Implementation of SAP® Applications Using Oracle® Maximum Availability Architecture Best Practices

With Oracle RAC, Oracle Data Guard, Oracle Recovery Manager & Oracle GRID Control

Mahesh Pakala

Global Infrastructure Consulting Services
Dell Inc.

Version 01

January 2009

CONTENTS

Executive Summary

This document contains the suggested Oracle 10g Release 2 EE Database Real Application Cluster (RAC) server architecture and implementation of a SAP7.0 Landscape for a large commercial customer. The cluster storage is configured using OCFS2 and Oracle Data Guard is being implemented to provide disaster recovery capabilities. Oracle GRID Control is being used to manage the total landscape. The server infrastructure consisted of Dell PowerEdge R900 and 2950 servers all running 64bit RHEL Linux AS 4.6. Storage is provisioned from an EMC Storage Array. The implementation has been influenced by Oracle's documented Maximum Availability Architecture (MAA) best practices for implementing highly available systems.

SAP Applications have numerous configuration options that can be chosen to suit particular business scenarios, hardware capability, and availability requirements. This document only describes how to configure the Oracle RAC to support a SAP 7 environment that will utilize the physical standby feature of Oracle Data Guard.

This document does not include the High Availability setup for the SAP Application server components SCS(Central Service) – Batch, Dialog, Update and Spool. The HA setup for the Enqueue and Message server is a different exercise and achieved by using the Oracle CRS and Oracle SAPCTL.

A number of conventions are used in this document:

Convention	Meaning
Application Tier	Machines running SAP Applications and other servers. Also called middle tier.
Database Tier	Machines running SAP Applications database.
Production System	Primary SAP Database system, which will used to create a standby system.
Standby System	SAP Applications system created as a copy of the production system.
oracle	User account that owns the Oracle CRS and Oracle cluster file system (database CRS_HOME and OCFS & GRID files).
Ora<DBSID>	Oracle Database and relevant files
ORACLE_BASE	/oracle/<DBSID>
NLS_LANG	As set in the old system
< >	Text enclosed in angle brackets represents a variable. Substitute a value for the variable text. Do not type the angle brackets.
****	The backslash character is entered at the end of a command line to indicate continuation of the command on the next line.
SAP Solution Manager	SAP Solution Manager 7.0
SAP System	SAP Solution Manager 7.0 system
ABAP + Java system or system based on AS ABAP and AS Java	Two stacks (ABAP + Java) of solution manager

Variable	Description
<SAPSID>	SAP system ID in uppercase letters
<sapsid>	SAP system ID in lowercase letters
<DBSID>	Database ID in uppercase letters
<dbsid>	Database ID in lowercase letters
<host_name>	Name of the corresponding host
<user_home>	Home directory of the user performance
<INSTDIR>	Installation directory for the SAP System
ORACLE_HOME	/oracle/<dbsid>/1020_64
ORACLE_SID	<dbsid>
ora<dbsid>	User account that owns the database file system (database ORACLE_HOME and files).

Users and Groups that need to be created for both Application servers and DB servers

User	Primary Group
<sapsid>adm	sapsys, oper, dba, sapinst
ora<dbsid>	dba, oper, sapinst
oracle	dba, oper ,root
Groups	**Members**
sapsys	<sapsid>adm
oper	<sapsid>adm, ora<dbsid>
dba	<sapsid>adm, ora<dbsid>
sapinst	<sapsid>adm, ora<dbsid>

Phase 1: Initial Planning

The reader of this document should be familiar with the Oracle10g database server, and have at least a basic knowledge of Real Application Clusters (RAC) and Data Guard configurations.

1.1 High Availability Terminology - Overview

It is important to understand the terminology used in a high availability environment. Key terms include the following.

Real Application Clusters (RAC) is an Oracle database technology that allows multiple nodes in cluster to access one database for greater application scalability or to increase speedup of large batch jobs through parallelism. A RAC environment also offers resilience if one or more machines become temporarily unavailable as a result of planned or unplanned downtime.

The advantages of using Real Application Clusters include:

- High availability
- Rapid and automatic recovery from node failures or an instance crash
- Increased scalability

Standby Database assists with disaster recovery by providing a completely automated framework to maintain one or more transactionally-consistent copies of the primary database. Changes can be transmitted from the primary database to the standby databases in a synchronous manner, which avoids any data loss, or in an asynchronous manner, which minimizes any performance impact on the production system. The standby database technology includes an automated framework to switch over to the standby system in the event of a physical disaster, data corruption, or planned maintenance at the production (primary) site.

A standby database can be either a physical standby database or a logical standby database. The only SAP supported database is Physical Standby Database.

- **Physical standby database**
  Provides a physically identical copy of the primary database, with on-disk database structures that are identical to the primary database on a block-for-block basis. The database schemas, including indexes, are the same. A physical standby database is kept synchronized with the primary database by recovering the redo data received from the primary database.
- **Logical standby database**
  Contains the same logical information as the production database, although the physical organization and structure of the data can be different. It is kept synchronized with the primary database by transforming the data in the redo logs received from the primary database into SQL statements and then executing the SQL statements on the standby database. A logical standby database can be used for other business purposes in addition to disaster recovery requirements. This allows users to access a logical standby database for queries and reporting purposes at any time. Thus, a logical standby database can be used concurrently for data protection and reporting.

Caution: At present, only physical standby databases are supported with SAP Applications; logical standby databases are not supported.

Oracle Data Guard is a set of services that create, manage, and monitor one or more standby databases to enable a production database to survive disasters and data corruption. If the production database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch a standby database to the production role, minimizing the downtime.

The advantages of using Data Guard include:

- Disaster protection and prevents data loss
- Maintains transactional consistent copies of primary database
- Protects against data corruption and user errors
- Does not require expensive and complex hardware or software mirroring

Data Guard offers three modes of data protection:

- **Maximum Protection**
  This mode offers the highest level of data protection. Data is synchronously transmitted to the standby database from the primary database, and transactions are not committed on the primary database unless the redo data is available on at least one standby database configured in this mode. If the last standby database configured in this mode becomes unavailable, processing stops on the primary database. This mode guarantees no data loss.
- **Maximum Availability**
  This mode is similar to the maximum protection mode, including no data loss. However, if a standby database becomes unavailable (for example, due to network connectivity problems), processing continues on the primary database. When the fault is corrected, the standby database is resynchronized with the primary database. If there is a need to fail over before the standby database is resynchronized, some data may be lost.
- **Maximum Performance**
  This mode offers slightly less data protection on the primary database, but higher performance than maximum availability mode. In this mode, as the primary database processes transactions, redo data is asynchronously shipped to the standby database. The commit operation on the primary database does not wait for the standby database to acknowledge receipt of redo data before completing write operations on the primary database. If any standby destination becomes unavailable, processing continues on the primary database, and there is little effect on primary database performance.

This note describes the use of Data Guard Maximum Performance protection mode with a physical standby database in this note. For more details on High Availability architectures, refer to the *Maximum Availability Architecture* white paper on the Oracle Technology Network. The system architects need to choose which option is the best suited for their environment.

Phase 2: Design Considerations

1.2 Requirements for Secondary Site

The site for your standby environment should:

- Be physically separate from the primary site, to protect against local and regional disasters. In terms of distance, it would be a best practice to locate a standby facility outside of any regional threat zone (hurricane, tornado, flood, etc) of the production facility is typical practice for an organization to locate its standby data center in a different place (may be a different city or state) from the production data center.
- Validate that you have sufficient system resources to meet your application SLAs in case of periodic role transitions or a disaster that requires Data Guard failover operation.
- Validate that there sufficient network redundancy and bandwidth on both sites and the network bandwidth can accommodate peak redo rates and user traffic.

For more detailed information, refer to HA Best Practices book and various MAA DG best practices papers.

Phase 3: Example Environment

SAP Netweaver 7.0 with an Oracle 10.2.0.1 database server was installed using '*Typical Mode*' Install. The database was then upgraded to 10.2.0.4. The standby system will use the same version of the same products as installed on the primary system. Oracle GRID Control 10.2.0.4 was installed on an independent server to manage the environment.

1.3 Software and Hardware Configuration

The following versions of software and hardware were used in this installation. The architecture described in this note is a sample configuration. For more details regarding supported architectures, refer to Sapnote - 527843

Software Component	Version
SAP Application	SAP Netweaver Version 7.0
Oracle10g Release2 database server	Release 10.2.0.4 (Production release)
Oracle Cluster CRS for RAC	Release 10.2.0.4 (Production release)
Oracle GRID Control	Release 10.2.0.4 (Production release)
Linux Operating System	RedHat AS4 U6 Enterprise 64 bit
OCFS2	ocfs2-2.6.9-55.0.12.ELsmp-1.2.9-1.el4.x86_64.rpm
PowerPath	EMCPowerpath.Linux-5.1.0.194.rhel.x86_94

1.4 Prerequisite Execution Tasks

1.) Step – I
- Setup the Cluster Hardware and Software on Production and Standby Sites
- Setup the Operating system and Environment
- Setup the shared storage and network
- Understanding the SAP architecture requirements
- Understanding the implementation challenges

2.) Step –II (Database Servers)
- Install Oracle Cluster File System(OSFS2)
- Install CRS and Oracle in separate homes
- Complete NIC teaming, Multipath and Storage setup

3.) Step – III
- Install Oracle GRID control
- Upgrade to 10.2.0.4
- Setup agents in all the relevant servers

4.) Step – IV (Application Server)
- Install the Linux Operating system on the app servers
- Setup the SAP Java environment

- Setup the storage and hardware configuration requirements

5.) Step – V (Installation)

- Install the SAP application using the Solution manager which also was used to install Oracle Server.
- After the SAP installation perform the Oracle DB upgrade and modifications to RAC
- Perform the Oracle Data Guard Setup and Oracle GRID Control
- Modify tnsnames and listener.ora files for SAP Application server

Warning: System administrators are strongly advised to take a complete backup of the environment before executing these procedures, and to take additional backups between stages of this migration. The procedures in this document should be tested in non-production environments before using them in production environments. All users must log off your system while the migration procedure is being carried out.

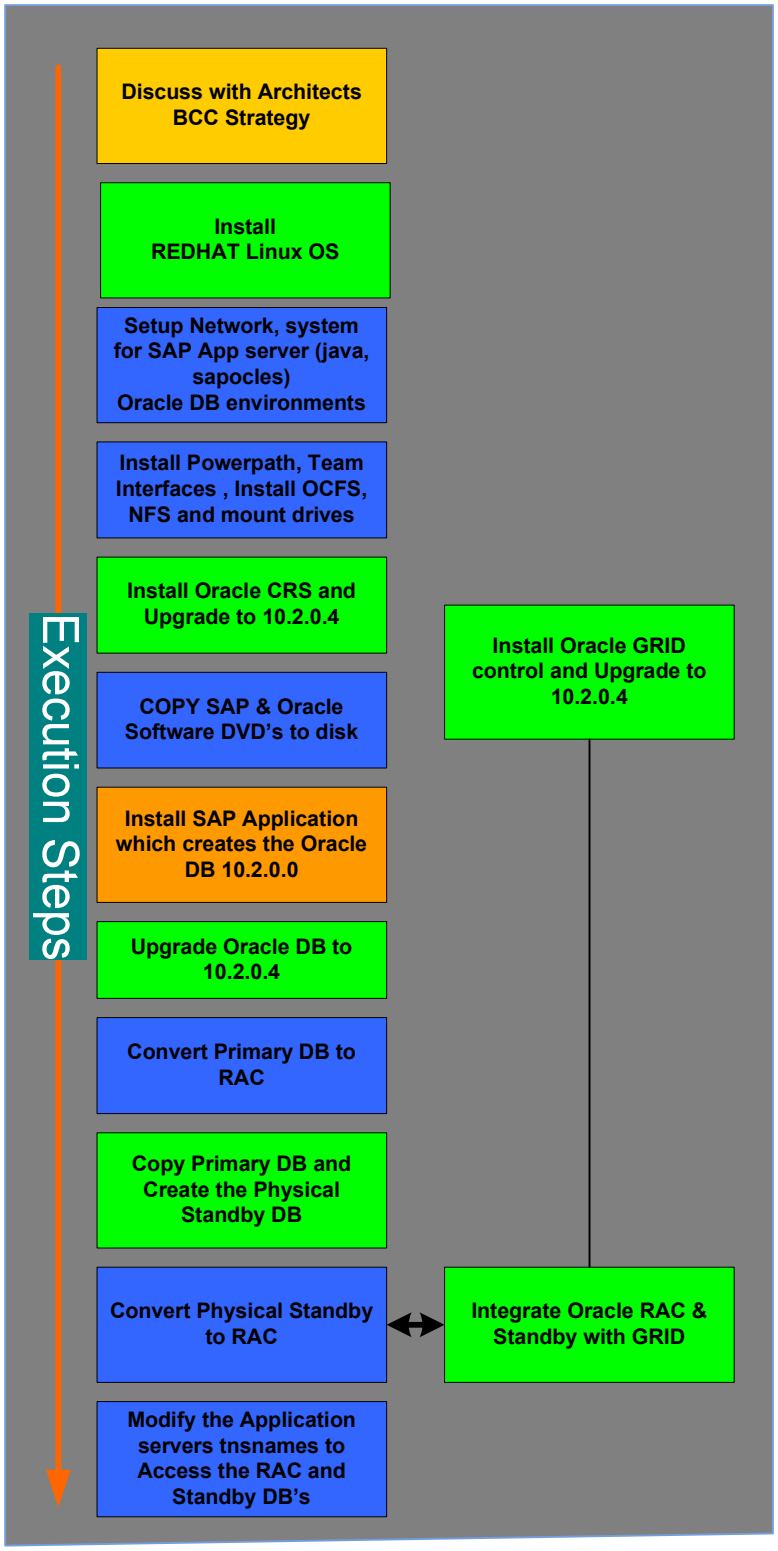
Execution Steps
Discuss with Architects BCC Strategy
Install REDHAT Linux OS
Setup Network, system for SAP App server (java, sapocles) Oracle DB environments
Install Powerpath, Team Interfaces , Install OCFS, NFS and mount drives
Install Oracle CRS and Upgrade to 10.2.0.4
COPY SAP & Oracle Software DVD's to disk
Install SAP Application which creates the Oracle DB 10.2.0.0
Upgrade Oracle DB to 10.2.0.4
Convert Primary DB to RAC
Copy Primary DB and Create the Physical Standby DB
Convert Physical Standby to RAC
Modify the Application servers tnsnames to Access the RAC and Standby DB's
Install Oracle GRID control and Upgrade to 10.2.0.4
Integrate Oracle RAC & Standby with GRID

