures. The tests also evaluated the user experience during a node failure while interactively working with a JD Edwards EnterpriseOne application as an end user. Behavior of submitted batch jobs (UBE) processed were also observed and documented.

Note: The applications in this testing represent a generic sample to test operations and failover characteristics of HAA with JD Edwards EnterpriseOne. Testing was implemented in a controlled environment with the generally available versions of JD Edwards EnterpriseOne.



HAA Environment for Oracle JD Edwards EnterpriseOne

HAA Design for Oracle JD Edwards EnterpriseOne

The HAA environment contained the same components as a standard JD Edwards EnterpriseOne environment except that each item is installed on top of another component that provides HA. The environment discussed in this document contained the following items:

- » JD Edwards EnterpriseOne Database was run on top of the Oracle RAC 12c
- » JD Edwards EnterpriseOne Server and Batch Server was run on top of Oracle Clusterware 11g
- » JD Edwards EnterpriseOne JAS Server was run on multiple WebLogic Server 11g servers with an OHS front end
- » Deployment and Server Manager instance had no special HAA configuration



In the HAA environment the database is installed on RAC with both nodes active. If one node fails the database stays up and running on the other node. The JD Edwards EnterpriseOne Server in the HAA environment is set up to run on one active node with another standby node ready to take over. The JD Edwards EnterpriseOne application is installed on a disk shared between the two nodes. If a failure occurs the standby node takes over the processing for JD Edwards EnterpriseOne. This is also the configuration of the Batch Servers. The JAS or HTML layer can be configured in either a vertical cluster or horizontal cluster with separate Oracle HTTP Server (OHS) routing transactions to either server.

- » A vertical cluster has multiple JD Edwards EnterpriseOne JAS instances on a single server. This allows for a controlled performance on a single server.
- » A horizontal cluster has multiple JD Edwards EnterpriseOne JAS instances on separate servers.

This configuration provides HA clustering. The Deployment Server and the Server Manager Console are not required for end user access of the JD Edwards EnterpriseOne application and are not configured on Clusterware. Continual backups and virtualizations of the Deployment Server and the Server Manager Console should be considered for recovery of these servers.

Oracle RAC

Oracle RAC extends the Oracle Database so that you can store, update, and efficiently retrieve data using multiple database instances on different servers at the same time. Oracle RAC provides the software that manages multiple servers and instances as a single group. The data files that comprise the database must reside on shared storage that is accessible from all servers that are part of the cluster. Each server in the cluster runs the Oracle RAC software.

Each database instances in an Oracle RAC database uses its own memory structures and background processes. Oracle RAC uses Cache Fusion to synchronize the data stored in the buffer cache of each database instance. Cache Fusion moves current data blocks (which reside in memory) between database instances, rather than having one database instance write the data blocks to disk and requiring another database instance to reread the data blocks from disk. When a data block located in the buffer cache of one instance is required by another instance, Cache Fusion transfers the data block directly between the instances using the interconnect, enabling the Oracle RAC database to access and modify data as if the data resided in a single buffer cache.

Oracle RAC is also a key component for implementing the Oracle enterprise grid computing architecture. Having multiple database instances accessing a single set of data files prevents the server from being a single point of failure. If a node in the cluster fails, the Oracle Database continues running on the remaining nodes. Individual nodes can be shut down for maintenance while application users continue to work.

Any packaged or custom application that scales on an Oracle Database scales well on Oracle RAC without requiring changes to the application code.

Single-instance Oracle databases have a one-to-one relationship between the Oracle database and the instance. Oracle RAC environments, however, have a one-to-many relationship between the database and instances. An Oracle RAC database can have up to 100 instances, all of which access one set of database files. All database instances must use the same interconnect, which can also be used by Oracle Clusterware.

Oracle RAC databases differ architecturally from single-instance Oracle databases in that each Oracle RAC database instance also has:

- » At least one additional thread of redo for each instance
- » An instance-specific undo tablespace



The combined processing power of the multiple servers can provide greater throughput and Oracle RAC scalability than is available from a single server.

For more information on the set up of Oracle RAC see:

- » JD Edwards EnterpriseOne Oracle 11gR2 Database Data Guard with RAC (Doc ID 1315368.1)
- » RAC Guide: 12.1.0.1 RAC Installation on Linux [Video] (Doc ID 1600316.1)

Oracle Clusterware

Oracle Clusterware is the integrated foundation for RAC and the HA and resource management framework for all applications on any platform. Oracle Clusterware 11g provides comprehensive multi-tiered HA and resource management for consolidated environments. The HAA configuration described in this document uses Oracle Clusterware to provide enterprise-class resiliency for the JD Edwards EnterpriseOne Server.

The Oracle Grid Infrastructure provides the necessary components to HA for the JD Edwards EnterpriseOne application. The JD Edwards EnterpriseOne application is integrated within the Oracle Grid Infrastructure with the Oracle Grid Infrastructure Agent. Oracle Grid Infrastructure Agents enable workload management and HA for Oracle applications. This agent is available for download from [Oracle Technology Network](http://www.oracle.com/technetwork/database/database-technologies/clusterware/downloads/index.html) (OTN) and can be configured in a user-defined agent home. This Agent is a standalone deployment model and offers greater flexibility in choosing the agent home directory structure and for subsequent agent upgrades.

Download for Clusterware JD Edwards EnterpriseOne Agent:

<http://www.oracle.com/technetwork/database/database-technologies/clusterware/downloads/index.html>

Customers install the JD Edwards EnterpriseOne application into a shared disk configured and managed by Clusterware Automatic Storage Management (ASM). Clusterware ASM provides disk configuration and management to the Clusterware component. In the HAA JD Edwards EnterpriseOne environment ASM provides the shared disk where the JD Edwards EnterpriseOne application resides. This disk is shared between the two nodes in the HAA environment. One of the nodes is the active node providing cpu for the execution of the JD Edwards EnterpriseOne application. The other node is the standby node ready to take over processing. Please refer to the Oracle Grid Infrastructure Agents guide on the download site for the Clusterware Agent for installation instruction.

Conventions

For the purpose of this document the following configuration is used:

Deployment Server

Tools Release = 9.1.4.6

Server Manager

densrvman.domain.com

Tools Release = 9.1.4.6

Database RAC Servers

Oracle RAC 12c

OS Release = OEL 6.5

Enterprise Servers

Oracle Clusterware 11g



OS Release = Oracle Enterprise Linux 5.9

Tools Release = 9.1.4.6

EnterpriseOne Batch Servers

Oracle Clusterware 11g

OS Release = Oracle Enterprise Linux 5.9

Tools Release = 9.1.4.6

HTML Servers

OS Release = OEL 6.5

WebLogic Server 12c (12.1.2.0)

Tools Release = 9.1.4.4

Note: Oracle RAC is used for the database servers and the Enterprise Server used Clusterware 11g because of the dependencies of ASM, Clusterware agent, and the Linux 5.9 operating system.

Note: We recommend the RAC Guide: 12.1.0.1 RAC Installation on Linux Document ID 1600316.1 when installing and configuring the RAC and Clusterware environments.

Installing and Configuring the Deployment Server

No special procedures were used to install a Deployment Server in the HAA environment. At this point in the document we assume that a Deployment Server has been installed and you are ready to run the platform packs for the Database and Enterprise Servers. Once the Database RAC setup is complete and the EnterpriseOne Server is setup in the Clusterware, run the installation workbench.

Note: The Deployment Server is not necessary for normal user activity and does not need a scaling or clustering solution for HAA.

Installing and Configuring the Oracle Database on RAC

The EnterpriseOne Database was setup using the following documents:

- » JD Edwards EnterpriseOne Oracle 11gR2 Database Data Guard with RAC (Doc ID 1315368.1)
- » RAC Guide: 12.1.0.1 RAC Installation on Linux [Video] (Doc ID 1600316.1)

Installing and Configuring EnterpriseOne Using Clusterware 11g for an Active - Passive Cluster

The following section describes how to setup a Linux system to act as a Clusterware system. At a minimum, two Linux nodes are needed. For this example we will use the following machine names:

- » nodelx01
- » nodelx02.

To setup a Linux system to act as a Clusterware system:

1. Pre-configure the system for the installation of Clusterware.
2. Install Clusterware
3. Configure ASM
4. Install the Clusterware agent updated to handle the Oracle JD Edwards EnterpriseOne Server.
5. Install the Oracle JD Edwards EnterpriseOne specific installation
 - a. Install the Server Manager Agent
 - b. Install EnterpriseOne Platform Pack
 - c. Install the Tools Release Update
 - d. Configure the Clusterware Agent
 - e. Start the EnterpriseOne application under Clusterware control.

Pre-configuration for Clusterware Install

The Enterprise and Batch Servers all require pre-configuration before the Clusterware installation. The systems in this document all are setup with Oracle Enterprise Linux 5.9. All RPM installs are available when a system is registered with the Oracle Enterprise Network. Use the yum command to upgrade the patch level of the Linux system. Unless otherwise specified all changes are made on both nodes of the cluster.

1. Add and setup users, groups and create installation directories.

Use the following commands:

```
/usr/sbin/groupadd -g 501 oinstall
/usr/sbin/groupadd -g 502 dba
/usr/sbin/groupadd -g 504 asmadmin
/usr/sbin/groupadd -g 506 asmdba
/usr/sbin/groupadd -g 507 asmoper
/usr/sbin/groupadd -g 508 jde910

/usr/sbin/useradd -u 501 -g oinstall -G
asmadmin,asmdba,asmoper grid
/usr/sbin/useradd -u 502 -g oinstall -G dba,asmdba oracle
/usr/sbin/useradd -u 503 -g jde910 -G oinstall,dba jde910

mkdir -p /u01/app/grid
mkdir -p /u01/app/<product version>/grid
mkdir -p /u01/app/oracle
mkdir -p /u01/app/oracle/product/<product version>/dbhome_1
mkdir -p /u01/oraInventory
chown -R grid:oinstall /u01/app/<product version>
chmod -R 777 /u01/app
```

Note: On the Batch and Enterprise servers the product level was 12.1.0.1 .

2. Make the following changes on both nodes of the cluster:
 - a. Edit the file /etc/sysconfig/network set NOZEROCONF=yes on both nodes

```
#cat /etc/sysconfig/network
```

```
NETWORKING=yes
NETWORKING_IPV6=no
HOSTNAME=node1x01.clustr.com
GATEWAY=101.39.194.1
NOZEROCONF=yes
```

- b. Update the settings in the `/etc/sysctl.conf` file with at least these values:

```
#Added for Oracle Grid and EnterpriseOne
kernel.shmmni=4096
kernel.sem=1024 32000 100 142
fs.aio-max-nr=1048576
fs.file-max=6815744
```

```
net.ipv4.ip_local_port_range=9000 65500
net.core.rmem_default=262144
net.core.rmem_max=4194304
net.core.wmem_default=262144
net.core.wmem_max=1048576
```

- c. Use the `sysctl` command to update the system:

```
sysctl -p /etc/sysctl.conf
```

- d. Make sure the following RPMs are loaded on the system. For the system registered with the Oracle Enterprise Linux use the `yum` command to update the RPMs.

```
binutils-2.20.51.0.2-5.36.el6 (x86_64)
compat-libstdc++-33-3.2.3-69.el6 (x86_64)
elfutils-libelf-0.152-1.el6 (x86_64)
elfutils-libelf-devel-0.152-1.el6 (x86_64)
gcc-4.4.7-4.el6 (x86_64)
gcc-c++-4.4.7-4.el6 (x86_64)
glibc-2.12-1.132.el6 (x86_64)
glibc-2.12-1.132.el6 (i686)
glibc-common-2.12-1.132.el6 (x86_64)
glibc-devel-2.12-1.132.el6 (x86_64)
glibc-headers-2.12-1.132.el6 (x86_64)
ksh-20120801-10.el6 (x86_64)
libaio-0.3.107-10.el6 (x86_64)
libaio-devel-0.3.107-10.el6 (x86_64)
libgcc-4.4.7-4.el6 (x86_64)
libgcc-4.4.7-4.el6 (i686)
libstdc++-4.4.7-4.el6 (x86_64)
libstdc++-devel-4.4.7-4.el6 (x86_64)
make-3.81-20.el6 (x86_64)
sysstat-9.0.4-22.el6 (x86_64)
unixODBC-2.2.14-12.el6_3 (x86_64)
unixODBC-devel-2.2.14-12.el6_3 (x86_64)
```

- e. Install the following RPMs:

```
oracleasm-2.6.18-8.el5-2.0.4-1.el5
oracleasm-lib-2.0.4-1.el5
oracleasm-support-2.1.8-1.el5
```

- f. Edit the PAM system files.

- i. Make sure the file `/etc/pam.d/login` contains the following line:

```
session include system-auth
```

- ii. Make sure the file `/etc/pam.d/system-auth` contains the following line:

- session required pam_limits.so
- g. Edit the file `/etc/security/limits.conf` file. Add the following lines to the file:

```
# grid user limits
grid soft nofile 1024
grid hard nofile 65536
grid soft nproc 16384
grid hard nproc 16384
grid soft stack 10240
grid hard stack 32768

# oracle user limits
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft nproc 16384
oracle hard nproc 16384
oracle soft stack 10240
oracle hard stack 32768
```

Note: Clusterware uses the NTP system. NTP stands for Network Time Protocol, and is an Internet protocol used to synchronize the clock of computers. NTP requires a time server for configurations.

- h. Enter these servers in the NTP server in the file `/etc/ntp.conf`:
- ```
ntp server <NTP SERVER>
ntp server <NTP SERVER> prefer
```
- i. Edit the file `/etc/sysconfig/ntpd` adding the `-x` parameter to the options line:
- ```
OPTIONS="-x -u ntp:ntp -p /var/run/ntpd.pid"
```
- j. Configure the NTP server to run on reboot.
- Use the command:
- ```
chkconfig ntpd on 35
```
- k. Check this with the command:
- ```
#chkconfig --list | grep ntpd
```
- ```
ntpd 0:off 1:off 2:on 3:on 4:on 5:on
6:off
```

---

*Note: Verify that 2, 3, and 5 have the "on" setting.*

---

3. Start the server using the command:
- ```
service ntpd start
```

Note: Oracle Enterprise Linux 5 comes with Avahi, IP 6 Tables and Secure Linux. This system can conflict with Clusterware.

4. To shutdown these programs do the following:
- a. To shutdown and prevent Avahi from restarting do the following:
- ```
service avahi-daemon stop
service avahi-dnssconfd stop
```

```
chkconfig avahi-daemon off
chkconfig avahi-dnssconfd off
```

- b. To shutdown and prevent IP tables from running do the following:

```
service iptables stop
service ip6tables stop
chkconfig iptables off
chkconfig ip6tables off
```

- c. To disable SELinux system edits the /etc/selinux/config file. Change the SELINUX line to "disabled"

```
cat /etc/selinux/config
This file controls the state of SELinux on the system.
SELINUX= can take one of these three values:
enforcing - SELinux security policy is enforced.
permissive - SELinux prints warnings instead of
enforcing.
disabled - SELinux is fully disabled.
SELINUX=disabled
```

- d. Reboot the system.

## Configure Disks for ASM

In this paper we used a fiber connected disk array configured with RAID1 into a set of four device files.

```
ls -l /dev/sd*
brw-r----- 1 root disk 8, 0 Aug 22 10:00 /dev/sda
brw-r----- 1 root disk 8, 16 Aug 22 10:00 /dev/sdb
brw-r----- 1 root disk 8, 32 Aug 22 10:00 /dev/sdc
brw-r----- 1 root disk 8, 48 Aug 22 10:00 /dev/sdd
```

Use the fdisk command to setup the sda, sdb, sdc, and sdd device files. For this paper we created three separate 4 gigabyte disk partitions; sda1,sdb1 and sdc1 for the voting disks. Then we created disk partitions; sda2,sdb2 and sdc2, using the remaining space on the disks. In addition, the whole sdd disk was partitioned into the sdd1 partition. This disk configuration is sufficient to demonstrate the Clusterware and ASM configuration. It should not be used as an optimal solution for any customer. This configuration is only created on one node. The other node will see the changes in the external disk, if the fiber connect is wired and configured correctly.

```
ls -l /dev/sd*
brw-r----- 1 root disk 8, 0 Aug 22 15:33 /dev/sda
brw-r----- 1 root disk 8, 1 Aug 22 15:42 /dev/sda1
brw-r----- 1 root disk 8, 2 Aug 22 15:42 /dev/sda2
brw-r----- 1 root disk 8, 16 Aug 22 15:34 /dev/sdb
brw-r----- 1 root disk 8, 17 Aug 22 15:42 /dev/sdb1
brw-r----- 1 root disk 8, 18 Aug 22 15:42 /dev/sdb2
brw-r----- 1 root disk 8, 32 Aug 22 15:36 /dev/sdc
brw-r----- 1 root disk 8, 33 Aug 22 15:42 /dev/sdc1
brw-r----- 1 root disk 8, 34 Aug 22 15:42 /dev/sdc2
brw-r----- 1 root disk 8, 48 Aug 22 15:37 /dev/sdd
brw-r----- 1 root disk 8, 49 Aug 22 15:42 /dev/sdd1
```

The disk partitions need to be labeled for the ASM system. Run the oracleasm configure command and set the default user to grid and the group to asmadmin. Select the Oracle ASM library driver to start on boot and scan the ASM disks on boot.

```
/etc/init.d/oracleasm configure
Configuring the Oracle ASM library driver.
```

This will configure the on-boot properties of the Oracle ASM library driver. The following questions will determine whether the driver is loaded on boot and what permissions it will have. The current values will be shown in brackets ('[]'). Hitting <ENTER> without typing an answer will keep that current value. Ctrl-C will abort.

```
Default user to own the driver interface []: grid
Default group to own the driver interface []: asmadmin
Start Oracle ASM library driver on boot (y/n) [n]: y
Scan for Oracle ASM disks on boot (y/n) [y]: y
Writing Oracle ASM library driver configuration: done
Initializing the Oracle ASMLib driver: [
OK]
Scanning the system for Oracle ASMLib disks: [
OK]
```

On one node run the `oracleasm createdisk` command labeling the voting disks and the data disks. Mark them as `sda1`, `sdb1` and `sdc1` disks as `OCR_VOTEX` where X is 1 for “sda1”, 2 for “sdb1” and 3 for “sdc1”. Mark the rest as data disks.

```
/etc/init.d/oracleasm createdisk OCR_VOTJD EDWARDS
ENTERPRISEONE /dev/sda1
Marking disk "OCR_VOTJD EDWARDS ENTERPRISEONE" as an ASM
disk: [OK]
/etc/init.d/oracleasm createdisk OCR_VOTE2 /dev/sdb1
Marking disk "OCR_VOTE2" as an ASM disk: [
OK]
/etc/init.d/oracleasm createdisk OCR_VOTE3 /dev/sdc1
Marking disk "OCR_VOTE3" as an ASM disk: [
OK]
/etc/init.d/oracleasm createdisk OCR_DATA1 /dev/sda2
Marking disk "OCR_DATA1" as an ASM disk: [
OK]
/etc/init.d/oracleasm createdisk OCR_DATA2 /dev/sdb2
Marking disk "OCR_DATA2" as an ASM disk: [
OK]
/etc/init.d/oracleasm createdisk OCR_DATA3 /dev/sdc2
Marking disk "OCR_DATA3" as an ASM disk: [
OK]
/etc/init.d/oracleasm createdisk ASM_DATA /dev/sdd1
Marking disk "ASM_DATA" as an ASM disk: [
OK]
```

On the other node run the `oracleasm scandisk` command to import the configuration from the other node.

```
/etc/init.d/oracleasm scandisks
```

Use the `oracleasm listdisks` command to make sure the `scandisk` command worked.

```
/etc/init.d/oracleasm listdisks
OCR_VOTJD EDWARDS ENTERPRISEONE
OCR_VOTE2
OCR_VOTE3
OCR_DATA1
```

```
OCR_DATA2
OCR_DATA3
ASM_DATA
```

## Install Clusterware

1. Download the grid infrastructure install software from the Clusterware site:  
<http://www.oracle.com/technetwork/database/database-technologies/clusterware/downloads/index.html>
2. Download to a temporary directory to extract the Clusterware installer.  

```
mkdir -p /u01/software/direct_gridInfra
ls -lrt
total 1908284
-rw-r--r-- 1 root root 201673595 Jul 18 10:51
linuxamd64_12c_grid_2of2.zip
-rw-r--r-- 1 root root 1750478910 Jul 18 11:08
linuxamd64_12c_grid_1of2.zip
```
3. Unzip the files, this will create a grid directory.  

```
ls -lrt
total 1908284
drwxr-xr-x 7 root root 4096 Jun 10 2013 grid
-rw-r--r-- 1 root root 201673595 Jul 18 10:51
linuxamd64_12c_grid_2of2.zip
-rw-r--r-- 1 root root 1750478910 Jul 18 11:08
linuxamd64_12c_grid_1of2.zip
```
5. Change to the grid directory and run the sshUserSetup.sh scripts both the oracle and grid users on one node.  

```
su - oracle
#cd /u01/software/direct_gridInfra/grid/sshsetup
./sshUserSetup.sh -user oracle -hosts "nodex1 nodex2" -advanced
-exverify -confirm

#su - grid
#cd /u01/software/direct_gridInfra/grid/sshsetup
./sshUserSetup.sh -user grid -hosts "nodex1 nodex2" -advanced -
exverify -confirm
```
6. Install cvuqdisk rpm on both nodes.  

```
cd /u01/software/direct_gridInfra/grid/rpm
CVUQDISK_GRP=asmadmin; export CVUQDISK_GRP
rpm -ivh ./cvuqdisk-1.0.9-1.rpm
Preparing...
[100%]
 1:cvuqdisk
[100%]
```
7. Run the runcluvfy.sh from the cvuqdisk-1.0.0-1.rpm script as grid user.  

```
#su - grid
./runcluvfy.sh stage -pre crsinst -n nodex1,nodex2 -osdba
asmdba -asm -asmdev ORCL:OCR_VOTJD EDWARDS
```



```
ENTERPRISEONE,ORCL:OCR_VOTE2,ORCL:OCR_VOTE3 -asmgrp asmadmin -
presence local -crshome /u01/app/12.1.0.1/grid
```

8. Setup /etc/hosts file. Add the ip and names for the node and the private network ip and names.

```
Do not remove the following line, or various programs
that require network functionality will fail.
127.0.0.1 nodelx01 localhost.localdomain localhost
#::1 localhost6.localdomain6 localhost6

101.39.194.89 nodelx01 nodelx01.domain.com
101.39.194.90 nodelx02 nodelx02.domain.com

192.168.0.101 nodelx01-priv
192.168.0.102 nodelx02-priv
```

## Installation of Oracle Clusterware Software

```
[grid@dnptlx110 grid]$ export DISPLAY= 101.39.115.53:0.0
[grid@dnptlx110 grid]$ ls
install rpm runInstaller stage
response runcluvfy.sh sshsetup welcome.html
[grid@dnptlx110 grid]$./runInstaller
```

1. Start the Oracle Clusterware installer.

```
[grid@dnptlx110 grid]$./runInstaller
Starting Oracle Universal Installer...

Checking Temp space: must be greater than 120 MB. Actual 110741 MB Passed
Checking swap space: must be greater than 150 MB. Actual 9919 MB Passed
Checking monitor: must be configured to display at least 256 colors. Actual 1
6777216 Passed
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2014-08-27_10
-34-09AM. Please wait ...
```

### Clusterware Installer Running

For this install we are skipping the software updates. Customers will want to update the software.