1.4.1System Architecture for setup

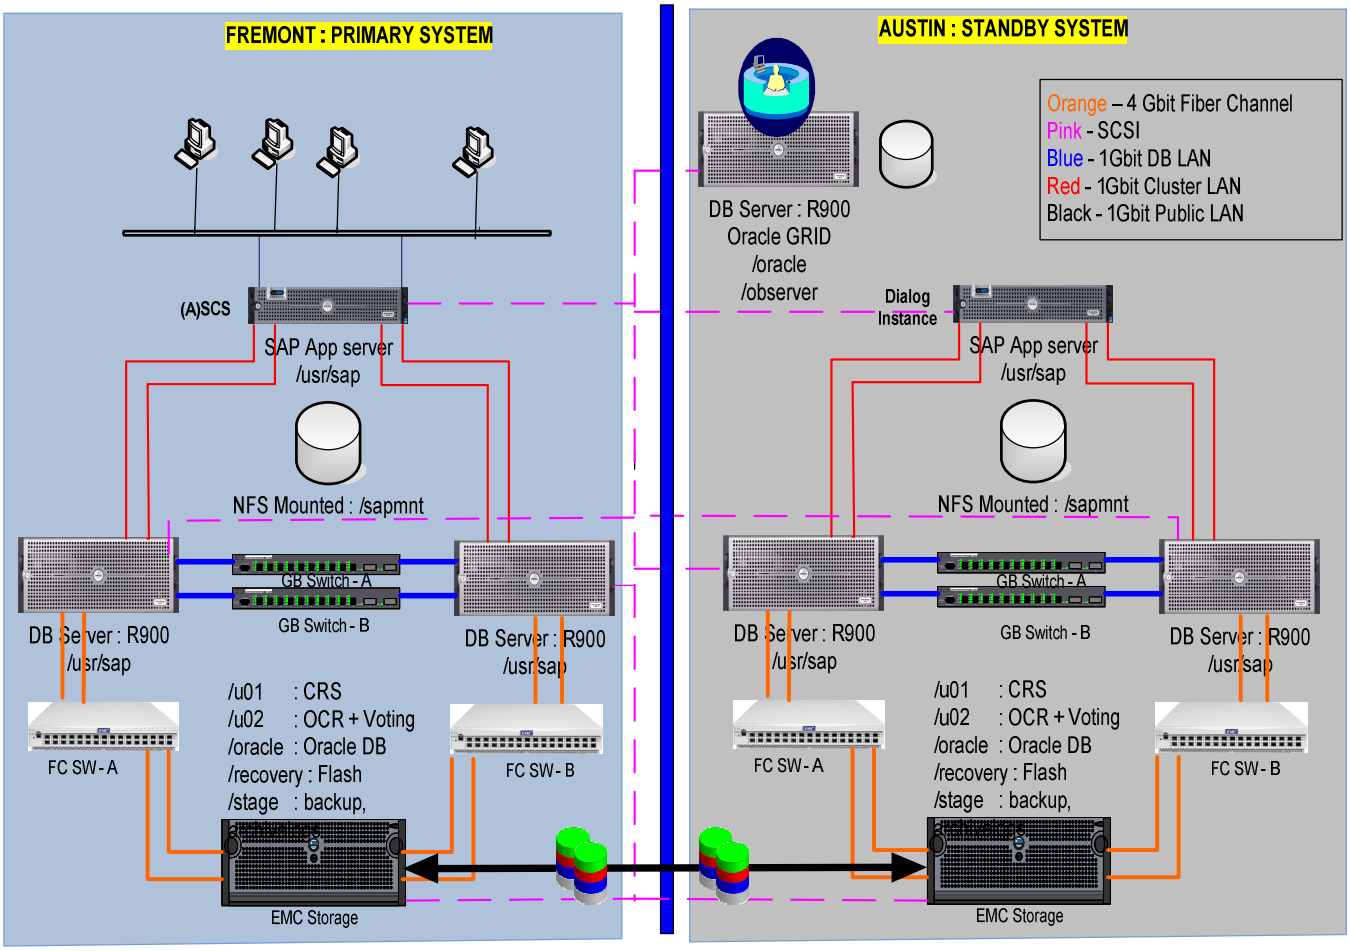

.

Task #1 REDHAT Linux OS Installation

1. Boot the server using the "Install CD".
    - You may need to change the BIOS settings to allow booting from the CD and to adjust other functionality.

2. The boot screen appears with the `boot:` prompt at the bottom of the screen.
    - Press **Enter** to continue with a graphical install on the console.
    - The installer scans your hardware, briefly displays the Red Hat splash screen, and then begins a series of screen prompts.

3. The "CD Found" pop up window appears, press **TAB** to highlight **SKIP** and press **ENTER** to continue. Booting continues and the graphical installation environment is started.

4. Splash Screen
    - Click **NEXT** to continue.

5. Language Selection Screen
    - Accept the default (English), click **NEXT** to continue.

6. Keyboard Configuration Screen
    - Accept the default (U.S. English), click **NEXT** to continue.
    - Installation Number Pop Up
        - Enter **Installation Number**

        - or –

        - Click **Skip entering Installation Number**
        - Click **OK**
        - Skip Pop Up Screen
            - Click **Skip**

7. Disk Partitioning Setup Screen
    - A thorough treatment of disk partitioning is beyond the scope of this guide, which assumes that you are familiar with disk partitioning methods.
    - Click on **Review and modify partitioning layout**.

    The partitions were created in the local drives –
    - / - 45 GB
    - /home - 8GB
    - /swap - 30 GB
    - /opt - 10
    - /tmp - 5GB

8. Boot Loader Configuration Screen
    - Accept the defaults by clicking **NEXT** to continue.

9. Network Configuration Screen
    - It is necessary to configure database servers with a static IP address. To do so, select the desired Ethernet interface and click on **Edit** .

- A pop-up window appears. Check the **Manual configuration** box, and enter the IP Address and Netmask for the server. Deselect the **Enable IPv6 support**. Click **OK**.
- In the Hostname box, select **manually** and enter the hostname.
- In the Miscellaneous Settings box, enter the remaining network settings.

10. Time Zone Selection Screen
    - Choose the time settings that are appropriate for your area.
    - Setting the system clock to UTC is a best practice for database servers. To do so, click on **System clock uses UTC**.
    - Click on **NEXT** to continue.

11. Set Root Password Screen
    - Enter a password for root, and enter it again to confirm.
    - Click **NEXT** to continue.

12. Package Installation Defaults Screen
    - Select **Customize now**.
    - Click **NEXT** to continue.

13. Package Group Selection Screen
    - Select only the package sets shown here and leave all others unselected.
    - Desktop
        - Gnome
    - Applications
        - Editors
        - Graphical Internet
    - Development
        - Development Libraries
        - Development Tools
        - GNOME Software Development
        - Legacy Software Development
            - Click **Optional packages**
            - gtk+
        - X Software Development
            - Click Optional packages
            - Select openmotif-devel
        -
    - Servers
        - Legacy Network Server
            - Click **Optional packages**
            - rsh-server
        - Network Servers
        - Printing Support (optional)
        - Server Configuration Tools
    - Base System
        - Administration Tools
        - Base
            - Click **Optional packages**
            - iscsi-initiator-utils
        - Java
        - Legacy Software Support
            - Click **Optional packages**

        - compat-db
        - openmotif22
        - qt4
      - System Tools
        - Click **Optional packages**
        - gnome-nettool
        - nmap-frontend
        - sysstat
        - wireshark-gnome
      - X Window System
    - Click on **Next** to continue.

14. About to Install Screen
    - Click **NEXT** to begin the installation.

15. Installing Packages Progress Screen
    - Software is now installed to the hard disk.

16. Congratulations Screen
    - Remove the installation media from the system, and click on **Reboot**.

The server now reboots.

1.5 First Boot

The first boot program runs and completes the installation process.

1. Welcome Screen
   - The system automatically reboots and presents a new welcome screen.
   - Click on **Next** to continue.

2. License Agreement Screen
   - Read the license agreement. If you agree to the terms, select **Yes, I agree to the License Agreement.**
   - Click on **Next** to continue.

3. Date and Time Screen
   - Set the Date and Time.
   - Do NOT set up NTP at this time.
   - Click on **Next** to continue.

4. Display Screen
   - Accept the defaults.
   - Click **Next** to continue.

5. Red Hat Login Screen
   - Click on **Tell me why I need to register and provide a Red Hat login**.
   - Click on **Next** to continue.
   - Click on **I can not complete registration at this time. Remind me later**.
   - Click on **Next** to continue.

6. System User Screen
   - Create an account for yourself.

- Do not create an account for oracle at this time. Creating the oracle account is covered later in this section.
- Click on **Next** to continue.

7. Additional CDs Screen
   - Click on **Next** to continue.

8. Finish Setup Screen
   - Click on **Next** to continue.

Congratulations, the core, default Red Hat operating system is now installed. Login now using the root account and complete the tasks in the following section. Repeat this step for all nodes in the cluster.

1.6 Linux Environment Configuration

Run Level

The default run level for Red Hat is to come up in multi-user and X11 mode. It is best to not have the full X11 desktop running and use those precious CPU cycles for real work; like running database queries.

Edit the /etc/inittab by change the following line:

```
id:5:initdefault:
```

to read,

```
id:3:initdefault:
```

When rebooted, the system will come up into multi-user mode and not invoke the X11 environment

Task#2 Network Configuration

From the desktop, launch the **Applications>System Settings>Network** tool. Adjust the network interface(s) as needed for your network. Test that network connectivity is working (ping, dig, nslookup, and so forth).

You can configure all the NIC's using the files in /etc/sysconfig/network-scripts/ifcfg-ethXX

Use eth0 for the public interface for public and VIP traffic; this is an Oracle requirement. Use eth1 for private interconnect traffic. When setting the hostname; use a fully-qualified name; this too is an Oracle requirement.

1.7 Edit /etc/hosts

The hosts file should always – as rule of thumb for Oracle – resolve all Oracle-related host names rather than relying solely on DNS. DNS is a known single point of failure because of outages or user-error.

So, make sure that each local server IP address is listed and that both the short and fully-qualified hostnames are listed in /etc/hosts. It is also wise to list other servers needed by the database (e.g. RMAN catalog, DB links, standbys, RAC nodes) as well.

For a dual interface, two (2) node RAC configurations, the hosts file may look something as follows:

```
# Do not remove the following line, or various programs
# that require network functionality will fail.

##PRIMARY###
10.67.40.68 fremont01.dell.com                  fremont01
10.67.40.69 fremont01-sap.dell.com              fremont01-sap
10.67.58.40 fremont01-priv

10.67.40.70 fremont02.dell.com                  fremont02
10.67.40.71 fremont02-sap.dell.com              fremont02-sap
10.67.58.41 fremont02-priv

####STANDBY####

10.67.40.74 austin01.dell.com                   austin01
10.67.40.75 austin01-sap.dell.com               austin01-sap
10.67.58.42 austin01-priv

10.67.40.76 austin02.dell.com                   austin02
10.67.40.77 austin02-sap.dell.com               austin02-sap
10.67.58.43 austin02-priv

####GRID##
10.67.40.78 dallas.dell.com                     dallas
```

```
10.67.40.79 dallas-sap.dell.com                 dallas-sap
10.67.58.44 dallas-priv

127.0.0.1   localhost.localdomain   localhost
```

Please also make sure(SAP requirement) that when you type hostname you get the following results for all servers -

```
#hostname
fremont01

#hostname -f
fremont01.dell.com
```

1.8 Teaming NICS(BONDING)

```
Check if the all the Ethernet interfaces are up and connected -

[root@fremont01 bin]#
[root@fremont01 bin]# mii-tool
eth0: negotiated 100baseTx-FD, link ok
SIOCGMIIPHY on 'eth1' failed: Resource temporarily unavailable
SIOCGMIIPHY on 'eth2' failed: Resource temporarily unavailable
SIOCGMIIPHY on 'eth3' failed: Resource temporarily unavailable
eth4: negotiated 100baseTx-FD, link ok
eth5: negotiated 100baseTx-FD, link ok
eth6: negotiated 100baseTx-FD, link ok
eth7: no link
[root@fremont01 bin]#
```

Next most important task after the OS has been installed and environment setup is to bond the NICS or called teaming for providing higher availability for the interfaces.

```
System Configuration(/etc/modprobe.conf):
alias eth0 bnx2
alias eth1 bnx2
alias eth2 bnx2
alias scsi_hostadapter megaraid_sas
alias scsi_hostadapter1 ata_piix
alias scsi_hostadapter2 qla2xxx

alias scsi_hostadapter3 qla2400
###BEGINPP
include /etc/modprobe.conf.pp
###ENDPP
alias peth0 bnx2
alias peth1 bnx2
#### Interconnect NIC bonding ###
alias bond0 bonding
options bond0 mode=1 miimon=100
alias bond1 bonding
options bond1 mode=1 miimon=100
```

Modify the Network Configuration files-

Cat /ifcfg-bond0

```
DEVICE=bond1
IPADDR=10.67.40.76
ONBOOT=yes
TYPE=Ethernet
BOOTPROTO=none
NETMASK=255.255.255.0
USERCTL=no
```

Cat /ifcfg-eth0

```
## Private Network Interconnect Bonding

DEVICE=eth0
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
MASTER=bond1
SLAVE=yes
BOOTPROTO=static
```

Cat /ifcfg-eth6

```
## Private Network Interconnect Bonding

DEVICE=eth6
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
MASTER=bond1
SLAVE=yes
BOOTPROTO=static
```

Cat /ifcfg-bond0

```
DEVICE=bond1
IPADDR=10.67.40.76
ONBOOT=yes
TYPE=Ethernet
BOOTPROTO=none
NETMASK=255.255.255.0
USERCTL=no
```

Cat /ifcfg-eth1

```
## Private Network Interconnect Bonding

DEVICE=eth1
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
MASTER=bond1
SLAVE=yes
BOOTPROTO=static
```

Cat /ifcfg-eth4

```
## Private Network Interconnect Bonding

DEVICE=eth4
ONBOOT=yes
TYPE=Ethernet
USERCTL=no
MASTER=bond1
SLAVE=yes
BOOTPROTO=static
```