2. Select Installing and Configure Oracle Grid Infrastructure for a Cluster.
3. Select Configure a Standard Cluster.
4. Select the Advanced Configuration.
  - a. English Language
5. The Grid Plug and Play screen requires the cluster name and the SCAN name. For the Oracle RAC install, the SCAN name is the name of the database to attach to. The SCAN name is a valid name which resolves to an IP address on the domain name server for the network. This is not necessary for the simpler

Clusterware install. The cluster name and IP will be configured later with the install of the cluster agent and configuration of the JD Edwards EnterpriseOne resource. For now enter a default value. For this project we used nodelxJD EDWARDS ENTERPRISEONE for the clustername.

- a. Cluster name nodelxTMP
  - b. Scan name nodelxTMP
  - c. Scan port 1521
  - d. Don't configure GNS
6. Add the nodelx02 and nodelx02-vip hostnames.
  7. Select the eth1 to be private and not ASM.
  8. Select Yes to Configure Grid Infrastructure Management Repository.
  9. Select Use Standard ASM for storage.
  10. Set the disk group name to Vote and select the each of the previously configured Voting disks.
  11. Set up the ASM password.
  12. Do not use IPMI.
  13. Use the asmadmin, asmdba and asmoper account and groups.
  14. Select the previously created install location.
  15. Create Inventory
  16. Select Automatically run configuration scripts and enter the root password for your system.

The installer will now run the verifying process. If this report contains errors, they must be fixed. The installer may report a warning for the ASM device; in our configuration we were able to ignore these warnings.

17. The summary screen will display the configuration. Select the Install button to begin the install.

At some point during the install you will be prompted to run all the configuration scripts as root.

18. Select Close and the install is completed.

### Configuring the Shared Disk Array with the ASMCA Utility

1. Run asmca to set up shared disks.

As the grid user on nodelx01 run the asmca tool. The asmca tool is found at /u01/app/12.1.0.1/grid/bin.

---

*The asmca shows the VOTE group that was previously created.*

---

2. Click the Create button to add a disk group.

For this paper we selected the disk partition that remained on the three voting disks and created a Disk Group Name of DATA\_GRP\_1 with external Redundancy selecting the OCR\_DATA1, OCR\_DATA2 and OCR\_DATA3 labeled disk partitions and a DATA\_GRP\_2 with the full disk partition from the /dev/sdd disk.

3. Click OK.

The disk group will now appear in the previous Disk Groups tab of the asmca screen.

4. Mount the disk group, select the Mount All
5. Create a volume group from the DATA\_GRP\_2 and leave the DATA\_GRP\_1 group in reserve, in case it is needed later. To add a new volume group:
  - a. Select the Volumes tab.
  - b. Click Create, select the Disk Group Name DATA\_GRP\_2 from previous screen, fill in Volume name and set the Size Usable to the Usable Space minus 1 G bytes.
  - c. Click OK button.

- ASM will create the volumes group and display this in the Volumes tab.
6. Enable the volume group.
    - a. Select volume group
    - b. Click the enable all button.
  7. The asmca tool will display a confirmation pop up window. Click the “Yes” button.
  8. Make the volume group into an ASM cluster file system.
    - a. Select the ASM Cluster File Systems tab
    - b. Click the Create button
    - c. Select a “Cluster File System” for the Type of ACFS.
    - d. Enter a mount point of “/u02”.
    - e. Make sure the “Auto Mount” is checked.
 

The user name and user group should be “jde910”. This file system will host the JD Edwards EnterpriseOne application on either the JD EDWARDS ENTERPRISEONE Server or Batch server.
    - f. Enter a description for the file system.
    - g. Select the DG2\_Vol group created in the previous step.
    - h. Click the “OK” button.
 

The asmca will prompt you to run a script.
  9. Open a session on the machine and run the acfs\_scripts.sh.  
The script will mount the file system.
  9. Return to the asmca screen and click close on the screen prompting you to run the script. The “ASM Cluster File Systems” will now show the /u02 file system.
    - a. Exit the asmca.
 

The /u02 file system is known to both nodes in the cluster and is a resource in the clusterware system. The file system is mounted on the nodelx2 system. The crsctl command will show the file system under the control of the Clusterware.

## Installing the JD Edwards EnterpriseOne Platform Pack

To install the JDE Platform Pack into the shared disk

1. On the /u02 file system as the jde910 user create a directory /u02/jdedwards/e910 and /u02/jdedwards/PlatPack directories.

```
#su - jde910
$mkdir -p /u02/jdedwards/e910
...
$unzip V30261-01.zip
...
$unzip V35081-01.zip
...
$cd install
$ls
addLangs.sh detachHome.sh oneclick.properties resource
unzip
addNode.sh instImages oraparam.ini runInstaller
attachHome.sh lsnodes oraparamsilent.ini runSSHSetup.sh
$chmod a+x *.sh
$chmod a+x unzip
$./runInstaller &
```

2. Download the Platform Packs part files V30261-01.zip and V35081-01.zip into the /u02/jdedwards/e910 directory.
  - a. Unzip the part files.
  - b. Change to the install directory.
  - c. Change the script files and the unzip utility to be executable.
  - d. Execute the runInstaller program.
3. On the Platform Pack installer, click the Next button.
4. For this paper, we only installed the Development Environment, select the Custom option.
5. Set the path for the install to /u02/jdedward/e910 and click the Next button.
6. The installer will be aware that this is a cluster and request which nodes to install on. Select the nodelx02 as well as the nodelx01.
7. For this paper we only used the Development Environment and have the database on the RAC node. Select only the Enterprise Foundation and Development Environment, and click Next.
8. Select the Install option and click Next.
9. Enter the type of database and database scan name from the RAC install, click Next.
10. Enter the database connect string. For this paper we used, jde910. Enter the value and click Next.
11. The summary screen will appear, click the Install button.

The install will start and prompt to run a script.

```
[root@ nodelx01 SharedScripts]# addacct.sh
-bash: addacct.sh: command not found
[root@ nodelx01 SharedScripts]# ./addacct.sh
Add user for running services
.
groupadd: group jde910 exists
useradd: user jde910 exists
Failed to add user jde910
Change password for jde910
Changing password for user jde910.
New UNIX password:
BAD PASSWORD: it is based on a dictionary word
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
Running on Linux
Adding user oracle to group jde910
Running on Linux
Edit .profile for jde910
.
Please be patient - changing ownership of /u02/jdedwards/e910
to jde910 - this could take several minutes.
.
End of addacct.sh processing
[root@ nodelx01 SharedScripts]#
```

12. Return to the command prompt and run the script. When the script finishes, return to the Install Screen and click the OK button.

The install will complete.

## Install the JDE Server Manager Agent

**Important:** Do not install the JD Edwards Server Manager Agent to the /u02 shared disk. An install of Server Manager Agent on the share drive will not work. Install the Server Manager Agent on a file system local to each node.

1. Download the Server Manger Agent from a Server Manager Console. As the jde910 user, unzip the linux.zip file. Change to the Disk1/install directory extracted from the zip file. If necessary change the script and unzip files to execute with the change mode command.

```
$chmod a+x *.sh unzip
```

2. Run the runInstaller program as jde910 user.  
The Server Manager Agent installer starts.
3. Install on a local file system. In this case we selected /u01 and the jde\_sm directory, click Next.
4. The installer will detect the Clusterware software. Select both nodes, click Next.
5. Enter the local and the HTTP port for the Server Manager Console, click Next.
6. The summary of what will be installed is displayed, click Install.  
The install of the Server Manager Agent will begin.  
The install completes.
7. Navigate to the Server Manager Console and create a managed instance.

The screenshot shows the Oracle JD Edwards EnterpriseOne Server Manager interface. The top navigation bar includes "ORACLE JD Edwards EnterpriseOne Server Manager", "Server Manager Documentation", "EnterpriseOne Tools Documentation", and "Sign Out". The main content area is titled "Managed Homes and Managed Instances". It features a sidebar with navigation options: "What do you want to do?" (INSTALL: Management Agents, Manage Software, Database Drivers; CONFIGURE: Server Manager Users, Server Groups; TRACK: User Activity, Server Activity, Table Cache), a "Select Instance..." dropdown, and a "Select View" dropdown set to "Managed Homes and Managed Instances". The main area displays a table of Managed Homes and Managed Instances.

| Managed Home Location                                              | Managed Instances                     |
|--------------------------------------------------------------------|---------------------------------------|
| <input type="checkbox"/> node1x01.domain.com<br>/u01/jde_910/SCFHA | No managed instances.                 |
| <input type="checkbox"/> denovm38e<br>/u01/Oracle/smc_home/SCFMC   | home<br>Management Console<br>Running |

8. Click the node1x01 Managed Home location.
9. Click the Create New Managed Instance button
10. Select the Register an Existing Enterprise Server selection.
11. For this example we named the instance JD EDWARDS ENTERPRISEONE\_01\_instance. Enter the path to the shared JD EDWARDS ENTERPRISEONE install and click Continue.
12. The Database and network configuration should be picked up from the existing configuration. Click the Continue button.

The new instance will be displayed on the Managed Homes and Managed Instances screen. This same procedure should be repeated on the other node. Register the instance with the name JD EDWARDS ENTERPRISEONE\_02\_instance on the node1x02 node. The Server Manager Console will now display two

managed home locations, one for nodelx01 and one for nodelx02. In addition there will be a separate managed instance for each of these nodes.

---

**Important:** Do not use the Server Manger Console to start or stop the JD EDWARDS ENTERPRISEONE server on the Clusterware. If you start a JD EDWARDS ENTERPRISEONE instance using Server Manager and then start another JD EDWARDS ENTERPRISEONE instance with Clusterware the two instances will corrupt the database and spec files. If you stop a JD EDWARDS ENTERPRISEONE instance outside of Clusterware it is interrupted as software fail by Clusterware and will be restarted. It is critically important that you do not use Server Manager to start or stop JD EDWARDS ENTERPRISEONE servers. However, with this configuration you can apply tool releases to the instances. Distribute the tools software using the Managed Software Components screen to both managed locations. Select one of the instances and use the Change button to apply the change to the instance. Appendix A contains a script change to the RunOneWorld.sh that will prevent anything but the Clusterware from starting the EnterpriseOne Server.

---

At this point you can run the workbench installation of the regular EnterpriseOne install procedure.

### Installing the Clusterware Agent

These instructions are for a system on which a grid infrastructure has previously been installed. This example has the nodelx01 and nodelx02. Your system may have more nodes of different names. The owner of the grid infrastructure is the grid user with a group of oinstall.

1. Download the agent from: <http://www.oracle.com/technetwork/database/database-technologies/clusterware/downloads/index.html>
2. Create a directory on one of your cluster nodes and download the file there.
3. Download or copy the xagpack5.zip file to this location. Make sure the file and directory are owned by the grid user.

```
chown -r grid:oinstall /u01/software/agent
 a. As the grid user unzip the file.
 su - grid
 unzip xagpack_5.zip
 Afterwards there will be an install and xag directory.
```

```
[root@nodelx01 agent5]# ls -lrt
total 148
drwxrwxr-x 3 grid oinstall 4096 Jul 23 01:42 install
drwxrwxr-x 7 grid oinstall 4096 Jul 23 01:42 xag
-rw-r--r-- 1 grid oinstall 137309 Jul 29 11:01 xagpack_5.zip
[root@nodelx01 agent5]# █
```

Extracted xagpack\_5.zip

4. Before installing, create an install directory on each node in your cluster owned by the grid user.

```
mkdir -p /u01/cluster_agent
ssh grid@nodelx01 'mkdir -p /u01/cluster_agent'
```
5. Change to the /u01/software/agent5/xag directory and run the xagsetup.sh command with install, directory and all nodes parameters.

```
./xagsetup.sh --install --directory /u01/cluster_agent --
all_nodes
```

---

Running the `xagsetup.sh`, will install the new `agctl` agent on the nodes.