After the Oracle clusterware is installed you should see the below information -

```
[root@fremont01 bin]# ./oifcfg getif
bond1  172.24.22.0  global  public
bond0  192.168.10.0  global  cluster_interconnect
To confirm the bonding is setup correctly please verify -(Sample)
[root@fremont01 ~]# ping -I bond0 192.168.10.52
PING 192.168.10.52 (192.168.10.52) from 192.168.10.52 bond0: 56(84) bytes
of data.
64 bytes from 192.168.10.52: icmp_seq=1 ttl=64 time=0.038 ms
64 bytes from 192.168.10.52: icmp_seq=2 ttl=64 time=0.037 ms
64 bytes from 192.168.10.52: icmp_seq=3 ttl=64 time=0.041 ms
64 bytes from 192.168.10.52: icmp_seq=4 ttl=64 time=0.027 ms
64 bytes from 192.168.10.52: icmp_seq=5 ttl=64 time=0.024 ms
64 bytes from 192.168.10.52: icmp_seq=6 ttl=64 time=0.031 ms
64 bytes from 192.168.10.52: icmp_seq=7 ttl=64 time=0.045 ms
64 bytes from 192.168.10.52: icmp_seq=8 ttl=64 time=0.025 ms
64 bytes from 192.168.10.52: icmp_seq=9 ttl=64 time=0.030 ms
64 bytes from 192.168.10.52: icmp_seq=10 ttl=64 time=0.020 ms

--- 192.168.10.52 ping statistics ---
10 packets transmitted, 10 received, 0% packet loss, time 9001ms
rtt min/avg/max/mdev = 0.020/0.031/0.045/0.010 ms
[root@sage ~]#
```

```
Cat /proc/net/bonding/bond0

Ethernet Channel Bonding Driver: v3.1.2 (January 20, 2007)

Bonding Mode: fault-tolerance (active-backup)
Primary Slave: None
Currently Active Slave: eth1
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

Slave Interface: eth1
MII Status: up
Link Failure Count: 1
Permanent HW addr: 00:1e:4f:37:a7:be

Slave Interface: eth2
MII Status: up
```

```
Link Failure Count: 1
Permanent HW addr: 00:10:18:37:44:b2
```

1.9 EMC PowerPath Installation

Download the latest powerpath from EMC Powerlink website based on the certification matrix. Install the Powerpath package and the Naviagent.

#Rpm –ivh <package_name.rpm>

After installation register the package -

#emcpreg *install or add*

Use "emcpadm" to change the LUN order incase its different across nodes. It is stored in the /etc/emcp_devices.dat and idx files

#emcpadm rename -s emcpowera -t emcpowerq

To Display PowerPath pseudo device names in use:

#emcpadm getused

Verify that emcpowerX powerpath devices exist.

#powermt *display dev=all*

#powermt *config*

Also do fdisk –l and check the directory /dev

Warning : Please note the following caveats if using EMC Powerpath

- We cannot use power path, if any SAP programs are running on the server. In other words an Oracle only database server can run with power path, but this may not be the case in a SAP ERP environment (Like using BR Tools instead of RMAN etc). Secondly, in case you plan to use sapctl and enqueue replication these SAP processes are located on the database nodes.
- It has to be a complete 3 –tier no portion of the SAP solution stack, including the Enqueue service may reside on a RH Linux server system, which is connected to the storage via PP (as opposed to MPIO).
- These restrictions are not there if you use RH Native MPIO, which is recommended by SAP.

1.10 NTP Configuration

It is ***critical*** that Oracle products have accurate and identical date and timestamps. Oracle RAC will fail if the time between nodes drifts more than a few seconds apart. Trace, alert and log files are difficult to use if the timestamps between database node instances and clients are too far apart.

Setup the ntpd daemon to automatically maintain time synchronization. Ask your system administrator for the name(s) of your company's ntp server(s) before performing this step. As the root user, perform the following:

```
/etc/init.d/ntpd stop
cd /etc
```

```
vi ntp.conf
```

Add the correct ntp server name entries for your network. Find the server section and add the server entries as follows:

```
server 0.pool.ntp.org
```

Now, manually test that ntp is working by using the following command:

```
ntpdate 0.pool.ntp.org
11 Oct 12:14:03 ntpdate[3909]: adjust time server 66.187.224.4 offset
0.025833 sec
```

If this command fails, troubleshoot and resolve as this would indicate a network connectivity problem. Once the manual command is working, start the ntp daemon using the following command:

```
/etc/init.d.ntp start
ntpd: Synchronizing with time server: [  OK  ]
Starting ntpd: [  OK  ]
```

Finally, make sure the NTP daemon starts automatically at every reboot by using the following command:

```
chkconfig ntpd on
```

TASK#3 SAP Oracle User Environment

1.11 Enable sudo

Using the visudo editor, uncomment the following line:

```
%wheel        ALL=(ALL)             ALL
```

1.12 A Word About UID's and GID's

It is a best practice to have a global list of group id (gid) and user id (uid) from which all servers share. In a clustered database this is a requirement or the nodes of the cluster can not communicate and will step on each other's file ownership and permissions. From the manageability aspect having the same uid/gid for each user simplifies managing access and permissions across many servers. Pick a set of UID and GID values and stick to using them globally.

1.13 Create the dba Group

The Oracle documentation goes to great lengths to discuss two groups: oinstall and dba. In the real world, oinstall is less than useful and can often get in the way of the dba getting things done. Create only the dba group and make oracle the only member of that group. Do not use the oinstall group at all.

As the root user:

```
groupadd -g 501 dba
```

1.14 Create the oracle User Account

Create the oracle user account only for installing the CRS. The command shown creates the $HOME in /opt/oracle.

The other accounts need not be created manually as the SAPINST will create them automatically based on the supplied SID. (This step is described in more detail later in the document)

The manual scripts are more required for creating of the standby servers.

##APP server

groupadd -g 500 sapinst

groupadd -g 503 sapsys

useradd -c "SAP System Administrator" -m -g "sapsys" -G "sapinst" -s "/bin/csh" -u 500 zlsadm

###DB Server

groupadd -g 502 sapinst

groupadd -g 503 sapsys

groupadd -g 504 oper

useradd -c "SAP System Administrator" -m -g "sapsys" -G "sapinst,dba,oper" -s "/bin/csh" -u 500 zlsadm

useradd -c "SAP Database Administrator" -M -d "/oracle/ZLS" -g "dba" -G "oper,sapinst" -s "/bin/csh" -u 503 orazls

1.15 Set the oracle User Account Password

The DBA's are ultimately responsible for hardening the oracle account password. But when delivering new machines, it is best to set the delivered default password to a known standardized delivery password.

As the root user:

```
passwd oracle
Changing password for oracle.
New Password: manager1
Reenter New Password: manager1
Password changed.
```

1.16 Create the ORACLE_BASE Directory

The ORACLE_BASE (typically /opt/oracle) is the "base" directory from which all Oracle 10g Database products, log files and inventory are stored. It is mandatory as per SAP Note 527843 to have shared ORACLE_HOME, shared CRS Home, Voting Disk, OCR and ORACLE home directory mounted using OCFS2 or a certified shared NAS NFS solution..

As root:

```
mkdir /oracle
chown ora<SID>:dba /u01/oracle
```

1.17 Create the CRS_BASE Directory

The CRS_BASE (typically /u01/oracleCRS) is the "base" directory from which all Oracle 10g ClusterWare products and logfiles are stored.

As root:

```
mkdir /u01/oracleCRS
chown oracle:dba /opt/oracleCRS
```

1.18 Kernel Parameters

The following script snippet sets the needed kernel parameters.

```
cat >> /etc/sysctl.conf <<EOF
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.file-max = 76800
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default=4194394
net.core.wmem_default=262144
net.core.rmem_max=4194304
```

```
net.core.wmem_max=262144
EOF
```

1.19 Additional /etc/pam.d Entries

The following script snippet sets the needed PAM parameters.

```
cat >> /etc/pam.d/login <<EOF
session required pam_limits.so
EOF
```

1.20 Configure /etc/security/limits.conf

The following script snippet sets the needed limits parameters.

```
cat >> /etc/security/limits.conf <<EOF
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
sapsys hard nofile 32800
```
sapsys soft nofile 32800
EOF

1.21 Create the /etc/profile.local

Place the following into /etc/profile.local, this file will need to be created.

```
cat >> /etc/profile <<EOF
if [ \$USER = "oracle" ]; then
 if [ \$SHELL = "/bin/ksh" ]; then
 ulimit -p 16384
 ulimit -n 65536
 else
 ulimit -u 16384 -n 65536
 fi
 umask 022
fi
EOF
```

Task#4 OCFS Installation

The installation and setup of OCFS for shared storage is a standard installation process. The only difference is that for a **SAP installation we will be sharing the following mandatory mount points across the cluster** –

- Oracle User Home
- Oracle CRS cluster home
- Oracle HOME for oracle binaries

Sharing these directories assists in not redoing the same setup exercise across different nodes in the cluster and is recommended by SAP.

Please download the relevant Linux packages from the Oracle website -
http://oss.oracle.com/projects/ocfs2/files/RedHat/RHEL4/x86_64/1.2.9-1/

[root@fremont01 ocfs]# ls
ocfs2-2.6.9-55.0.12.EL-1.2.9-1.el4.x86_64.rpm
ocfs2-2.6.9-55.0.12.ELsmp-1.2.9-1.el4.x86_64.rpm
ocfs2console-1.2.7-1.el4.x86_64.rpm
ocfs2-tools-1.2.7-1.el4.x86_64.rpm
ocfs2-tools-debuginfo-1.2.7-1.el4.x86_64.rpm
ocfs2-tools-devel-1.2.7-1.el4.x86_64.rpm

[root@fremont01 ocfs]# rpm -ivh ocfs2-tools-debuginfo-1.2.7-1.el4.x86_64.rpm \ ocfs2-tools-1.2.7-1.el4.x86_64.rpm \ ocfs2console-1.2.7-1.el4.x86_64.rpm \ ocfs2-2.6.9-55.0.12.ELsmp-1.2.9-1.el4.x86_64.rpm

```
Preparing...                ########################################### [100%]
   1:ocfs2-tools            ########################################### [ 25%]
   2:ocfs2-tools-debuginfo  ########################################### [ 50%]
   3:ocfs2console           ########################################### [ 75%]
   4:ocfs2-2.6.9-55.0.12.ELs########################################### [100%]
```

Configure Storage Using OCFS2

On the first node:

1 Log in as `root`.
2 Perform the following steps:

a. Start the X Window System by typing:
   `startx`

b. Generate the OCFS2 configuration file /etc/ocfs2/cluster.conf with a default cluster name of ocfs2 by typing the following in a terminal window:

ocfs2console

c. From the menu, click Cluster Configure Nodes. If the cluster is offline, the console will start it. A message window appears displaying that information. Close the message window. The Node Configuration window appears.

d. To add nodes to the cluster, click Add. Type the node name (same as the host name) and the private IP. Retain the default value of the port number. After typing all the details, click OK.

Repeat this step to add all the nodes to the cluster.

e. When all the nodes are added, click Apply and then click Close in the Node Configuration window.

NOTE: If you get the error message: `Unable to access cluster service`, delete the file: /etc/ocfs2/cluster.conf

and try again.

f. From the menu, click Cluster Propagate Configuration.

The Propagate Cluster Configuration window appears. Wait until the message `Finished` appears in the window and then click Close.

g. Select File Quit.

3 On all the nodes, enable the cluster stack on startup by typing:

```
/etc/init.d/o2cb enable
```

4 Change the O2CB_HEARTBEAT_THRESHOLD value on all the nodes using the following steps:

a. Stop the O2CB service on all the nodes by typing:

```
/etc/init.d/o2cb stop
```

b. Edit the O2CB_HEARTBEAT_THRESHOLD value in /etc/sysconfig/o2cb to 81 on all the nodes.

c. Start the O2CB service on all the nodes by typing:

```
/etc/init.d/o2cb start
```

5 On the first node, for a Fibre Channel cluster, create one partition on each of the other two external storage devices with fdisk:

a. Create a primary partition for the entire device by typing:

```
fdisk /dev/emcpowerX
```

Type `h` for help within the fdisk utility.

b. Verify that the new partition exists by typing:

```
cat /proc/partitions
```

c. If you do not observe the new partition, type:

```
sfdisk -R /dev/<device name>
```

The following sample values are used in the steps used below -:

- Local storage for Local cluster

```
/u01 - ocr, voting - 27GB
/u02 – crs and Oracle GRID Agent sw - 216GB
/oracle – Oracle software, datafiles, redo - 434GB
/flash - flashback - 650GB
/stage - stage  - 216GB
```

- Fiber Channel Storage devices: `emcpowera`, `emcpowerb`, and `emcpowerc`

6 On any one node, format the external storage devices with 4 K block size, 128 K cluster size, and 4 node slots (node slots refer to the number of cluster nodes) using the command line utility mkfs.ocfs2 as follows:

Oracle Cluster Home & Oracle GRID Agent

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L u01 /dev/emcpowerc1
```

OCR.dbf and Voting Disk

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L ocrvoting /dev/emcpowere1
```

Oracle User Home, Oracle Software Binaries & Database Files

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L oracle /dev/emcpowerb1
```

Flash Recovery Area & Staging

```
mkfs.ocfs2 -b 4K -C 128K -N 4 -L flash /dev/emcpowera1
mkfs.ocfs2 -b 4K -C 128K -N 4 -L stage /dev/emcpowerd1
```

NOTE: For more information about setting the format parameters of clusters, see www.oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html.

7 On each node, perform the following steps:

a. Create mount points for each OCFS2 partition. To perform this procedure, create the target partition directories and set the ownerships by typing:

```
mkdir -p /u01 /u02 /oracle /flash /stage
chown -R oracle.dba /u01 /u02 /oracle /flash /stage
```

b. On each node, modify the **/etc/fstab** by adding the lines for each device:

```
#ORACLE_HOME
```

```
/dev/emcpowerc1      /u01 ocfs2   _netdev   0 0

#Oracle Files
/dev/emcpowere1     /ocrvoting  ocfs2   _netdev,datavolume,nointr  0  0
/dev/emcpowerd1     /stage      ocfs2   _netdev,datavolume,nointr  0  0
/dev/emcpowerb1     /oracle     ocfs2   _netdev,datavolume,nointr  0  0
/dev/emcpowera1     /flash      ocfs2   _netdev,datavolume,nointr  0  0
```

If the PowerPath pseudo devices do not appear with exactly the same device name across all the nodes, modify the /etc/fstab file on each node(or use emcpadm to rename from one specific node) to ensure that all the shared directories on each node access the same disks. Make appropriate entries for all OCFS2 volumes.

c. On each node, type the following to mount all the volumes listed in the /etc/fstab file:
   `mount -a -t ocfs2`

d. On each node, add the following command to the /etc/rc.local file:
   `mount -a -t ocfs2`

OCFS2 has one configuration file, namely, /etc/ocfs2/cluster.conf. In it, one needs to specify all the nodes in the cluster. This file should be the same on all the nodes in the cluster. Whereas one can add new nodes to the cluster dynamically, any other change, like name, ip address, requires the cluster to be restarted for the changes to take effect.

We used ocfs2console to setup and propagate the cluster.conf to all the nodes in the cluster. This needs to be done only on one of the nodes in the cluster. This creates the same /etc/ocfs2/cluster.conf on all nodes in the cluster. Please note that differing cluster.conf(s) will lead to unpredictable behavior.

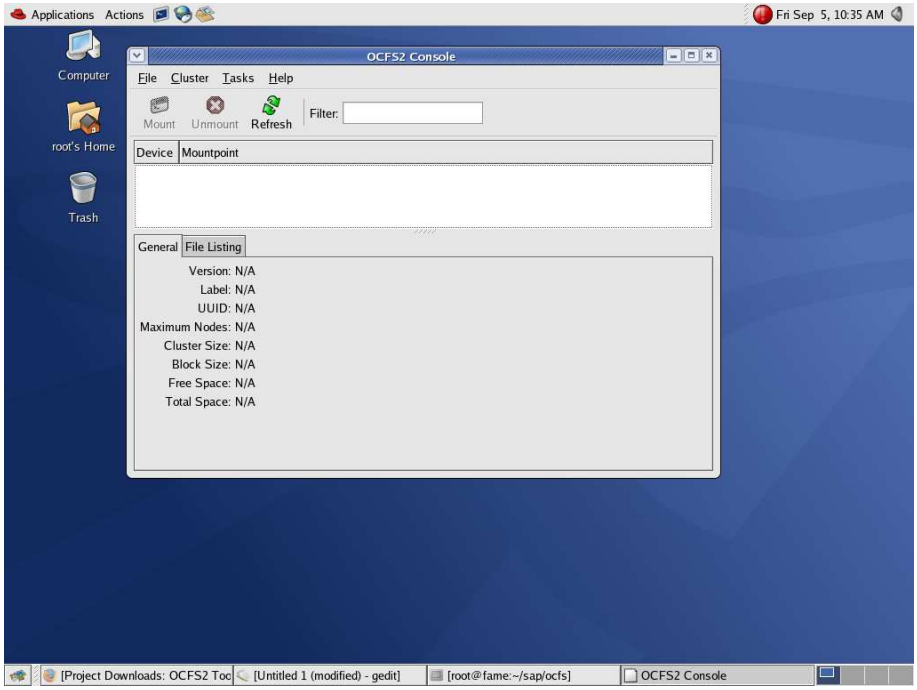

Start ocfs2console and click on menu item Cluster followed by Configure Nodes. If the cluster is offline, the console will start it and display a message to that effect. If cluster.conf is not present, the console will create one with a default cluster name ocfs2

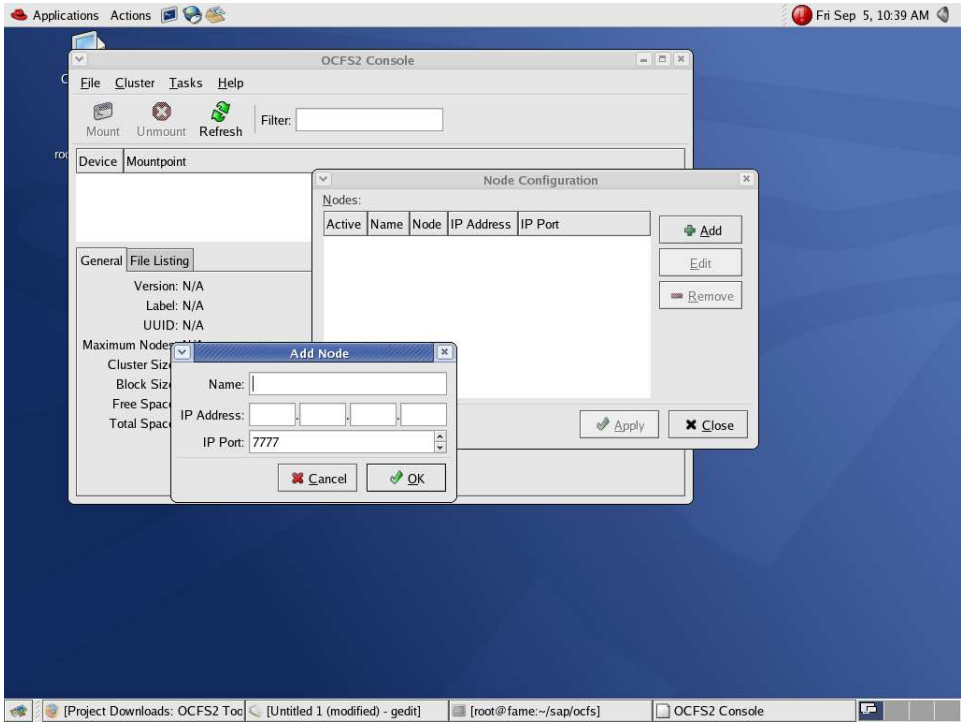

To verify that each node can detect each storage LUN or logical disk, perform the following steps:

1. Visually verify that the storage devices and the nodes are connected correctly to the Fiber Channel switch ..
2. Verify that you are logged in as root.
3. On each node, type:
    a. more /proc/partitions
    b. fdisk –l

The node detects and displays the LUNs or logical disks, as well as the partitions created on those external devices.

A list of the LUNs or logical disks that are detected by the node is displayed, as well as the partitions that are created on those external devices. Multipath pseudo devices appear in the list, such as

Task #5 Oracle Software Installation

After the shared storage and the OS server environment have been setup then the other key tasks is of setting up of the SSH for user equivalence across the cluster.

1.22 SSH SETUP

Log on to the first node as user oracle and the user's home is shared across the cluster. Create directory ~/.ssh, if it does not exist already. Go to this directory.

$ mkdir .ssh

$ cd .ssh

Create private and public key files for RSA and DSA authorization:

$ ssh-keygen –t rsa –f <path_to_homedirectory>/.ssh/id_rsa_<nodename>

Do not specify a passphrase, just hit Enter when asked.

$ ssh-keygen –t dsa –f <path_to_homedirectory>/.ssh/id_dsa_<nodename>

Once again, do not specify a passphrase. Hit Enter.

Repeat the above steps on all nodes. Remember that you must log on to every node, as key generation is dependent on the node itself. After all private and public key files have been generated; you must create a config file containing pointers to the Identity file for every node.

These Identity files contain private keys for the respective nodes.

You must specify the full pathname to these files. Example for the file ~/.ssh/config:

```
ForwardX11 no
PasswordAuthentication no
Host fremont01
IdentityFile /oracle/.ssh/id_rsa_fremont01
IdentityFile /oracle/.ssh/id_dsa_fremont01
Host fremont02
IdentityFile /oracle/.ssh/id_rsa_fremont02
IdentityFile /oracle/.ssh/id_dsa_fremont02
Host austin01
IdentityFile /oracle/.ssh/id_rsa_austin01
IdentityFile /oracle/.ssh/id_dsa_austin01
Host austin02
IdentityFile /oracle/.ssh/id_rsa_austin02
IdentityFile /oracle/.ssh/id_dsa_austin02
```

Concatenate all files containing public keys to the file authorized_keys

$ cat *.pub > authorized_keys

Change the file permissions as shown.
$chmod 600 config
$chmod 600 authorized_keys
$cd
$chmod 700 .ssh

Chmod 755 $HOME (that is the oracle home directory)

You should now be able to open a ssh connection to all hosts without being prompted for a password. Try this by executing
$ssh <nodename> date

Repeat this on all nodes in order to have the fingerprints recorded in the file known_hosts.

1.23 Oracle CRS Installation

This involves six steps –

- Step#1 : Create the Oracle/CRS environment
- Step#2 : Install the cluster ware using OUI
- Step#3 : Install the Oracle Software using the SAP Solution Manager
- Step#4 : Upgrade the Oracle Software using the Oracle RunInstaller after the solution manager completes its installation

Before installing the Oracle RAC 10*g* Release 2 database software, you must first install Oracle Clusterware.