---

6. You have just installed the stand alone agent. Make sure to execute the `/u01/cluster_agent/bin/agctl` command and add it to the front of your `PATH` environment variable `/u01/cluster_agent/bin`
7. Add this line to the end of the grid `.bashrc` file

```
export PATH=/u01/cluster_agent/bin:${PATH}
```
8. Exit and re-enter the grid account and do the following:

```
which agctl
```

This should report the `/u01/cluster_agent/bin/agctl` path and will replace the use of the new Clusterware agent instead of the agent installed with the Clusterware software.

### Start the EnterpriseOne Application under Clusterware Control

1. Change to the root user. This is required to create the shared ip for the JD EDWARDS ENTERPRISEONE server. Run the command to install the JD EDWARDS ENTERPRISEONE resources.

```
[root@node1x01 bin]# ./agctl add jde_enterprise_server jde910 --jde_home /u02/jdedwards/e910 --network 1 --ip 101.39.194.81 --user jde910 --group oinstall
[root@node1x01 bin]#
```

Running of the `agctl` command to install the EnterpriseOne resources.

The shared IP is 101.39.194.81. This IP will be shared between the two nodes of the cluster. This will be the IP address of the JD EDWARDS ENTERPRISEONE application. The `/u02/jdedwards/e910` is the location of the installed JD EDWARDS ENTERPRISEONE server. The user and group is the user that will start and stop the JD EDWARDS ENTERPRISEONE application. The `jde_enterprise_server` is a fixed parameter. The `jde910` is the name the cluster resource will be given.

2. You can check that the install went correctly by listing the resource. Use the command `crsctl` to show the status of the resource.

```
[root@node1x01]$crsctl status res -w 'NAME st xag.'
NAME=xag.jde910-vip.vip
TYPE=app.appvip_net1.type
TARGET=OFFLINE
STATE=OFFLINE

NAME=xag.jde910.jde
TYPE=xag.jde.type
TARGET=OFFLINE
STATE=OFFLINE

[root@node1x01]$
```

Installation of the EnterpriseOne Application

The `-w` parameter allows you to filter the results. The filter shown here shows all Clusterware resource names `NAME` that start `st` with the `xag` string. Notice that there are two resources. The first, named `xag.jde910-vip.vip` is the virtual IP (VIP) that was specified in the `agctl add` command. Note that the name `jde910` was used. The

second resource named xag.jde910.jde is the JD EDWARDS ENTERPRISEONE server resource. The TARGET and the STATE values are OFFLINE. This means the resource has been installed, but is not running.

3. To start the JD EDWARDS ENTERPRISEONE server and bring the system online use the agctl start command.

```
[root@node1x01]$agctl start jde_enterprise_server jde910
[root@node1x01]$crsctl status res -w 'NAME st xag.'
NAME=xag.jde910-vip.vip
TYPE=app.appvip_net1.type
TARGET=ONLINE
STATE=ONLINE on node1x01

NAME=xag.jde910.jde
TYPE=xag.jde.type
TARGET=ONLINE
STATE=ONLINE on node1x01

[root@node1x01]$
```

Starting jde910 with the agctl command.

By using this command both the VIP and the server are started.

4. Switch to the jde910 user and run portttest to test if the JD EDWARDS ENTERPRISEONE is configured correctly.

```
[root@node1x01]# su - jde910
The ORACLE_HOME has been set to /u02/app/oracle/product/12.1.0/client_1
The PATH has been set to /u02/app/oracle/product/12.1.0/client_1/bin:/usr/kerberos/bin:/usr/local/bin:/bin:/usr/bin:/u02/jdedwards/e910/system/bin32

$ portttest JDE JDE DV910
Running portttest for JDE on DV910,...

Initializing Environment DV910,...
Environment DV910 was initialized successfully.

Initializing Environment DV910 Using a Token,...
Initializing JDE (User),...
JDE (User) Initialized successfully.

Opening table F986110,...
Opened table F986110 successfully.

Closing table F986110,...
Closed table F986110 successfully.

Opening table F0902,...
Opened table F0902 successfully.
```

5. Switch to the jde910 user and run portttest

```
[3] 00017574 00
[2] 00017582 8500000
[1] 00017582 36300000
[0] 00017582 8956200
Total number of rows printed = 99.

Calling DataDictionary Validation function,...
Data Dictionary Validation Succeed for "CO 00001".

Closing table F0902,...
Closed table F0902.

Freeing user JDE,...
Freed user JDE successfully.

Cleaning up environment DV910,...
Cleaned up environment DV910 successfully.

Congratulations! Porttest completed successfully.

All Done!
Bye!
$ █
```

Porttest completes successfully.

---

*Warning: If you log onto the nodex02 as the jde910 user and run the porttest command it will fail.*

---

## Basic Commands to Control the Cluster and Resources

The command to start and stop the cluster node including resources on the node:

```
crsctl start cluster -n nodebt01
crsctl stop cluster -n nodebt01
```

The command to check the status of all nodes in the cluster:

```
crsctl check cluster -all
```

The commands to stop the EnterpriseOne server and the virtual IP address:

```
crsctl stop res xag.jde910.jde
crsctl stop res xag.jde910-vip.vip
```

The commands to start the EnterpriseOne server and the virtual IP address:

```
crsctl start res xag.jde910-vip.vip -n nodebt01
crsctl start res xag.jde910.jde -n nodebt01
```

Example of a Clusterware stop for a single node:

```

#crsctl stop cluster -n nodebt01
CRS-2673: Attempting to stop 'ora.crsd' on 'nodebt01'
CRS-2790: Starting shutdown of Cluster Ready Services-managed resources
on 'nodebt01'
CRS-2673: Attempting to stop
...
CRS-2677: Stop of 'ora.scan2.vip' on 'nodebt01' succeeded
CRS-2672: Attempting to start 'ora.scan2.vip' on 'nodebt02'
CRS-2677: Start of 'ora.scan2.vip' on 'nodebt02' succeeded
...
Checking all nodes in the cluster. The nodebt01 is down, clusterware and its resources are down.
#crsctl check cluster -all

nodebt01
CRS-4535: Cannot communicate with Cluster Ready Services
CRS-4530: Communications failure contacting Cluster Synchronization
Services daemon
CRS-4534: Cannot communicate with Event Manager

nodebt02
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronizaiton Services is online
CRS-4533: Event Manager is online

To stop the resource but not the clusterware use the "res" resource option
#crsctl status res -w 'NAME st xag.jde'
NAME=xag.jde910-vip.vip
TYPE=app.appvip_net1.type
TARGET=ONLINE
STATE=ONLINE on nodebt01

NAME=xag.jde910.jde

```



```
TYPE=xag.jde.type
```

```
TARGET=ONLINE
```

```
STATE=ONLINE on nodebt01
```

```
#crsctl stop res xag.jde910.jde
```

```
CRS-2673: Attempting to stop 'xag.jde910.jde' on 'nodebt01'
```

```
CRS-2677: Stop of 'xag.jde910.jde' on 'nodebt01' succeeded
```

```
#crsctl status res -w 'NAME st xag.jde'
```

```
NAME=xag.jde910-vip.vip
```

```
TYPE=app.appvip_net1.type
```

```
TARGET=ONLINE
```

```
STATE=ONLINE on nodebt01
```

```
NAME=xag.jde910.jde
```

```
TYPE=xag.jde.type
```

```
TARGET=OFFLINE
```

```
STATE=OFFLINE on nodebt01
```

The resource xag.jde910.jde the EnterpriesOne server is offline, however the vip is still online. The vip is the virtual IP address.

```
#crsctl stop res xag.jde910-vip.vip
```

```
CRS-2673: Attempting to stop 'xag.jde910-vip.vip' on 'nodebt01'
```

```
CRS-2677: Stop of 'xag.jde910-vip.vip' on 'nodebt01' succeeded
```

```
#crsctl status res -w 'NAME st xag.jde'
```

```
NAME=xag.jde910-vip.vip
```

```
TYPE=app.appvip_net1.type
```

```
TARGET=OFFLINE
```

```
STATE=OFFLINE on nodebt01
```

```
NAME=xag.jde910.jde
```



```
TYPE=xag.jde.type
```

```
TARGET=OFFLINE
```

```
STATE=OFFLINE on nodebt01
```

Both the EnterpriseOne server and virtual IP are now stopped. To restart on the nodebt02, start the vip first and the EnterpriseOne server next. If the EnterpriseOne server starts before the vip is started the server will fail to start.

## Setting Up a Batch Server

To set up a JD EDWARDS ENTERPRISEONE Batch Server repeat the same procedure that was used to set up a clustered Oracle JD Edwards EnterpriseOne Installation and Upgrade Documentation. This is in chapter 16, Adding a Server to an Existing Configuration at:

[https://docs.oracle.com/cd/E24902\\_01/](https://docs.oracle.com/cd/E24902_01/)

[https://docs.oracle.com/cd/E24902\\_01/doc.91/e23313/adding\\_a\\_server.htm#AddServer](https://docs.oracle.com/cd/E24902_01/doc.91/e23313/adding_a_server.htm#AddServer)

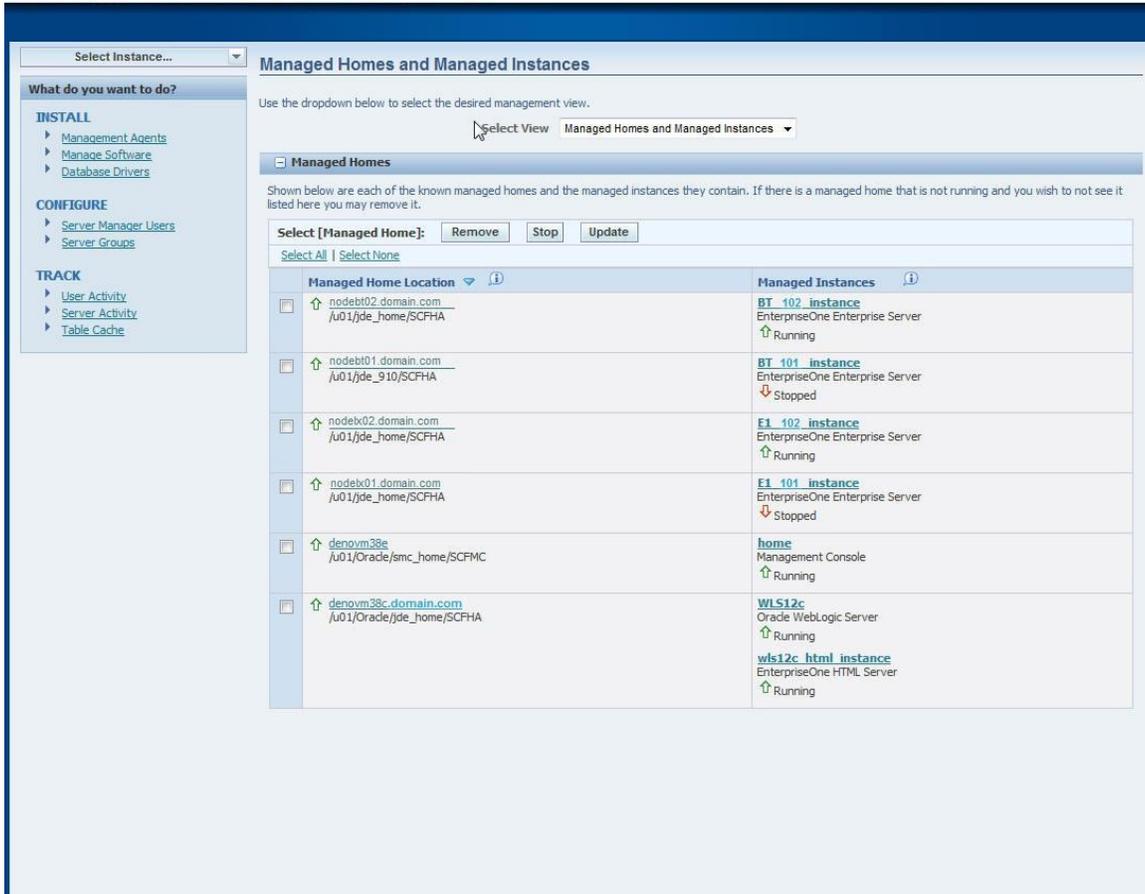
## Setting Up the Web Server for Clustering

---

*Note: The Web Server was setup using the document, JD Edwards EnterpriseOne 9.1.4 Clustering Best Practices with Oracle WebLogic Server 12.1.2 [Note:1935421.1](#)*

---

## Complete JD Edwards EnterpriseOne High Availability Architecture System



Server Manager Console Showing the HAA Environment.