```
cd /opt/stage/clusterware
./runInstaller
```

1. Welcome
   - Click on **Next**

2. Specify Inventory Directory and Credentials
   - Enter **/opt/oracle/oraInventory**
   - Group Name should be **dba**
   - Click on **Next**

3. Specify Home Details Installing Oracle RAC
   - Name**: OraCrs10g_home**
   - Path: **/opt/oracleCRS/product/10.2/crs_1**

4. Product-Specific Prerequisite Checks
    - Correct any problems found before proceeding.
    - Click on **Next**

5. Specify Cluster Configuration
    - Cluster Name: **sapcrs**
    - Add all nodes and fully-qualified names as needed.
    - Click on **Next**

6. Specify Network Interface Usage - Specify the Interface Type (public, private, or "do no use") for each interface

7. Specify Oracle Cluster Registry (OCR) Location
    - Choose **Normal Redundancy** and enter the raw file names
        - /u02/ocr01
        - /u02/ocr02

8. Specify Voting Disk Location
    - Choose **Normal Redundancy** and enter the raw file names
        - /u02/voting01
        - /u02/voting02
        - /u02/voting03

9. Summary
    - Click on **Install**

10. Execute Configuration Scripts
    - Execute the scripts as root on each node, one at a time, starting with the installation node.
    - Do not run the scripts simultaneously. Wait for one to finish before starting another.

11. Execute Configuration Scripts
    - Click **OK**

12. End of Installation
    - Click **Exit**

The snapshots are just a sample of the installation process that we followed for the CRS installation --

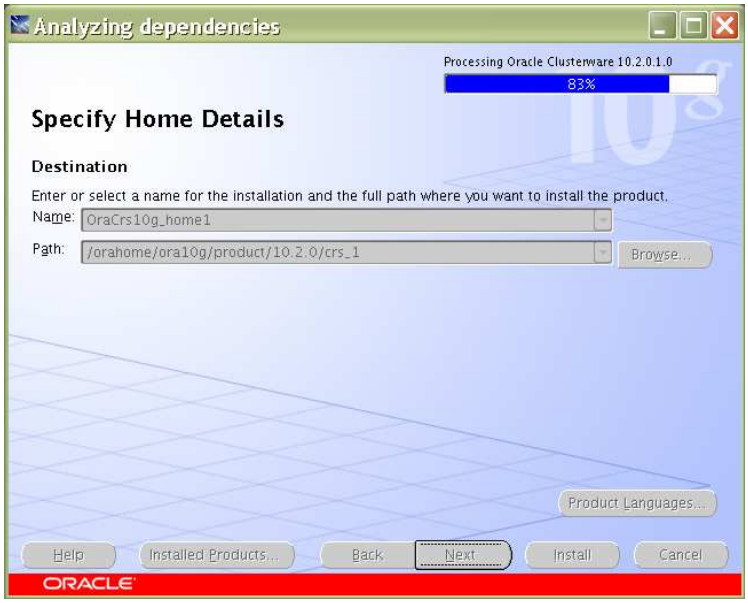

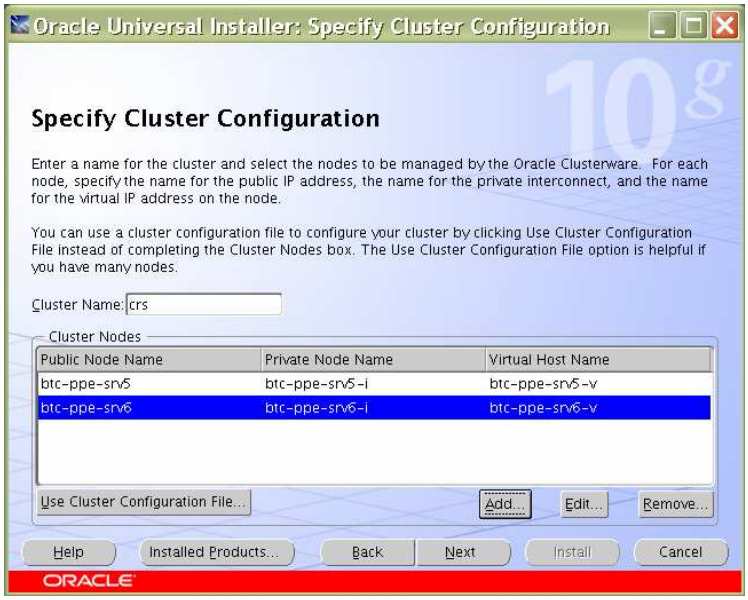

You make sure the Public, Private and VIP Hostname have the correct IP address and domain extensions etc

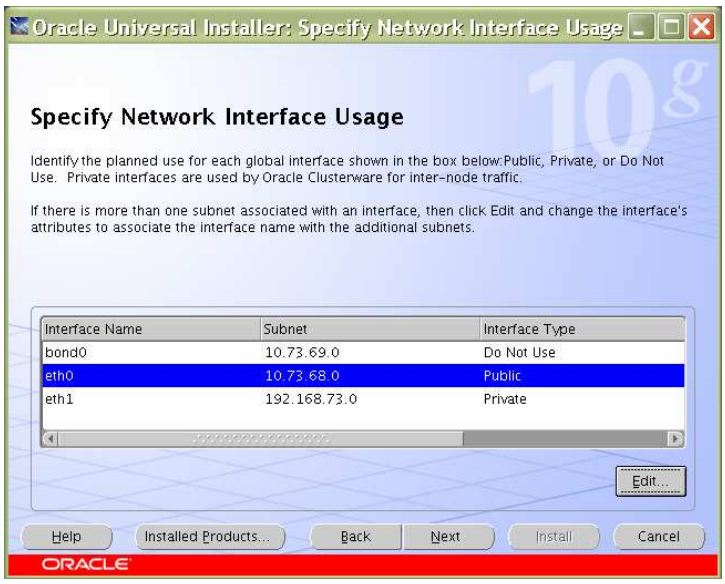

You modify the network interfaces – eth0, eth1 for PUBLIC and PRIVATE and make sure if there are any other ports you choose – Do Not Use

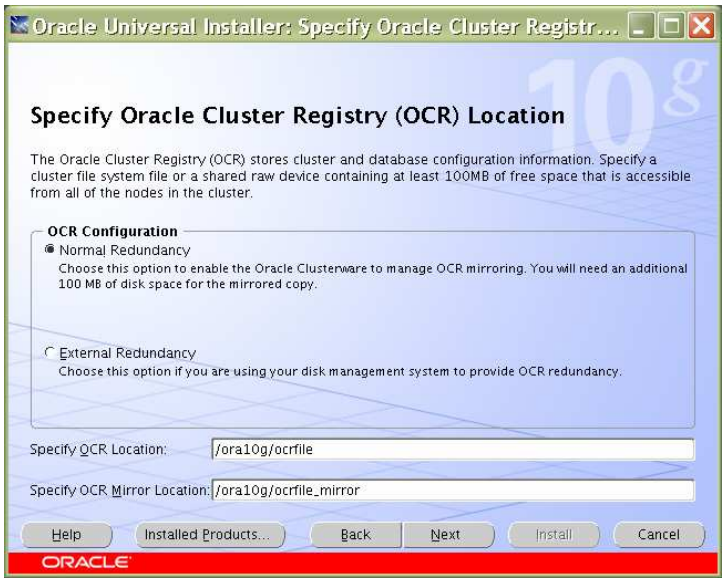

Define Normal redundancy with the correct path for the OCR disks to use – /ocrvoting/ocr01.dbf and /ocrvoting/ocr02.dbf

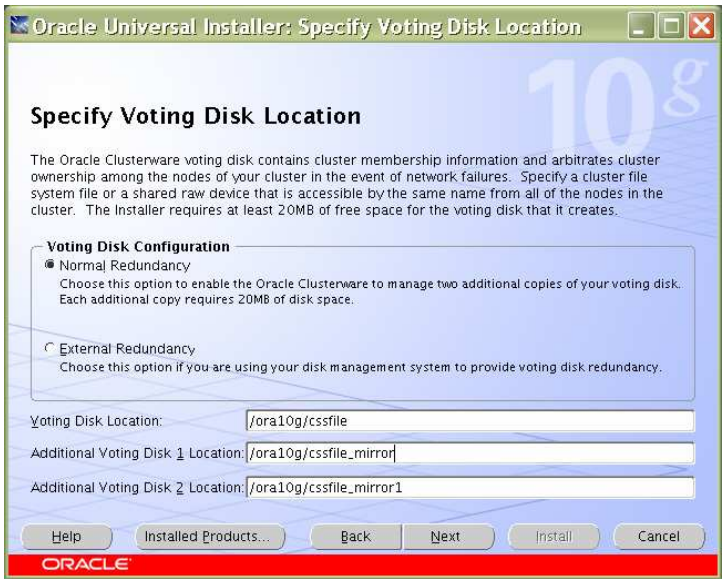

Define Normal redundancy with the correct path for the VOTING disks to use – /ocrvoting/voting01 , /ocrvoting/voting02 and /ocrvoting/voting03

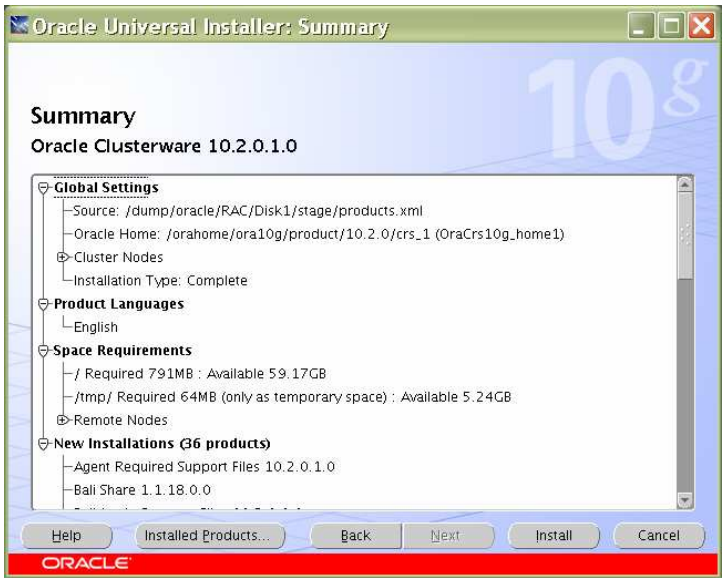

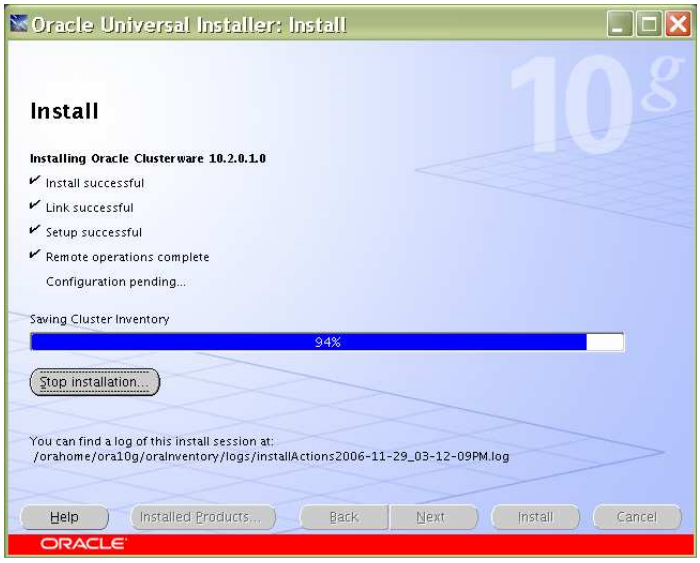

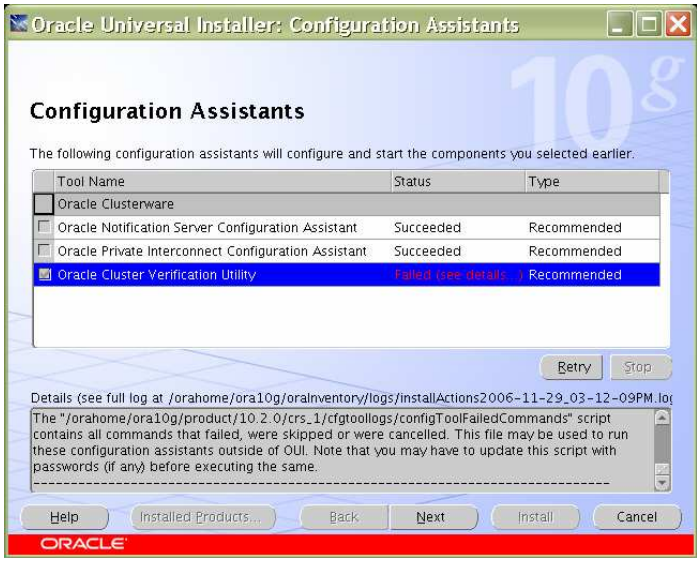

You will get the following address as the IP address used is one that is reserved for private networks, such as 10.0.0.0 or 192.168.0.0

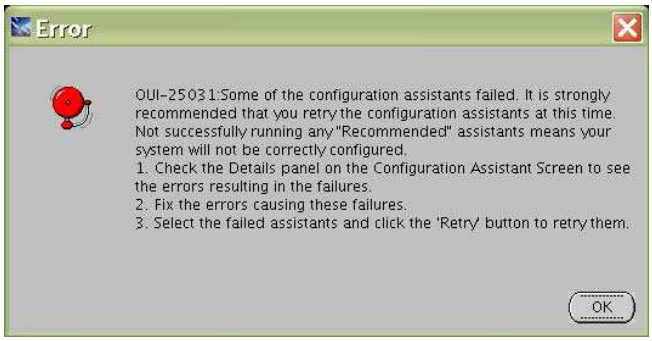

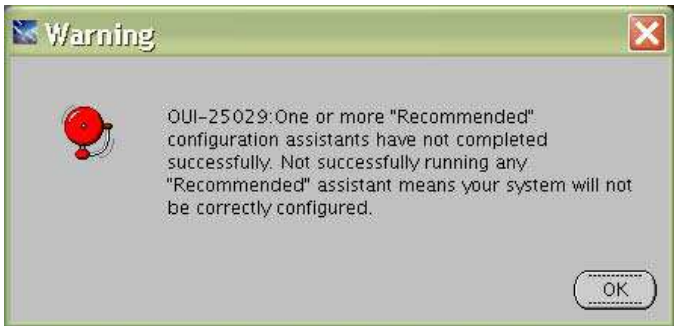

Run vipca utility from $ORA_CRS_HOME/bin directory as root user on the primary server.

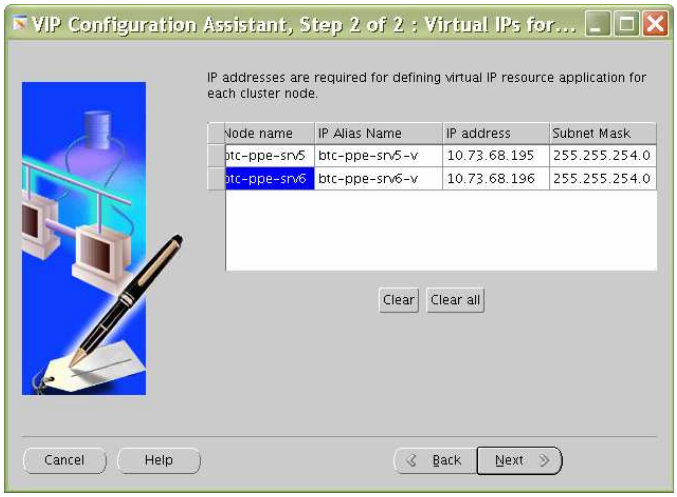

Pls make sure the VIP address is correctly entered in the table.

Then hit the retry the option for the configuration assistant if you have not already closed. This completes the CRS installation process.

The next logical step after verification is upgrading the CRS to 10.2.0.4 by using the Oracle OUI ***./runinstaller***

1.24 CRS Verification

At this stage, the following is true:

a. The clusterware has NOW been installed on node fremont01 and fremont02 and is part of the cluster.

b. Run the crs_stat –t command to make sure all the services are up and running.

c. Verify that OCR is aware of the cluster

```
[root@ fremont01 bin]# ./ocrcheck
Status of Oracle Cluster Registry is as follows :
         Version                  :          2
         Total space (kbytes)     :    1048100
         Used space (kbytes)      :       4368
         Available space (kbytes) :    1043732
         ID                       : 1117547110
         Device/File Name         : /u02/ocr01
                                    Device/File integrity check succeeded
         Device/File Name         : /u02/ocr02
                                    Device/File integrity check succeeded

         Cluster registry integrity check succeeded
```

Log Files for Clusterware

The log files are located in the $CRS_HOME/log/hostname directory. This directory contains subdirectories for crsd, client, cssd, evmd and racg.

For easy administration and navigation, you should define several different environment variables in the login profile for ORACLE_HOME, CRS_HOME, ORACLE_SID and the other standard oracle directories.

The next step is to install the Oracle GRID control.

Task#6 Oracle GRID Control Installation

It is recommend to physically locate the server in a third location and if it's not possible then the next best choice would be to locate it at the STANDBY site. For our example we will call it the DALLAS server.

Login to the DALLAS server and install the Oracle GRID Control and the agents.

The process will involve installing Enterprise Manager 10g Grid Control Using a New Database and then upgrading it to 10.2.0.4 by using the OUI. Do not attempt to use the "Installing Software-Only and configuring later" process though it's a faster and better approach as at times we may face the **bug#602750.1**

This describes how to install Enterprise Manager 10g Grid Control using a New Database

- Overview
- Prerequisites
- Installing EM10G Grid Control using a new DB

Using the Universal Installer, you will install Enterprise Manager 10g Grid Control using a new database.

To install the Oracle software you must use the GUI installer.

Login to the Linux box as user oracle and mount the Oracle Database 10g CD. Change directory to the CD and execute the script **runInstaller**.

```
./runInstaller
```

Make sure the Installation Type **Enterprise Manager 10g Grid Control Using a New Database** is selected and click **Next**.

(This is where you can choose to reuse an existing database)

The Universal Installer will make all the necessary prerequisite checks. When done, click **Next**.

Enter a password for the default super user SYSMAN and click **Next**.

Enter a password for all database accounts and click **Next**.

At this time, you will not setup access to Metalink and Proxies. You will configure Metalink in another section. Click **Next**.

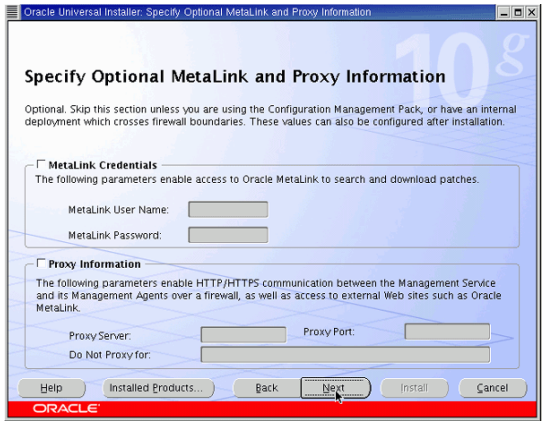

Enter **emrep.<yourdomain.com>** for the Global Database Name and **emrep** for the SID. This is the database that will be created where the OMS will reside. Then click **Next**.

At the Summary screen, review what will be installed and click **Install**.

You will see the progress window.

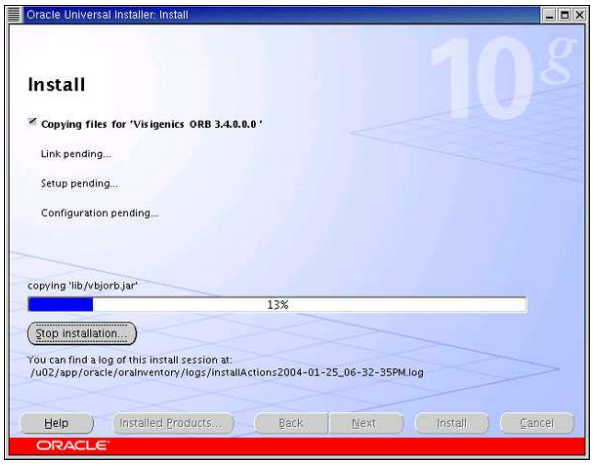

When the Setup Privileges window appears, open a new terminal window

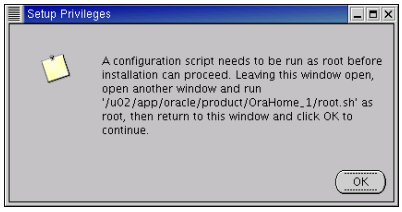

You need to execute **`root.sh`** as the **root** user. Open a terminal window and enter the following commands:

```
cd $ORACLE_BASE/product/OraHome_1
su
<rootpassword>
./root.sh
exit
exit
```

The Configuration Assistants Window will appear.

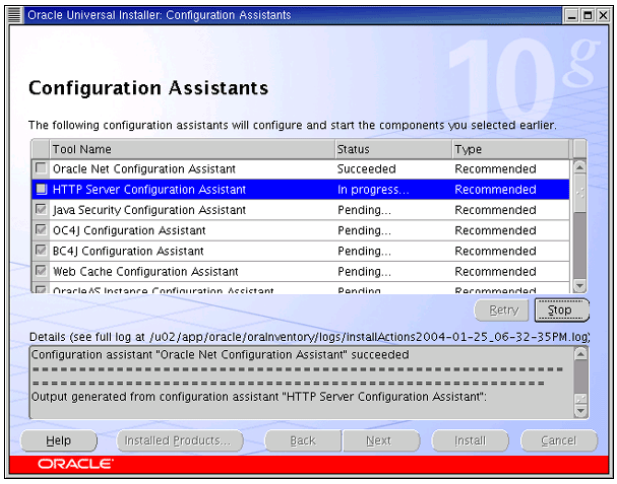

Your database is now being created.

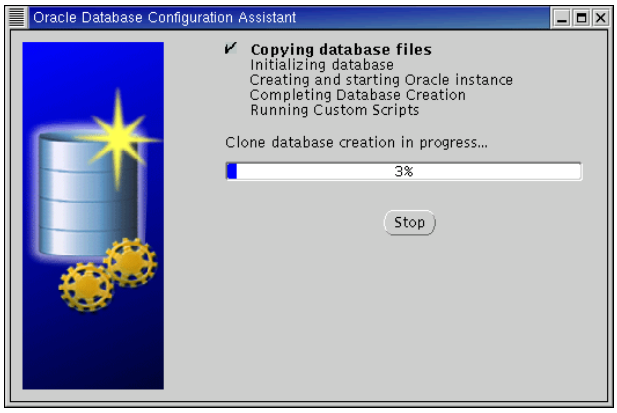

The Oracle Agent will now be installed on the same machine as the database that was just created. This database will then be available through Grid Control to manage its environment. Click **OK**.

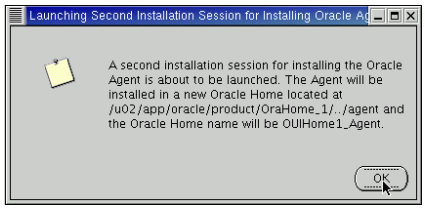

You will see the progress window.

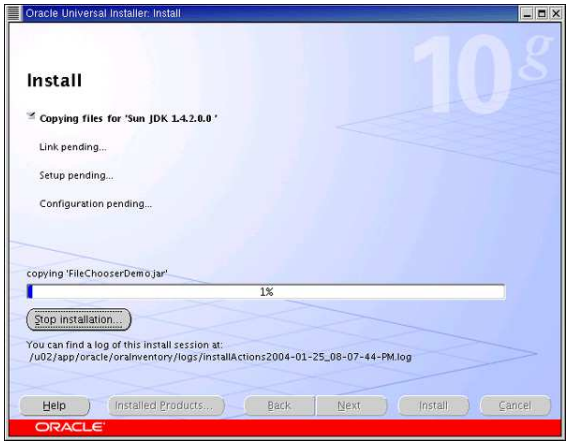

When the Setup Privileges window appears, open a new terminal window

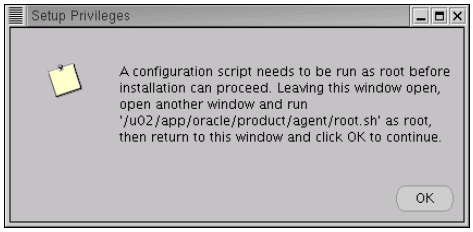

You need to execute **`root.sh`** as the **root** user. Open a terminal window and enter the following commands:

```
cd $ORACLE_BASE/product/agent
su
<rootpassword>
./root.sh
exit
exit
```

The Configuration Assistants Window will appear.

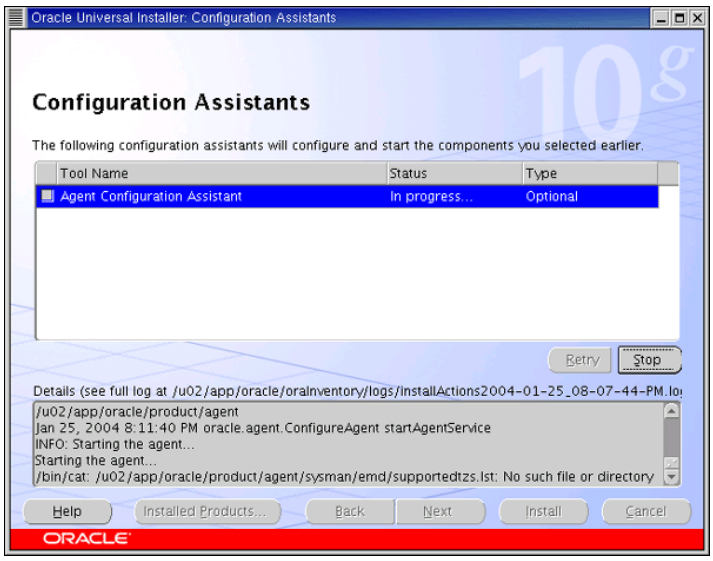

Your installation is now complete.

Click **Exit**.

Click **Yes** to exit.

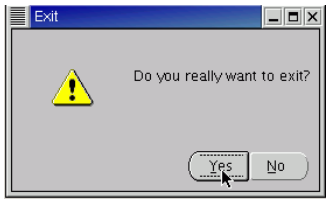

The Next Steps for the Oracle GRID control installation are:

1. **Upgrade the Oracle GRID control DB, OMS and Agent to its latest release 10.2.0.4 by using OUI and following the steps prompted by the installer.**

2. **Install the Oracle Enterprise Manager Agent on the DB servers. Please make sure to download the 64bit OEM Agent 10.2.0.4 for Linux**

Task#7 SAP Application Server Installation

The SAP Netweaver AS java architecture contains many features that are useful for building a HA Design.. The proven approach from the ABAP stack has been transferred to the JAVA stack in SAP NetWeaver. The JAVA Central Instance (Java CI) contains the Java dispatcher, which is similar to the ABAP dispatcher, receives client requests and forwards them to the appropriate serer processes. It is the Java server process that actually processes the requests and holds the session data.

Access to data in the underlying database of an SAP system is synchronized with a special lock system called the SAP Enqueue mechanism. This mechanism serializes access and prevents access from being changed for more than one requesting party.

The enqueue server usually runs as a service of the SAP central instance. If clients run in the same SAP instance, they can contact the enqueue server via the UNIX® Interprocess Communication (IPC) mechanism; if they are not part of the central instance, clients contact the enqueue server via the SAP message server. As opposed to all other components of the SAP system on the application layer, the enqueue server holds a state—an in-memory table of granted locks—that cannot be recovered gracefully if the service fails. The message server, which consequently plays an important role in contacting the enqueue server, holds no state; it receives only incoming connection requests and transfers them to the addressee. The message server can be restarted after failure, with no impact other than delayed communications. The enqueue server is a potential single point of failure in an SAP system, isolated from the failover provided at the database layer.

To secure the SAP system's services in a high-availability cluster, administrators must split the traditional central instance into dedicated instances because a large "service block" can be difficult to monitor. Furthermore, this large block makes restarting services difficult, because administrators must also restart parts of the central instance that have not failed. To run the enqueue server as a master/slave service, the enqueue service and the enqueue replication service should always reside on different hosts. The message server is not bound to a particular host. Because these are the two services that constitute a central instance, the cluster can run only those services, and all application servers must be outside the cluster. However, for systems management purposes, the message server can run together with a dialog service, and an application server can reside in the cluster, or close to it. In the example scenario, the Enqueue server; the enqueue replication server; the message server; and an application instance with dialog, update, batch, and spool work processes all run as services in the cluster. In the example scenario, the traditional central instance DVEBMGS is split into multiple instances as follows (the two numerals at the end of each instance name represent the system number):

- DVBGS00: Dialog, update, batch, gateway, and spool work processes
- DM01: Dialog service (for local administration) and message server
- E02: Enqueue server
- R02: Enqueue replication server

This splitting of the central instance and making the Enqueue resources highly available is not within the scope of this paper. For more information refer Oracle Whitepaper - Providing High Availability for SAP Resources, which discusses about SAPCTL etc.

The application server needs a central file share (/sapmnt/…) which is accessed by all servers in the environment. The NFS shared mount point is mounted across all the servers from the appserver, which follows the usual process of creating the mount points and identifying them in the fstab.

- NFS Mount the /sapmnt such that its **writeable** from all nodes
    - Modify the /etc/exports file on the **APPLICATION** server

    /sapmnt fremont01.dell.com(rw,no_root_squash,sync,wdelay)
    /sapmnt fremont02.dell.com(rw,no_root_squash,sync,wdelay)
    /sapmnt austin01.dell.com(rw,no_root_squash,sync,wdelay)
    /sapmnt austin02.dell.com(rw,no_root_squash,sync,wdelay)
    …

    - Start the NFS service
- Modify /etc/fstab on the DB Servers - fremont01, fremont02, austin01 & austin02
    - Mount <m/c name>:/sap/sapmnt /sapmnt

The following Software DVDs for ECC media should be available for installing ECC system –

- Installation Master DVD
- Kernel DVD
- Netweaver Java DVD
- Business Suite JAVA DVD
- RDBMS DVD
- RDBMS Client DVD
- RDBMS patch DVD
- Export DVD

Installing the Java Development Kit and JCE Policy
If the JDK is not already installed, you need to download from the **IBM site (not SUN)**
Refer SAP Note # 1172419

https://www14.software.ibm.com/webapp/iwm/web/reg/download.do?source=lxdk&S_PKG=amd64142sr10&S_TACT=105AGX05&S_CMP=JDK&lang=en_US&cp=UTF-8

Install it under /usr (as shown below -> folder java14_64 for 64 bit platform)

As of SAP NetWeaver 7.0 strong encryption is mandatory for the J2EE Engine and for all usage types that are deployed on it. You need to obtain the JCE policy files beforehand so that they can be installed by SAPinst.

1. Download the JCE policy files for your platform at:
   http://www6.software.ibm.com/dl/jcesdk/jcesdk-p.
   File name => unrestrict142.zip

2. Copy the JCE policy files to directory /usr/java14_64/jre/lib/security

Creating OS Groups and Users

The OS users and groups as described in the table can be precreated but it's recommended to use the SAP installer to create them for you and on the PRIMARY(FREMONT) servers while for the STANDBY servers you have to manually create the users and groups.

- Build two Linux Application servers (one at FREMONT and other at AUSTIN) . You can have multiple application servers running Dialog Instances but for this example we are assuming only one appserver per data center.
- Create two Linux Oracle RAC clusters (one at FREMONT and other at AUSTIN – two servers in each cluster)
- On the FREMONT **application** server create the following file systems and mount them to each of the servers in the SAP Linux architecture:
  - /sapmnt (atleast 10 GB) : NFS Mounted
  - /usr/sap (atleast 20 GB) : Local Drive

DBSID=SAPSID=ZLS

Users and Groups that need to be created for both Application servers and DB servers

User	Primary Group
<sapsid>adm	sapsys, oper, dba, sapinst
ora<dbsid>	dba, oper, sapinst
Groups	**Members**
sapsys	<sapsid>adm
oper	<sapsid>adm, ora<dbsid>
dba	<sapsid>adm, ora<dbsid>
sapinst	<sapsid>adm, ora<dbsid>

SAP File System for an ABAP+Java System (Unicode)

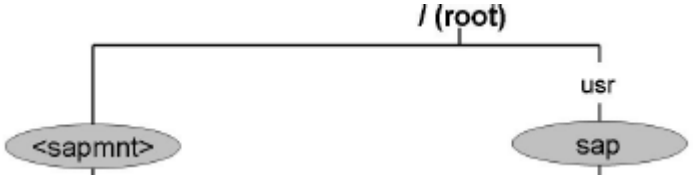

NFS Mounted to all servers Local to the servers

/sapmnt => 5.0 GB
/usr/sap => 5.0 GB

Sample Oracle File System Requirement and in reality you can have larger disk space.

	Mount Points	Space Requirement(minimum)
1	/oracle	500 MB
2	/oracle/client	500 MB
3	/oracle/stage/<102_64>	8.0 GB
4	/oracle/<DBSID>	8.0 GB
5	/oracle/<DBSID>/origlogA	1.0 GB
6	/oracle/<DBSID>/origlogB	1.0 GB
7	/oracle/<DBSID>/mirrlogA	1.0 GB
8	/oracle/<DBSID>/mirrlogB	1.0 GB
9	/oracle/<DBSID>/oraarch	10.0 GB
10	/oracle/<DBSID>/sapreorg	5.0 GB
11	/oracle/<DBSID>/sapdata1	45.0 GB
12	/oracle/<DBSID>/sapdata2	45.0 GB
13	/oracle/<DBSID>/sapdata3	45.0 GB
14	/oracle/<DBSID>/sapdata4	45.0 GB

 This is based on central system installation

** Make sure that file system is created as large file system

** Set default UMASK to 022

** Set ulimit for root/orasidadm/sidadm is set to unlimited

ulimit -d unlimited

ulimit -s unlimited

ulimit -m unlimited

The installation of the SAP ECC system is to be started using the solution manager's utility sapinst

cd /tmp

mkdir scsinstall

cd scsinstall

Start the installation process by invoking the sapinst from this directory. This directory will contain all the temporary installation data files.

<PATH>/installation master/sapinst

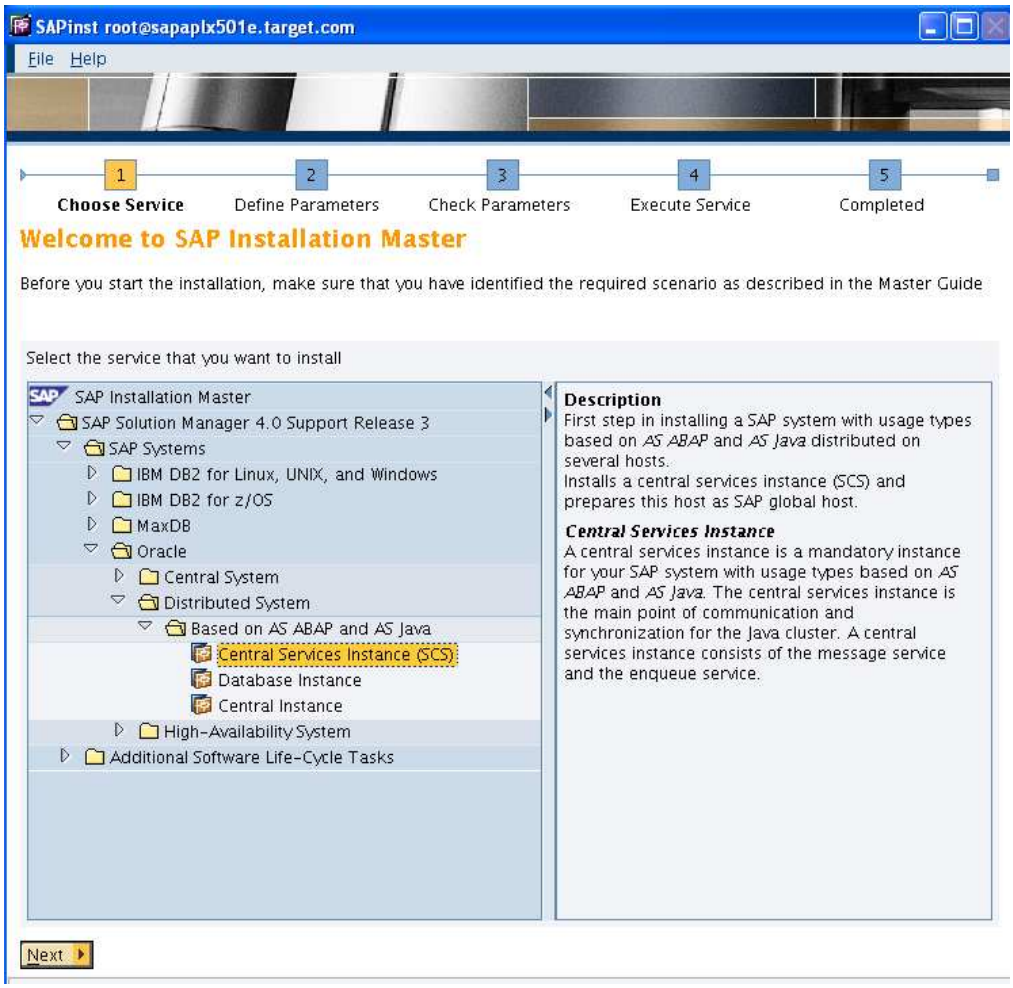

Note: The installation screen shots for the Central Services instance (SCS) and Central Instance (CI) are available in an Appendix X

The SAP Application installation process involves the following steps –

1. Central Services Instance (SCS) Installation
    a. This is a complete GUI driven process and the steps are in the appendix
2. Database Instance Installation (Refer Appendix D)
    a. This is a complete GUI driven process. This step will install the Oracle binaries and create a single instance database 10.2.0.1.
3. Central Instance Installation

After all the above three steps are completed we start with the tasks listed below.

Warning ::

- SAP Is not supported on 10.2.0.3

- SAP is not supported with ASM
- Do not precreate the Oracle database as then the SAPINST will fail.
- SAP Oracle Database needs to be upgraded to 10.2.0.4 and later relevant Oracle patches can be applied as per **SAP Note – 1137346**
- The Oracle Database can be upgraded using the installer from Oracle.

Task#8 Convert Single Instance Oracle Database to RAC

Step #1Startup mount
Run catclust.sql

Create addition UNDO tablespaces