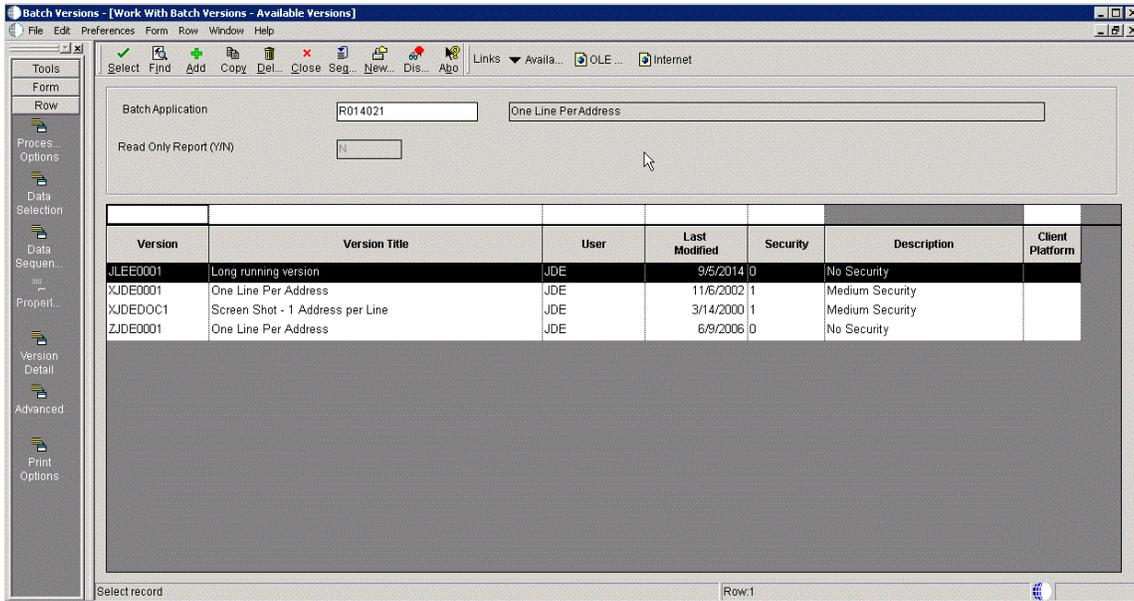
The complete HAA environment has the following items:

- » Clustered JD EDWARDS ENTERPRISEONE server nodelx01 and nodelx02 - service is running on JD EDWARDS ENTERPRISEONE\_102\_instance.
- » Clustered batch servers nodebt01 and nodebt02 - service is running on BT\_102\_instance.
- » Clustered web servers running on denovm38c.domain.com - service wls12c\_html\_instance.

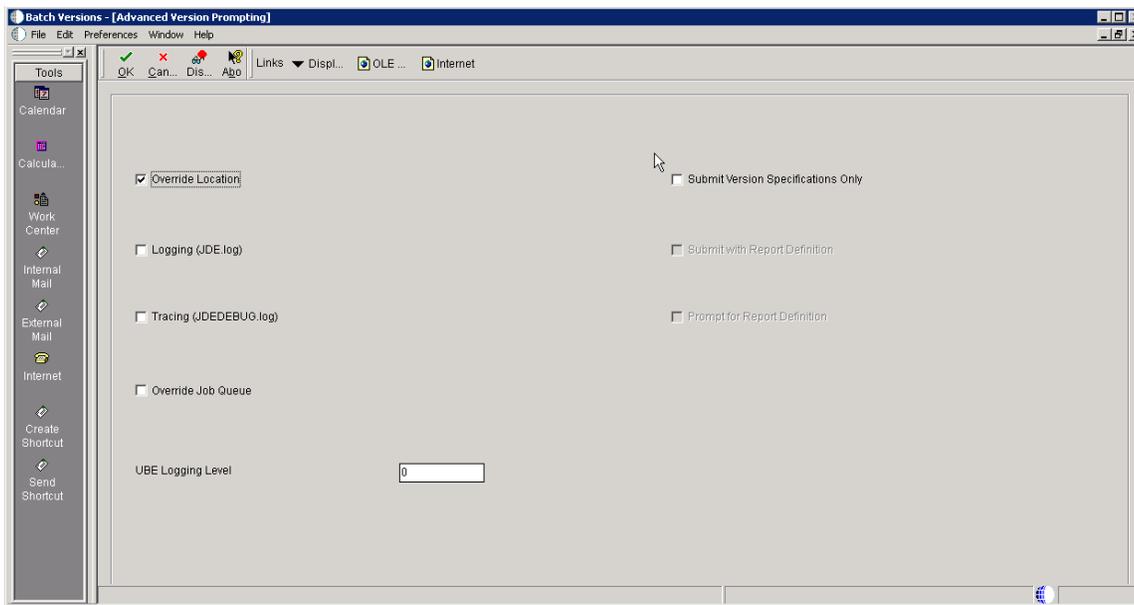
### Batch Behavior During Fail-over

- » Database Failover
- » JD EDWARDS ENTERPRISEONE Failover
- » Batch Failover

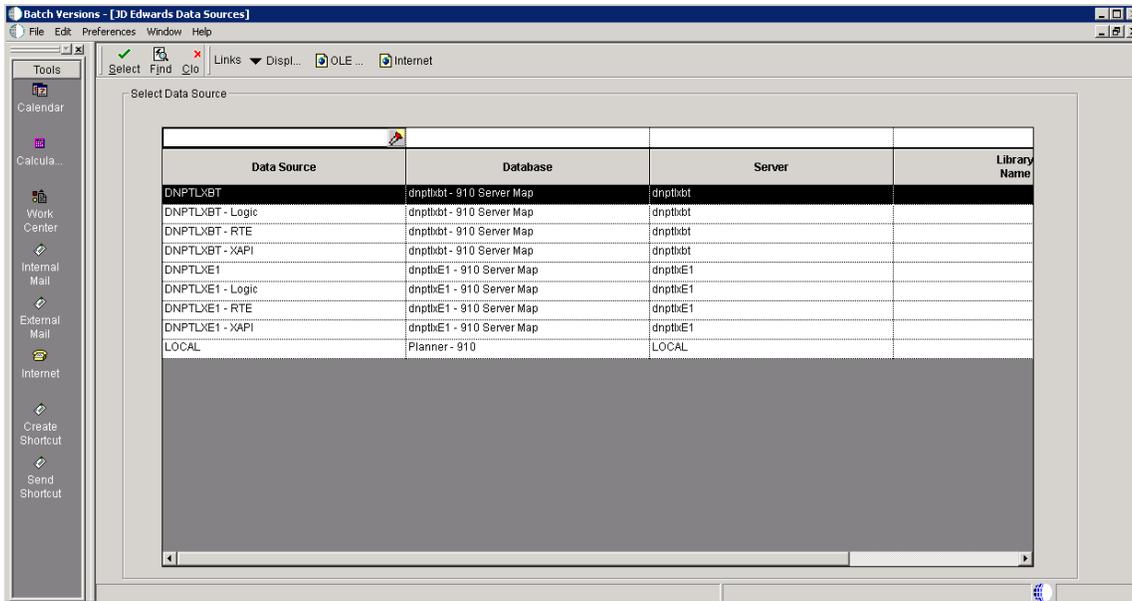
To test the failover of the Batch server a new version (JLEE0001) of the R014021 report was created. This version repeatedly listed the address book report for each entry in the address book. In other words, it ran an extremely long time without pausing. By doing this we could insure failing over the cluster while the report was processing. We expect that the report will fail, but we will not lose connectivity and will be able to simply resubmit the report and have it complete.



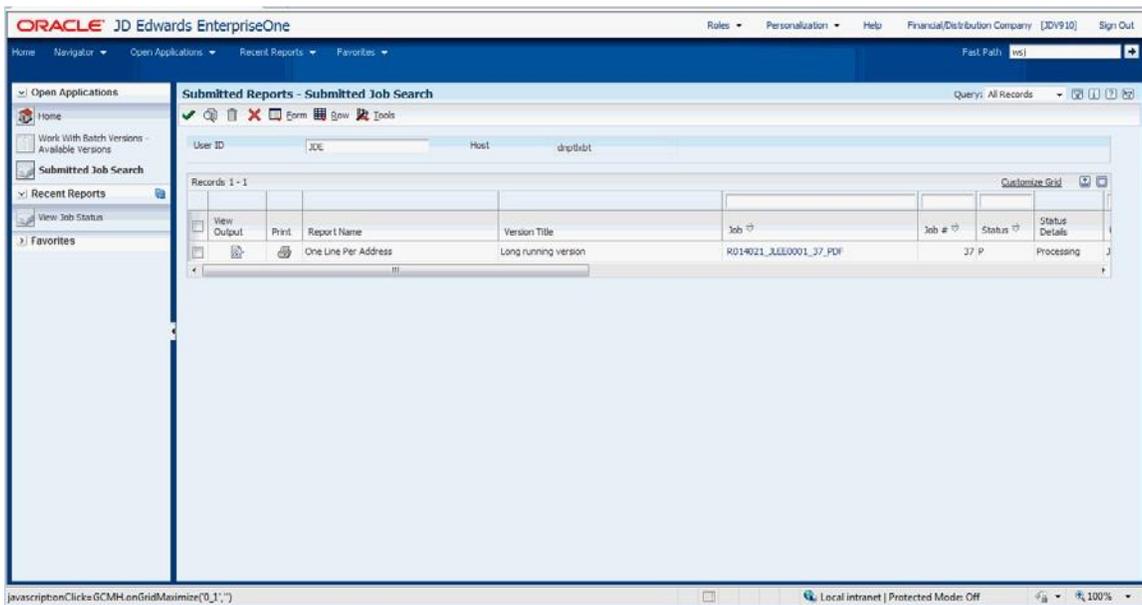
1. Select the JLEE0001 version of the R014021 report.