```
create undo tablespace PSAPUNDO_001 datafile
'/oracle/ZLS/sapdata1/undo/undo.data1 ' size 1000m reuse;

create undo tablespace PSAPUNDO_002 datafile
'/oracle/ZLS/sapdata2/undo/undo.data2 ' size 1000m reuse;
```

Step#2 : Modify the spfile for the newly created UNDO tablespaces

```
*.undo_management = auto
ZLS001.undo_tablespace = PSAPUNDO_001
ZLS002.undo_tablespace = PSAPUNDO_002
```

Step#3 : Create additional redo logfiles for the RAC instances

```
alter database add logfile thread 1 group 11
(´/oracle/ZLS/origlogA/log_g11m1.dbf´,
´/oracle/ZLS/mirrlogA/log_g11m2.dbf´) size 200M reuse;
alter database add logfile thread 1 group 12
(´/oracle/ZLS/origlogB/log_g12m1.dbf´,
´/oracle/ZLS/mirrlogB/log_g12m2.dbf´) size 200M reuse;
alter database add logfile thread 1 group 13
(´/oracle/ZLS/origlogA/log_g13m1.dbf´,
´/oracle/ZLS/mirrlogA/log_g13m2.dbf´) size 200M reuse;
alter database add logfile thread 1 group 14
(´/oracle/ZLS/origlogB/log_g14m1.dbf´,
´/oracle/ZLS/mirrlogB/log_g14m2.dbf´) size 200M reuse;


alter database add logfile thread 2 group 21
(´/oracle/ZLS/origlogA/log_g21m1.dbf´,
´/oracle/ZLS/mirrlogA/log_g21m2.dbf´) size 200M reuse;
alter database add logfile thread 2 group 22
(´/oracle/ZLS/origlogB/log_g22m1.dbf´,
´/oracle/ZLS/mirrlogB/log_g22m2.dbf´) size 200M reuse;
```

```
alter database add logfile thread 2 group 23
(´/oracle/ZLS/origlogA/log_g23m1.dbf´,
´/oracle/ZLS/mirrlogA/log_g23m2.dbf´) size 200M reuse;
alter database add logfile thread 2 group 24
(´/oracle/ZLS/origlogB/log_g24m1.dbf´,
´/oracle/ZLS/mirrlogB/log_g24m2.dbf´) size 200M reuse;
```

Step#4 : Enable the REDO log for the new instances

```
alter database enable public thread 1;
alter database enable public thread 2;
```

Step#5 : Drop the logfiles that are not required after verifying the status

```
select group#, archived, status from v$log;

REM switch to drop

alter system switch logfile;

alter database drop logfile group 101;
alter database drop logfile group 102;
```

Step#6

```
alter system set instance_number = 1
scope = spfile sid = 'RAC001';

instance_number = <thread>   01
thread = <thread>            01
instance_name = <dbsid><thread>  ZLS001
service_names = (<dbsid>, <dbsid><thread>     ZLS,ZLS001
*.undo_management = auto
undo_tablespace=<tablespacename>   PSAPUNDO_001
```

Step#7

After the initZLS.ora changes... create an spfile and delete all init.ora files

Step#8

Start the database instance by setting the ORACLE_SID=ZLS_001

Step#9

Create a LISTENER

LISTENER_ZLS_<NODE NAME IS AUTOMATICALLY ADDED>
Port - 1527

Create the TNS NAMES

Step#10

$ sqlplus / as sysdba
SQL> startup nomount

SQL> alter system set local_listener='LISTENER_ZLS_fremont01'
scope=spfile sid='ZLS001';
SQL> alter system set local_listener='LISTENER_ZLS_fremont02'
scope=spfile sid='ZLS002';

SQL> alter system set remote_listener='REMOTE_LISTENER_ZLS_fremont01'
scope=spfile sid='ZLS001';
SQL> alter system set remote_listener='REMOTE_LISTENER_ZLS_fremont02'
scope=spfile sid='ZLS002';

Shutdown immediate;

srvctl add database –d ZLS –o /oracle/ZLS/1020_64

The next step is the addition of all database instances to the configuration. In the example below, 2 instances are added:

srvctl add instance –d ZLS –i ZLS001 –n fremont01
srvctl add instance –d ZLS –i ZLS002 –n fremont02

srvctl add service –d ZLS –s D01_fremont01 –r ZLS001 –a ZLS002
srvctl add service –d ZLS –s D02_fremont02 –r ZLS002 –a ZLS001

Task# 9 : Create RAC Aware Oracle Data Guard Environment

The first step in setting up the Oracle GRID Data Guard environment is by to create the standby database using a manual process of backing up and copying the primary database to the remote site.

To setup the STANDBY environment shutdown the PRIMARY Oracle Database and then tar copy the **SAP User environments (UNIX Users, /usr/sap and /oracle)** to the Standby server. This does not include the Oracle DB as this would be copied using the RMAN process in the tasks listed below.

The groups and users are to be manually created on the standby database cluster. Please make sure the groupid, groups, userids are identical across all the environments.

##APP server

groupadd -g 500 sapinst

groupadd -g 503 sapsys

useradd -c "SAP System Administrator" -m -g "sapsys" -G "sapinst" -s "/bin/csh" -u 500 zlsadm

###DB Server

groupadd -g 502 sapinst

groupadd -g 503 sapsys

groupadd -g 504 oper

useradd -c "SAP System Administrator" -m -g "sapsys" -G "sapinst,dba,oper" -s "/bin/csh" -u 500 zlsadm

useradd -c "SAP Database Administrator" -M -d "/oracle/ZLS" -g "dba" -G "oper,sapinst" -s "/bin/csh" -u 503 orazls

Assumptions for implementing Oracle Data Guard:

- Shared Oracle Clusterware has already been setup identically on both primary and secondary sites
- Shared Oracle HOME already exists
- Directory structure is identical on both servers
- DBSID and Unix path for all files need to be the same
- The DB_UNIQUE_NAME needs to be different for identifying the database
- /sapmnt is NFS mounted to standby server

1.25 TASK 1: GATHER FILES AND PERFORM BACK UP

1. On the primary node, create a staging directory. For example:

[oracle@ZLS_host1 oracle]$ mkdir -p /oracle/stage

2. Create the same exact path on one of the standby hosts:

[oracle@ZLSSTBY_host1 oracle]$ mkdir -p /oracle/stage

3. On the primary node, connect to the primary database and create a PFILE from the SPFILE in the staging directory. For example:

SQL> CREATE PFILE='/oracle/stage/initZLS.ora' FROM SPFILE;

4. On the primary node, perform an RMAN backup of the primary database that places the backup pieces into the staging directory. For example:

   [oracle@ZLS_host1 stage]$ rman target /

   RMAN> BACKUP DEVICE TYPE DISK FORMAT '/oracle/stage/%U' DATABASE PLUS ARCHIVELOG;

   RMAN> BACKUP DEVICE TYPE DISK FORMAT '/oracle/stage/%U' CURRENT CONTROLFILE FOR STANDBY;

Note: Prior to backup you should delete the unnecessary archive logs and then perform -

   RMAN>crosscheck archivelog all

5. Place a copy of the listener.ora, tnsnames.ora, and sqlnet.ora files into the staging directory. For example:

[oracle@ZLS_host1 oracle]$ cp $ORACLE_HOME/network/admin/*.ora /oracle/stage

6. Copy the contents of the staging directory on the RAC primary node to the standby node on which the staging directory was created on in step 2. For example:

   [oracle@ZLS_host1 oracle]$ scp /oracle/stage/* \

   oracle@ZLSSTBY_host1:/oracle/stage

1.26 TASK 2: CONFIGURE ORACLE NET SERVICES ON THE STANDBY

1. Copy the listener.ora, tnsnames.ora, and sqlnet.ora files from the staging directory on the standby host to the $ORACLE_HOME/network/admin directory on all standby hosts.

2. Modify the listener.ora file each standby host to contain the VIP address of that host.

3. Modify the tnsnames.ora file on each node, including the primary RAC nodes and standby RAC nodes, to contain all primary and standby net service names. You should also modify the Oracle Net aliases that are used for the local_listener and remote_listener parameters to point to the listener on each standby host.
4. Start the standby listeners on all standby hosts.

1.27 TASK 3: CREATE THE STANDBY INSTANCES AND DATABASE

1. To enable secure transmission of redo data, make sure the primary and standby databases use a password file, and make sure the password for the SYS user is identical on every system. For example:

   $ cd $ORACLE_HOME/dbs

   $ orapwd file=orapwZLSSTBY password=oracle

   Note : Please verify by sqlplus>select * from v$pwfile_users;

2. Copy and rename the primary database PFILE from the staging area on all standby hosts to the $ORACLE_HOME/dbs directory on all standby hosts. For example:

[oracle@ZLSSTBY_host1 stage]$ cp initZLS1.ora $ORACLE_HOME/dbs/initZLSSTBY1.ora

3. Modify the standby initialization parameter file copied from the primary node to include Data Guard parameters

   Connect to the standby database on one standby host, with the standby in the IDLE state, and create an SPFILE in the standby DATA disk group: SQL> CREATE SPFILE='+DATA/ZLSSTBY/spfileZLSSTBY.ora' FROM PFILE='?/dbs/initZLSSTBY.ora';

4. In the $ORACLE_HOME/dbs directory on each standby host, create a PFILE that is named initoracle_sid.ora that contains a pointer to the SPFILE. For example:

   [oracle@ZLSSTBY_host1 oracle]$ cd $ORACLE_HOME/dbs

   [oracle@ZLSSTBY_host1 dbs]$ echo "SPFILE='+DATA/ZLSSTBY/spfileZLSSTBY.ora'" > initZLSSTBY1.ora

5. Create the dump directories on all standby hosts as referenced in the standby initialization parameter file. For example:

   [oracle@ZLSSTBY_host1 oracle]$ mkdir -p $ORACLE_BASE/admin/ZLSSTBY/bdump

   [oracle@ZLSSTBY_host1 oracle]$ mkdir -p $ORACLE_BASE/admin/ZLSSTBY/cdump

   [oracle@ZLSSTBY_host1 oracle]$ mkdir -p $ORACLE_BASE/admin/ZLSSTBY/udump

   [oracle@ZLSSTBY_host1 oracle]$ mkdir -p $ORACLE_BASE/admin/ZLSSTBY/adump

6. After setting up the appropriate environment variables on each standby host, such as ORACLE_SID, ORACLE_HOME, and PATH, start the standby database instance on the standby host that has the staging directory, without mounting the control file.

SQL> STARTUP NOMOUNT

7. From the standby host where the standby instance was just started, duplicate the primary database as a standby

   $ rman target sys/oracle@ZLS auxiliary /

   RMAN> DUPLICATE TARGET DATABASE FOR STANDBY;

   Note : You may have to use the option – *nofilename check*

8. Connect to the standby database, and create the standby redo logs to support the standby role. The standby redo logs must be the same size as the primary database online logs. The recommended number of standby redo logs is:

(maximum # of logfiles +1) * maximum # of threads

This example uses two online log files for each thread. Thus, the number of standby redo logs should be (2 + 1) * 2 = 6. That is, one more standby redo log file for each thread.