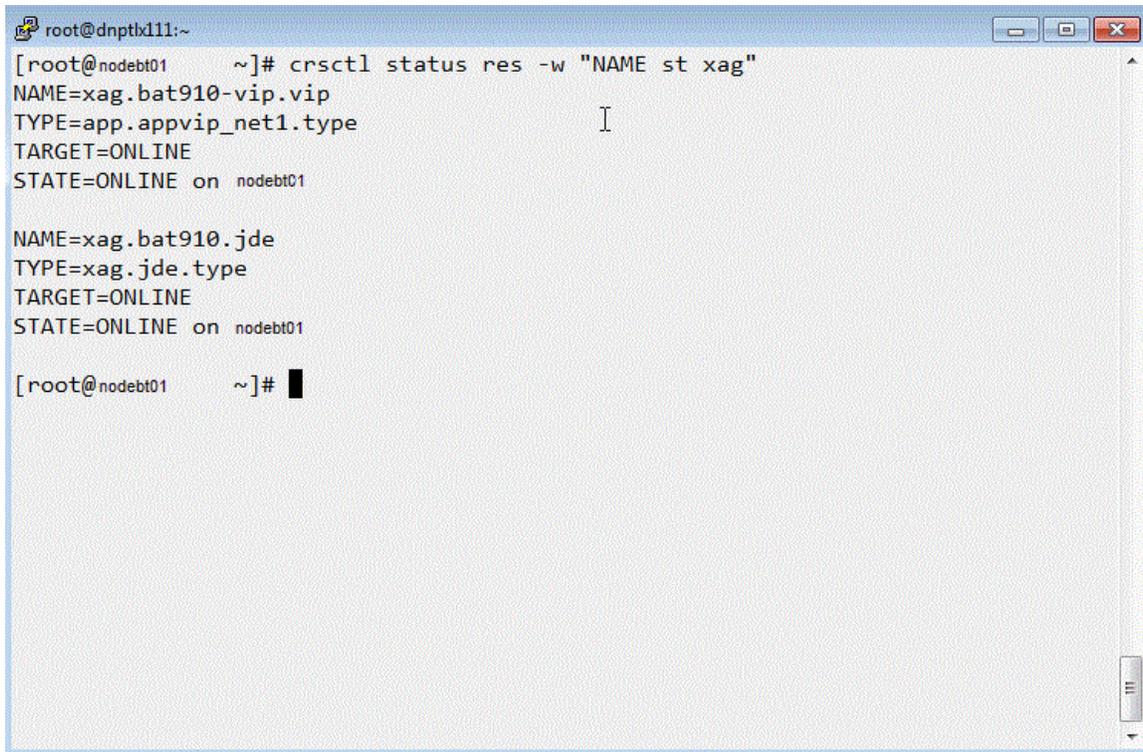
2. Select the advanced processing options and check Override Location.



3. Select the Batch Server DNPTLXBT from the Database Sources. This is the Data source that corresponds to the clustered batch server



4. Submit the batch job and then open the wsj application. Select the Form exit, Advanced then select the queue running on the DNPTLXBT server. This shows the R014021 report in the P status.

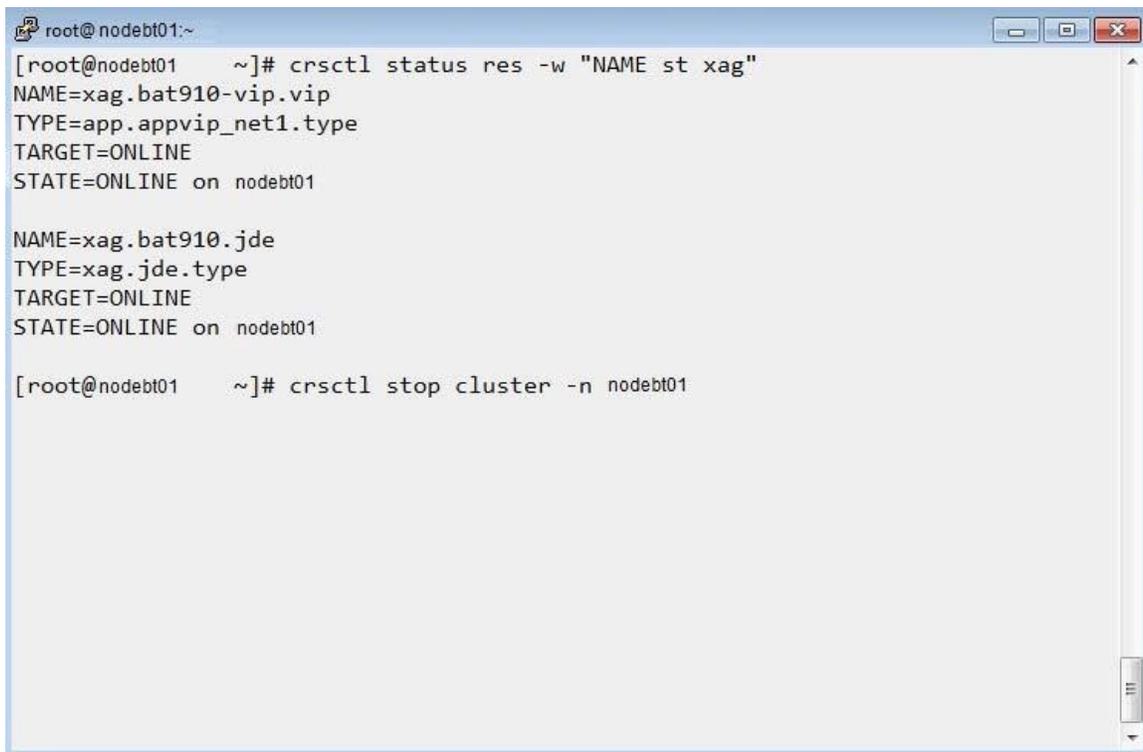


```
root@dnptx111:~
[root@nodebt01 ~]# crsctl status res -w "NAME st xag"
NAME=xag.bat910-vip.vip
TYPE=app.appvip_net1.type
TARGET=ONLINE
STATE=ONLINE on nodebt01

NAME=xag.bat910.jde
TYPE=xag.jde.type
TARGET=ONLINE
STATE=ONLINE on nodebt01

[root@nodebt01 ~]# █
```

The Batch Server resources, bat910 and VIP are running on the nodebt01 server.

A terminal window titled 'root@nodebt01:~' with standard window controls (minimize, maximize, close) in the top right. The terminal shows the following commands and output:

```
[root@nodebt01 ~]# crsctl status res -w "NAME st xag"
NAME=xag.bat910-vip.vip
TYPE=app.appvip_net1.type
TARGET=ONLINE
STATE=ONLINE on nodebt01

NAME=xag.bat910.jde
TYPE=xag.jde.type
TARGET=ONLINE
STATE=ONLINE on nodebt01

[root@nodebt01 ~]# crsctl stop cluster -n nodebt01
```

The crsctl stop command is used to stop the nodebt01 server. This simulates the failing of the nodebt01 hardware without pulling the plug.

```

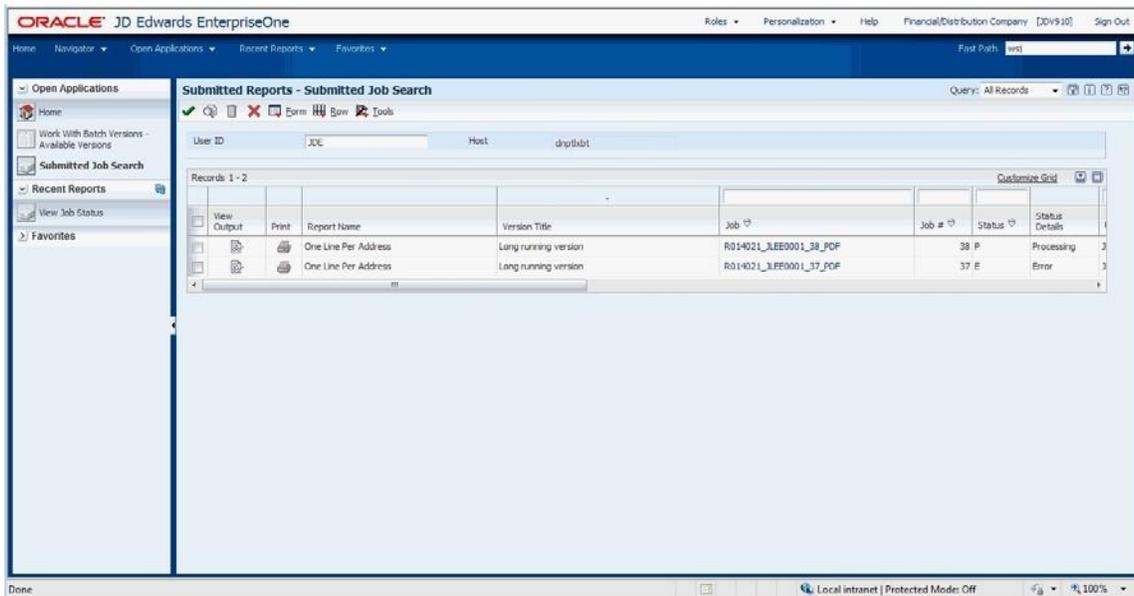
root@nodebt02:~
[root@nodebt02 ~]# crsctl status res -w "NAME st xag"
NAME=xag.bat910-vip.vip
TYPE=app.appvip_net1.type
TARGET=ONLINE
STATE=ONLINE on nodebt02

NAME=xag.bat910.jde
TYPE=xag.jde.type
TARGET=ONLINE
STATE=ONLINE on nodebt02

[root@nodebt02 ~]#

```

The Batch Server is running on the nodebt02 node. We now follow the same steps as previously and resubmit the job.



The failed batch job and the resubmitted job start to process. The job is now running on the dnptlxbt data source running on the nodebt02 node.

The screenshot displays the 'Submitted Reports - Submitted Job Search' window in Oracle JD Edwards EnterpriseOne. The window title is 'Submitted Reports - Submitted Job Search' and the user is 'JDE' on host 'dnp01bot'. The interface includes a navigation pane on the left with options like 'Open Applications', 'Submitted Job Search', 'Recent Reports', 'View Job Status', and 'Favorites'. The main area shows a table of submitted jobs. The table has columns for 'View Output', 'Print', 'Report Name', 'Version Title', 'Job #', 'Status', and 'Status Details'. Two records are shown: one for 'One Line Per Address' with version 'Long running version', job # 'R014021\_XFF0001\_38\_PDF', status '38 D', and 'Done'; and another for 'One Line Per Address' with version 'Long running version', job # 'R014021\_XFF0001\_37\_PDF', status '37 E', and 'Error'. The status 'Done' indicates the job has completed successfully, while 'Error' indicates a problem.

| View Output | Print | Report Name          | Version Title        | Job #                  | Status | Status Details |
|-------------|-------|----------------------|----------------------|------------------------|--------|----------------|
|             |       | One Line Per Address | Long running version | R014021_XFF0001_38_PDF | 38 D   | Done           |
|             |       | One Line Per Address | Long running version | R014021_XFF0001_37_PDF | 37 E   | Error          |

The job completes.

## High Availability (HA) Design for Oracle JD Edwards EnterpriseOne on the Oracle Public Cloud

When implemented in the Oracle Public Cloud, HA for JD Edwards EnterpriseOne is achieved by building in multiple levels of fault tolerance and load balancing capabilities. The purpose of this implementation is to prevent any loss of service. HA for JD Edwards EnterpriseOne in the Oracle Public Cloud is attained at three different layers:

- » Database layer – Oracle Real Application Clusters (RAC) with Oracle Database Cloud Service (DBCS)
- » Application layer – Aliasing virtual hostnames and the Network File System (NFS)
- » Web layer – Oracle Traffic Director (OTD) and Oracle HTTP Server (OHS)

### Oracle RAC with Database Cloud Service (DBCS)

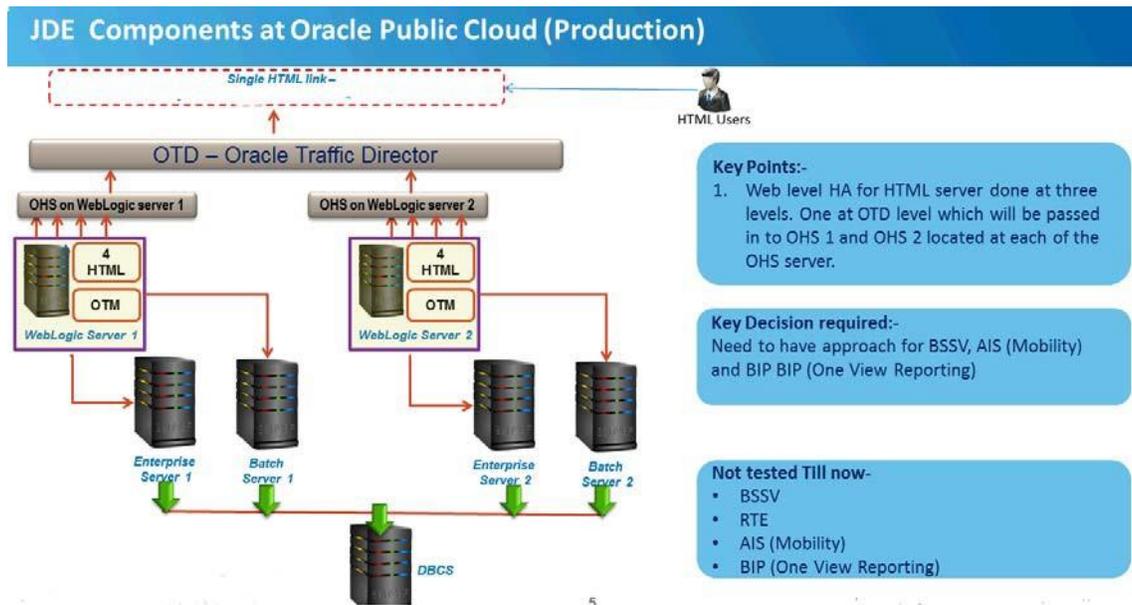
HA for the JD Edwards EnterpriseOne database is attained by using the Oracle DBCS with RAC.

#### See Also

- » Creating a DBCS RAC instance:  
<https://docs.oracle.com/en/cloud/paas/database-dbaas-cloud/csdbi/use-rac-this-service.html>
- » Deploying the JD Edwards EnterpriseOne Database as a Multi Tier Server in the Oracle Compute Cloud Service:  
[http://www.oracle.com/webfolder/technetwork/tutorials/obe/cloud/compute-iaas/deploying\\_jde\\_dbserver\\_multitier\\_92/dbcs.htm](http://www.oracle.com/webfolder/technetwork/tutorials/obe/cloud/compute-iaas/deploying_jde_dbserver_multitier_92/dbcs.htm)

### Oracle Traffic Director (OTD) and Oracle HTTP Server (OHS)

HA for the JD Edwards EnterpriseOne HTML Servers is attained by using the OTD and OHS as load balancers.



The above graphic illustrates the configuration tested by this technical brief:

- » A set of 4 HTML Servers that are deployed on two different Weblogic Servers residing on two different physical machines

- » HTML Servers that are present on WLS1 and are load balanced by OHS1
- » HTML Servers that are present on WLS2 and are load balanced by OHS2
- » OHS1 and OHS2 URLs are load balanced by OTD
- » OTD is configured with OAM to provide a single SSO URL

## Virtual Hostnames Setup and Configuration

HA for the JD Edwards EnterpriseOne application server and batch server is attained by setting up and configuring virtual hostnames.

For this technical brief the following hostnames were used for one of the production environments:

- » Enterprise Server Set 1 hostnames:

jdeprodapp1  
jdeprodbatch1

- » Enterprise Server Set 2 hostnames:

jdeprodapp2  
jdeprodbatch2

- » HTML Server hostnames:

jdeweb1  
jdeweb2

| Physical/Actual Hostnames | Virtual/Logical Hostnames |
|---------------------------|---------------------------|
| jdeprodapp1               | jdeprodapp                |
| jdeprodapp2               | jdeprodapp                |
| jdeprodube1               | jdeprodube                |
| jdeprodube2               | jdeprodube                |

### Aliasing for the Applications Server

To provide aliasing for the applications server:

13. Add the alias jdeprodapp in the /etc/hosts file of jdeprodapp1 and jdeprodapp2.  
For example, the host file of jdeprodapp1 may look like:

#### 10.0.0.9 jdeprodapp1 jdeprodapp

```
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.9 jdeprodapp1 jdeprodapp
~
~
~
```

14. Add the alias jdeprodapp in the /etc/hosts file of jdeprodapp2.  
For example, the host file of jdeprodapp1 may look like:

#### 10.0.0.10 jdeprodapp2 jdeprodapp

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.10 jdeprodapp2 jdeprodapp
~
~
~

```

15. Add the alias jdeprodapp for jdeprodapp1 in the /etc/hosts file of the host machine of WLS1.  
For example, the host file of jdeweb1 may look like:

#### 10.0.0.9 jdeprodapp1 jdeprodapp

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.9 jdeprodapp1 jdeprodapp
~
~
~

```

16. Add the alias jdeprodapp for jdeprodapp2 in the /etc/hosts file of the host machine of WLS2.  
For example, the host file of jdeweb2 may look like:

#### 10.0.0.10 jdeprodapp2 jdeprodapp

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.10 jdeprodapp2 jdeprodapp
~
~
~

```

### Aliasing for the Batch Server

To provide aliasing for the Batch Server:

17. Add the alias jdeprodube in the /etc/hosts file of jdeprodube1.  
For example, the host file of jdeprodube1 may look like:

#### 10.0.0.11 jdeprodube1 jdeprodube

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.11 jdeprodube1 jdeprodube
~
~
~

```

18. Add the alias jdeprodube in the /etc/hosts file of jdeprodube2.  
For example, the host file of jdeprodube2 may look like:

#### 10.0.0.12 jdeprodube2 jdeprodube

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.12 jdeprodube2 jdeprodube
~
~
~

```

19. Add the alias jdeprodube for jdeprodaube1 in the /etc/hosts file of host machine of WLS1.  
For example, the host file of jdeweb1 may look like:

#### 10.0.0.11 jdeprodube1 jdeprodube

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.11 jdeprodube1 jdeprodube
~

```

20. Add the alias jdeprodube for jdeprodube2 in the /etc/hosts file of host machine of WLS2. For example, the host file of jdeweb2 may look like:

### 10.0.0.12 jdeprodube2 jdeprodube

```

127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
10.0.0.12 jdeprodube2 jdeprodube
~

```

## Modifying JD Edwards EnterpriseOne to Listen to the Logical Name

To enable JD Edwards EnterpriseOne to listen to the logical name:

21. Create an Add-on server plan and run workbench to add virtual hostnames (jdeprodapp and jdeprodube) to the database table. The following sample screenshots provide an overview of how to create an ADD-ON plan for the Enterprise Server jdeprodapp:

The screenshot shows the 'Add-on Servers Plan - [Installation Planner]' window. The 'Plan Information' section contains the following fields:

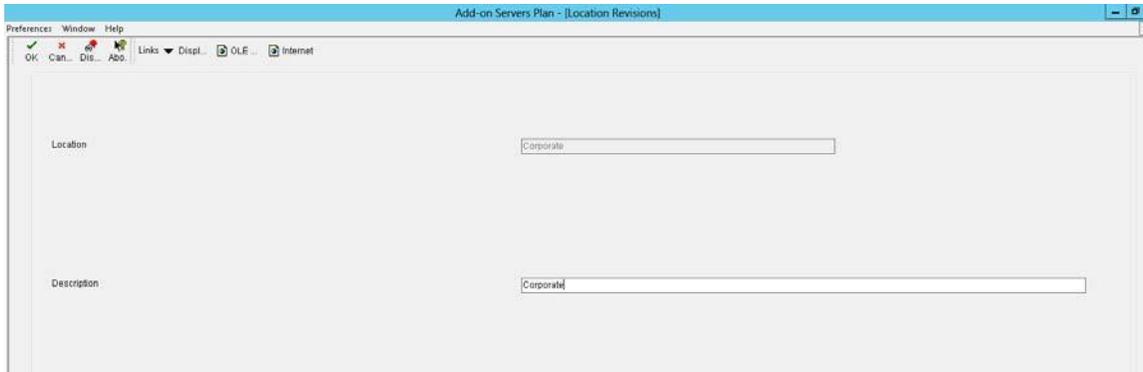
- Name: ADDONAPP
- Description: Adding JDEPRODAPP enterprise
- Status: 10 In Planning
- Release: To Release E920

### Plan Information

The screenshot shows the 'Add-on Servers Plan - [Location Search]' window. The table below displays the location information:

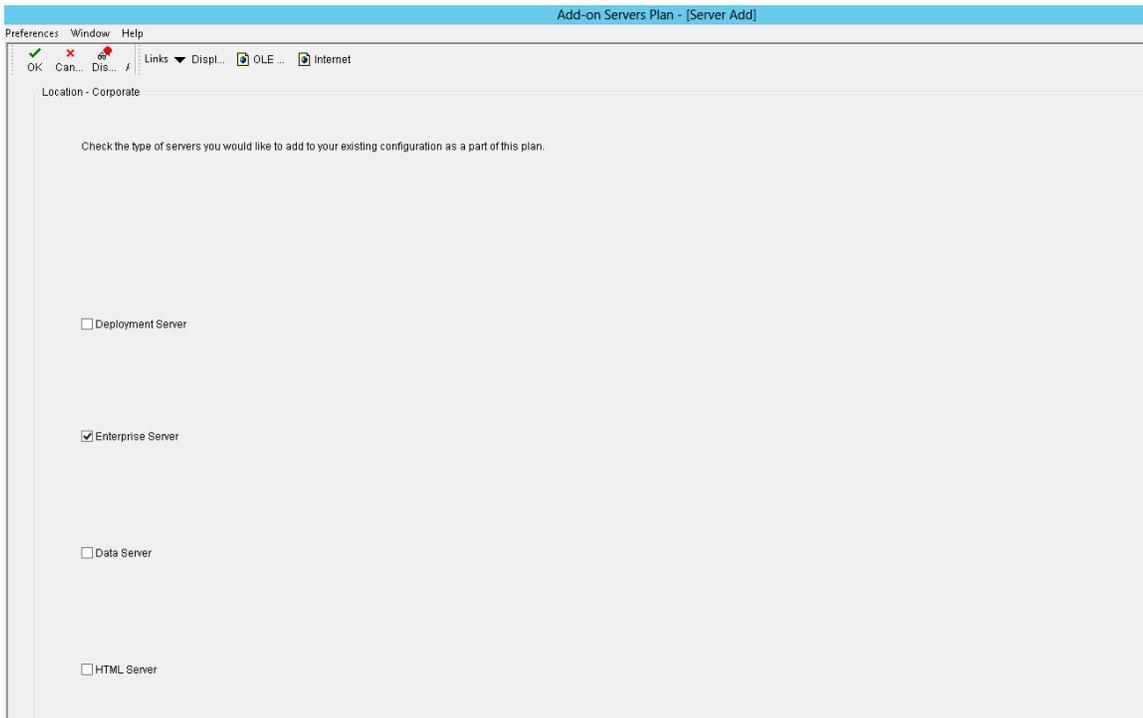
| Location  | Description     | Location Code | Parent Location | Local Development |
|-----------|-----------------|---------------|-----------------|-------------------|
| Corporate | Denver location |               |                 | N                 |

## Location Search Form



The screenshot shows a dialog box titled "Add-on Servers Plan - [Location Revisions]". It has a menu bar with "Preferences", "Window", and "Help". Below the menu bar is a toolbar with icons for "OK", "Cancel", "Dismiss", "Apply", "Links", "Display...", "OLE...", and "Internet". The main area contains two text input fields: "Location" and "Description". Both fields contain the text "Corporate".

## Location Revisions Form



The screenshot shows a dialog box titled "Add-on Servers Plan - [Server Add]". It has a menu bar with "Preferences", "Window", and "Help". Below the menu bar is a toolbar with icons for "OK", "Cancel", "Dismiss", "Apply", "Links", "Display...", "OLE...", and "Internet". The main area is titled "Location - Corporate" and contains the text: "Check the type of servers you would like to add to your existing configuration as a part of this plan." Below this text are four checkboxes: "Deployment Server", "Enterprise Server", "Data Server", and "HTML Server". The "Enterprise Server" checkbox is checked.

## Server Add Form

Machine Usage: 20 Enterprise Server Location: Corporate  
 Machine Name: jdeprodapp Primary User: jde920  
 Description: Enterprise Server for JDE App  
 Release: E920  
 Host Type: 35 Linux  
 Port Number: 6017  
 Logical Machine Name: JDEPRODAPP  
 Database Type: ORACLE  
 Server Map Data Source: jdeprodapp - 920 Server Map  
 Installation Path: Ju01jdedwardsppack/e920  
 Deployment Server Name: DEN06KSW

Enterprise Server Revisions Form

Data Source Name: jdeprodapp - 920 Server Map  
 Data Source Use: DB Local Data Source  
 Data Source Type: ORACLE  JDBNET Data Source?  
 Data Class: H Host Server Map  
 Platform: LINUX Linux  
 Database Server Name: jdedbs  
 Object Owner ID: jw920  
 Database Name (TNS): jdeorcl

Data Source Revisions Form

22. Deploy an already built server package to these virtual servers.

Build a server package to deploy on many virtual servers. This can be achieved by setting ***docompression=1*** in JDE.INI of the build machine.

23. Activate the respective OCM of the System and Server Map data sources for JDEPRODAPP and JDEPRODUBE.
24. Add the Security Server as JDEPRODAPP in the .ini files of all HTML Servers (JAS).

Internally with the setting previously performed as described in this technical brief, all of the WLS 1 set of HTML Servers will alias JDEPRODAPP as JDEPRODAPP1 and all of the WLS2 set of HTML Servers will alias JDEPRODAPP as JDEPRODAPP2.

## Disaster Recovery (DR) for Oracle JD Edwards EnterpriseOne on the Oracle Public Cloud

When implemented in the Oracle Public Cloud, data is replicated on a regular basis between production and standby sites for all layers:

- » Database tier – Activate Data Guard
- » Application layer – Using Rsync
- » Web layer – Using Rsync

For example, a customer might want disaster recovery implemented with these expectations:

- » Recovery Time Objective (RTO) — 2 hours to recover critical applications and 10 hours to recover non-critical applications
- » Recovery Point Objective (RPO) — 15 minutes for recovery of the complete database

## Active Data Guard

DR at the database layer is achieved using Active Data Guard, which is described in this documentation:

<https://docs.oracle.com/en/cloud/paas/database-dbaas-cloud/csdbi/use-data-guard-this-service.html>

## Setup and Configuration of RSYNC

For JD Edwards EnterpriseOne app and HTML Servers, follow these steps to perform rsync to replicate the required files from PD to DR on a regular basis:

25. Identify the required “jde” related folder under /u01 that needs to be backed up.
26. Verify that the hostname at PD and DR are the same.
27. Generate a SSH key pair for oracle/jde920 users on both machine 1 and machine 2.
28. Copy the public key contents of the oracle/jde920 user of machine 1 to machine 2 into the /home/oracle/.ssh/authorized\_keys file.
29. Copy public key contents of the oracle/jde920 user of machine 2 to machine 1 into the /home/oracle/.ssh/authorized\_keys file.
30. Create shell scripts with rsync commands to copy required changes in the file, folder, and so on from the PD machine to the DR machine using this command syntax:  

```
#!/bin/sh
sudo -u oracle rsync -pavzr oracle@141.145.114.149:/u01/jde_home/ /u01/jde_home
exit
```
31. Use the Crontab utility to automatically run this script every 10 to 15 minutes to synchronize with your OD and DR servers.

## HA and DR of Media Objects and Print Queue Directory

- » Set up the NFS server on the DBCS instance and JD Edwards EnterpriseOne Enterprise Server and HTML Server to act as the NFS Client.
- » Specify that all MOQUE and Print Queue files are to be stored in a file system on DBCS. We use DBCS instance because later we can move all files from DBCS file structure to DBFS. In case of disaster, Active Data Guard takes care of replicating DBFS on DR site and hence provision high availability of the file system. To move the files from file system to DBCS refer to [https://docs.oracle.com/cd/E18283\\_01/appdev.112/e18294/adlob\\_fs.htm](https://docs.oracle.com/cd/E18283_01/appdev.112/e18294/adlob_fs.htm)
- » You can use the below section to set up the NFS server on the DBCS and NFS client on the JD Edwards Enterprise Server.
- » Make sure you don't use NFS for Enterprise Server folders (<pathcode>/spec)

### Setting up a Print Queue Directory on a DBCS Instance File Structure

For this scenario:

**DBCS instance name: jdedbs92a (will act as NFS Server)**  
**Enterprise Server name: ent920 (will act as NFS Client)**

#### Setting up the NFS Server on the DBCS Machine jdedbs92a

32. Log in to the DBCS machine jdedbs92a as 'opc'.
33. Search for the NFS package:

**`$sudo yum list | grep nfs`**

```
[opc@jdedbs92a u01]$ sudo yum list | grep nfs
nfs-utils.x86_64 1:1.2.3-54.el6 @public_ol6_latest
nfs-utils-lib.x86_64 1.1.5-9.el6 @public_ol6_latest
nfs4-acl-tools.x86_64 0.3.3-6.el6 @public_ol6_latest
nfs-utils.x86_64 1:1.2.3-70.0.1.el6_8.1 public_ol6_latest
nfs-utils-lib.i686 1.1.5-11.el6 public_ol6_latest
nfs-utils-lib.x86_64 1.1.5-11.el6 public_ol6_latest
nfs-utils-lib-devel.i686 1.1.5-11.el6 public_ol6_latest
nfs-utils-lib-devel.x86_64 1.1.5-11.el6 public_ol6_latest
nfs4-acl-tools.x86_64 0.3.3-8.el6 public_ol6_latest
pcp-pmda-nfsclient.x86_64 3.10.9-6.el6 public_ol6_latest
sblim-cmpi-nfsv3.i686 1.1.1-1.el6 public_ol6_latest
sblim-cmpi-nfsv3.x86_64 1.1.1-1.el6 public_ol6_latest
sblim-cmpi-nfsv3-test.x86_64 1.1.1-1.el6 public_ol6_latest
sblim-cmpi-nfsv4.i686 1.1.0-1.el6 public_ol6_latest
sblim-cmpi-nfsv4.x86_64 1.1.0-1.el6 public_ol6_latest
sblim-cmpi-nfsv4-test.x86_64 1.1.0-1.el6 public_ol6_latest
```

34. Install the nfs util and nfs util lib packages:
  - `$sudo yum install nfs-utils.x86_64`**
  - `$sudo yum install nfs-utils-lib.x86_64`**
35. Start the rpcbind service:
  - `$sudo /etc/init.d/rpcbind start`**
36. Start the nfs service:
  - `$sudo /etc/init.d/nfs start`**

```

[opc@jdedbs92a u01]$ sudo /etc/init.d/rpcbind start
Starting rpcbind: [OK]
[opc@jdedbs92a u01]$ sudo /etc/init.d/nfs start
Starting NFS services: [OK]
Starting NFS quotas: [OK]
Starting NFS mountd: [OK]
Starting NFS daemon: [OK]
Starting RPC idmapd: [OK]

```

37. Edit the **/etc/exports** file as root by adding the following line:

```
/u01/jdeprint ent920(rw, sync, no_root_squash)
```

38. Add permissions to the jdeprint file:

```
$sudo chmod -R 777 /u01/jdeprint
```

39. Restart the nfs service:

```
$sudo /etc/init.d/nfs restart
```

#### Setting up the NFS Client on Enterprise Server ent920

40. Display the mount for jdedbs92a:

```
$showmount -e jdedbs92a
```

```

[opc@ent920 jdedwards]$ showmount -e jdedbs92a
Export list for jdedbs92a:
/u01/jdeprint ent920.compute-acmeccs.oraclecloud.internal
[opc@ent920 jdedwards]$

```

41. Switch to jde920 user and create a folder /u01/jdedwards and give **chmod -R 777** permissions to it.

42. Switch user to opc and mount the shared Jde directory as following on the directory created above:

```
$sudo mount -t nfs jdedbs92a:/u01/jdeprint /u01/jdedwards
```

43. Stop the Enterprise Server services from the Server Manager Console.

44. Edit the Print Queue Directory section of enterprise jde.ini with value /u01/jdedwards as follows and click the **Apply** button.



JD Edwards EnterpriseOne Server Manager

45. Restart the Enterprise Server from the Server Manager Console.
46. Log in to the HTML Server (JAS) and submit a UBE. All the PDF output will be generated in /u01/jdedwards on ent920 and automatically goes to the mount /u01/jdeprint on the DBCS machine. The following sample screenshots validate this condition on each machine.
  - a. On the DBCS jdedbs92a machine:

```
root@jdedbs92a ~]# cd /u01
root@jdedbs92a u01]# ls
app depo download ftp jde920 jdeprint lost+found vsftpd_old.conf
root@jdedbs92a u01]# cd jdeprint/
root@jdedbs92a jdeprint]# ls
PrintQueue test
root@jdedbs92a jdeprint]# cd PrintQueue/
root@jdedbs92a PrintQueue]# ls
R0008P_XJDE0001_6_PDF R0008P_XJDE0001_6_PDF.jdedebug.log R0008P_XJDE0001_6_PDF.jde.log
root@jdedbs92a PrintQueue]#
```

- b. On the ent920 machine:

```
jde920@ent920:/u01/jdedwards/PrintQueue
[jde920@ent920 ~]$ cd /u01/jdedwards/
[jde920@ent920 jdedwards]$ ls
PrintQueue test
[jde920@ent920 jdedwards]$ cd PrintQueue/
[jde920@ent920 PrintQueue]$ ls
R0008P_XJDE0001_6_PDF R0008P_XJDE0001_6_PDF.jdedebug.log R0008P_XJDE0001_6_PDF.jde.log
[jde920@ent920 PrintQueue]$
```

## Appendix A: Script

The following script is provided as is and with no guarantee. It is provided as an example of what can be done by customers to modify the RunOneWorld.sh script to work within a Clusterware environment. This change will prevent the Server Manager from starting or stopping EnterpriseOne in conflict with the Clusterware.

This script code would be inserted into the RunOneWorld.sh script on an EnterpriseOne Server installed on a shared disk within a Clusterware managed environment.

```
#####
Transparent High Availability
#####

AWK=/usr/bin/awk

find CRS_HOME
CRS_HOME=`$AWK /crs_home/ /etc/oracle/olr.loc | $AWK -F"=" '{print $2}'`
if [! X$CRS_HOME = "X"]
then
start resource using crsctl if DESCRIPTION attribute is not set to JDE_THA
 if ["X$CRS_DESCRIPTION" = "X"] && ["$_CRS_DESCRIPTION" = "JDE_THA"]
 then
 HOSTNAME=`/bin/hostname`
 PRINTF=/usr/bin/printf
 ECHO=/bin/echo

start resource
 MSGOUT=`$CRS_HOME/bin/crsctl start res -w "(TYPE st xag.jde910.type) AND
(JDE_HOME = ${JDE_ROOT}/)" -n $HOSTNAME -env "DESCRIPTION=JDE_THA"`

print JDE output messages
 $PRINTF %s "MSGOUT" | while read -r line
 do
 CRS_MSG=`$ECHO $line | grep CRS-`
 if ["X$CRS_MSG" = "X"]
 then
 print $line
 exit
 fi
 done
 fi
fi

#####
Transparent High Availability - END
```



#####

This code should be inserted before the lines:

```
print " Starting jdenet_n..." >> $LOGFILE
cd $SYSTEM/bin32
$SYSTEM/bin32/jdenet_n > $SYSTEM/bin32/jdenet_n.log 2>&1 &
```

This is at line number 292.



**Oracle Corporation, World Headquarters**  
500 Oracle Parkway  
Redwood Shores, CA94065, USA

**Worldwide Inquiries**  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200



---

CONNE C T W I T H US

-  [blogs.oracle.com/oracle](http://blogs.oracle.com/oracle)
-  [facebook.com/oracle](http://facebook.com/oracle)
-  [twitter.com/oracle](http://twitter.com/oracle)
-  [oracle.com](http://oracle.com)

**Hardware and Software, Engineered to Work Together**

Copyright © 2014, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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. 0515



Oracle is committed to developing practices and products that help protect the environment

---