```
SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 1
GROUP 5 SIZE 100M,
GROUP 6 SIZE 100M,
GROUP 7 SIZE 100M;
SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 2
GROUP 8 SIZE 100M,
GROUP 9 SIZE 100M,
GROUP 10 SIZE 100M;
```

You can check the number and group numbers of the redo logs by querying the V$LOG view:

```
SQL> SELECT * FROM V$LOG;
```

You can check the results of the previous statements by querying the V$STANDBY_LOG view:

```
SQL> SELECT * FROM V$STANDBY_LOG;
```

You can also see the members created by querying the V$LOGFILE view:

```
SQL> SELECT * FROM V$LOGFILE;
```

9. On only one standby host (and this is your designated Redo Apply instance), start managed recovery and real-time apply on the standby database:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE USING CURRENT LOGFILE
DISCONNECT;
```

10. On either node of the standby cluster, register the standby database and the database instances with the Oracle Cluster Registry (OCR) using the Server Control (SRVCTL) utility. For example:

```
$ srvctl add database -d ZLSSTBY –o /oracle/product/10g_db_rac
$ srvctl add instance -d ZLSSTBY -i ZLSSTBY1 -n ZLSSTBY_host1
$ srvctl add instance -d ZLSSTBY -i ZLSSTBY2 -n ZLSSTBY_host2
```

The following are descriptions of the options in these commands:

The -d option specifies the database unique name (DB_UNIQUE_NAME) of the database.

The -i option specifies the database insance name.

The -n option specifies the node on which the instance is running.

The -o option specifies the Oracle home of the database.

1.28 TASK 4: CONFIGURE THE PRIMARY DATABASE FOR DATA GUARD

1. Configure the primary database initialization parameters to support both the primary and standby roles.

   *.log_archive_config='dg_config=(ZLSSTBY,ZLS)'
   *.log_archive_dest_2='service=ZLSSTBY
   valid_for=(online_logfiles,primary_role)
   db_unique_name=ZLSSTBY'
   *.standby_file_management=auto
   *.fal_server='ZLSSTBY'
   *.fal_client='ZLS'
   *.service_names=ZLS

Note that all the parameters listed above can be dynamically modified. The directory structure between the primary and standby has to be identical and hence we do not need to use the convert (log_file_name_convert and db_file_name_convert) parameters. It is recommended to set the parameters with "scope=spfile" so that they can be put into effect upon the next role change.

2. Create standby redo logs on the primary database to support the standby role. The standby redo logs are the same size as the primary database online logs. The recommended number of standby redo logs is one more than the number of online redo logs for each thread. Because this example has two online redo logs for each thread, three standby redo logs are required for each thread.

   SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 1
   GROUP 5 SIZE 100M,
   GROUP 6 SIZE 100M,
   GROUP 7 SIZE 100M;
   SQL> ALTER DATABASE ADD STANDBY LOGFILE THREAD 2
   GROUP 8 SIZE 100M,
   GROUP 9 SIZE 100M,
   GROUP 10 SIZE 100M;

3. These statements create two standby log members for each group, and each member is 100MB in size. One member is created in the directory specified by the DB_CREATE_FILE_DEST initialization parameter, and the other member is created in the directory specified by DB_RECOVERY_FILE_DEST initialization parameter. Because this example assumes that there are two redo log groups in two threads, the next group is group five.

   You can check the number and group numbers of the redo logs by querying the V$LOG view:
   SQL> SELECT * FROM V$LOG;
   You can check the results of the previous statements by querying the V$STANDBY_LOG view:

SQL> SELECT * FROM V$STANDBY_LOG;

You can also see the members created by querying V$LOGFILE view:

SQL> SELECT * FROM V$LOGFILE;

1.29 TASK 5: VERIFY DATA GUARD CONFIGURATION

1. On the standby database, query the V$ARCHIVED_LOG view to identify existing files in the archived redo log. For example:

   SQL> SELECT SEQUENCE#, FIRST_TIME, NEXT_TIME
   FROM V$ARCHIVED_LOG ORDER BY SEQUENCE#;

2. On the primary database, issue the following SQL statement to force a log switch and archive the current online redo log file group:

   SQL> ALTER SYSTEM ARCHIVE LOG CURRENT;

3. On the standby database, query the V$ARCHIVED_LOG view to verify that the redo data was received and archived on the standby database:

   SQL> SELECT SEQUENCE#, FIRST_TIME, NEXT_TIME, APPLIED
   FROM V$ARCHIVED_LOG ORDER BY SEQUENCE#;

4. Check On the primary and standby servers at the V$ARCHIVE_DEST & alert.log files for any errors

Refer the following URL for more information --
http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

1.29.1 Configure RAC & Data Guard with Oracle GRID Environment

At this stage the Oracle GRID control has been installed in an independent server and Oracle Data Guard manually setup on both the RAC Server clusters..

The challenge now is to make Oracle GRID control aware that we have two RAC database operating in a Physical Standby mode. The other requirement is to set up the observer server.

https://dallas.dell.com:1159/em/
Login ---
username - sysman
password - manager1

The GUI screen will step through the configuration process. Access the Data Guard Web pages through the Oracle Enterprise Manager Grid Control using the following steps:

1. Click the Targets tab to go to the Targets page.(The database must have been registered)

2. Click Databases to go to the Databases page.
3. On the Databases page, you see a listing of all discovered databases. In this scenario, the primary database, FREMONT, has already been discovered. Click the FREMONT database to go to the primary database home page.
4. Click Maintenance.
5. In the High Availability section under Data Guard, click Setup and Manage and log in.

-Targets --> Databases -->ZLS ---> Maintenance --> Setup and Manage (Data Guard)

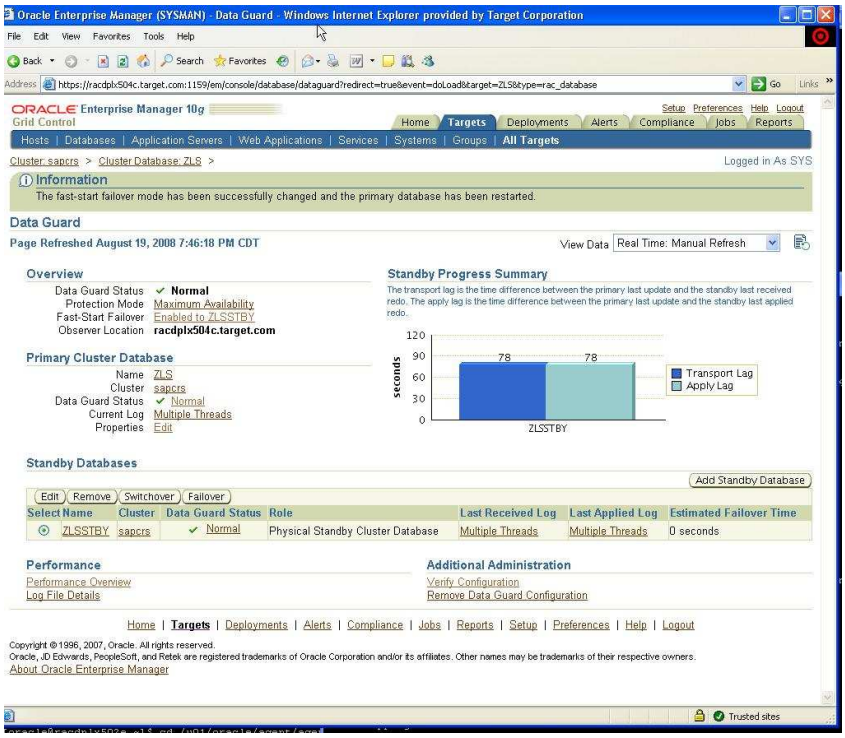

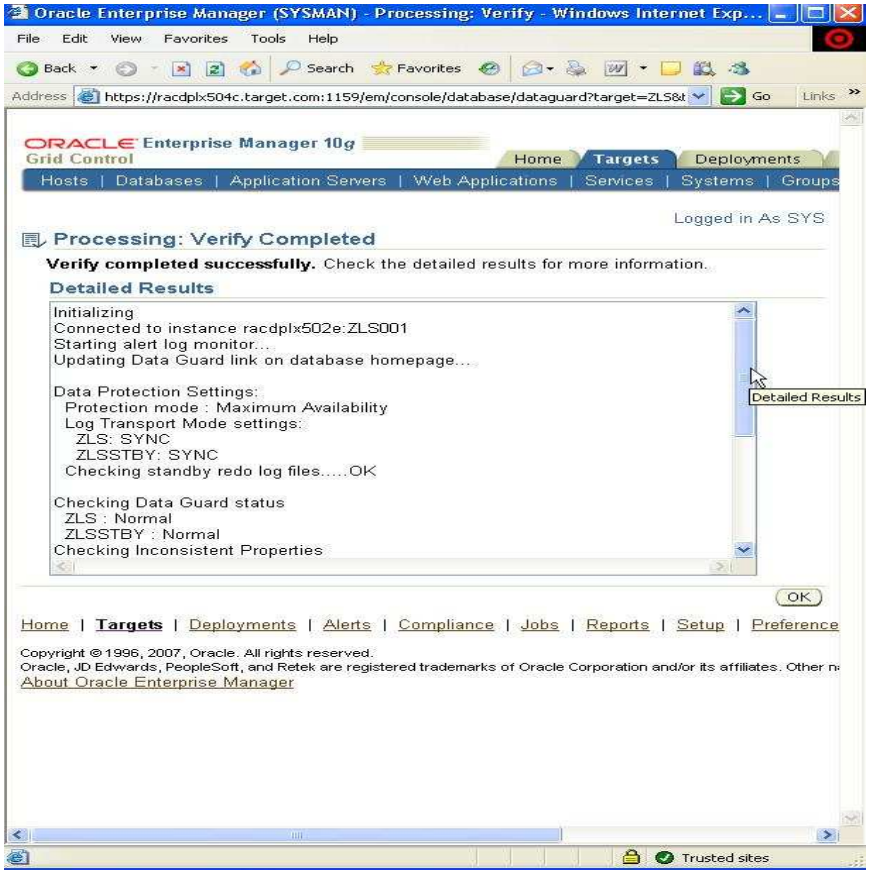

If the primary database is not already in a broker configuration, clicking Setup and Manage in Step 5 will go to the page shown in Figure . This screenshot shows information presented on the All Targets tab, containing graphs for all targets discovered for the cluster. From the list of managed targets, you can quickly assess the availability of the targets and click any graph to drill down for more information. The ***observer*** can be setup on a completely independent server to enable fast failover or switchover.

Note: Oracle Enterprise Manager Grid Control provides access to the complete set of Data Guard features. Limited monitoring-only Data Guard functionality is available in Enterprise Manager Database Control.

On the Data Guard Overview page, you can:

- Edit the protection mode—click Protection Mode in the Data Guard Overview section to view the current protection mode setting and, if necessary, change the protection mode.
- Enable Fast-Start Failover—click Disabled to invoke the Fast-Start Failover wizard that will guide you through the process of **enabling fast-start failover and the observer.**
- View a summary of standby progress—this chart shows the amount of data that each standby database has not yet received and applied.
- Retrieve information about the primary database—click Name, Host, Status, Current Log, or Properties (Edit) to view pertinent information about the primary database.

If you click Status or Edit, you are taken to the Edit Properties page where you can view and change the current state or properties of the primary database.

- View or change information on the standby databases:
  - **Add a standby database to the broker configuration—click Add Standby Database to add a physical or logical standby database.**
  - Change the state or properties—click Edit to go to the Edit Properties page to view and change the current state or properties of the standby database.
  - Discontinue Data Guard broker control—click Remove to remove the selected standby database from Data Guard broker control.
  - Switch the role from standby to primary—click Switchover to switch the database roles from standby to primary.
  - Transition the standby database to the role of primary database— click Failover when a catastrophic failure occurs on the primary database, and there is no possibility of recovering the primary database in a timely manner. The target standby database assumes the primary role, and the failed primary database is disabled by the broker.
- Perform additional administrative activities:
  - Click Verify to check the protection mode and properties, confirm that standby redo log files are present, and verify log switch.
  - Click Remove Data Guard Configuration to remove the broker configuration from Data Guard broker control.

For more information, please refer - Configuring Oracle Data Guard Broker (B14230-02)

1.29.2 Setup Oracle Network Layer for the SAP Application servers

The files that need to be modified are the listener.ora and the tnsnames.ora. The files are included as an attachment

Please refer to Appendix B & C for details on the implementation.

Basic Validation Testing

The following RAC verification checks should be performed on all Oracle RAC nodes in the cluster! Please note these examples are generic, use your installation specific parameters.

This section provides several srvctl commands and SQL queries you can use to validate your Oracle RAC 10*g* configuration.

There are five node-level tasks defined for SRVCTL:

- Adding and deleting node-level applications
- Setting and unsetting the environment for node-level applications
- Administering node applications
- Starting and stopping a group of programs that includes virtual IP addresses, listeners, Oracle Notification Services, and Oracle Enterprise Manager agents (for maintenance purposes).

Status of all instances and services

```
srvctl status database -d TEST
Instance TEST1 is running on node node1
Instance TEST2 is running on node node2
```

Status of a single instance

```
srvctl status instance -d TEST -i TEST2
Instance TEST2 is running on node node2
```

Status of a named service globally across the database

```
srvctl status service -d TEST -s TST_TAF
Service tst_taf is running on instance(s) TEST2, TEST1
```

Status of node applications on a particular node

```
srvctl status nodeapps -n linux1
VIP is running on node: linux1
GSD is running on node: linux1
Listener is running on node: linux1
ONS daemon is running on node: linux1
```

List all configured databases

```
srvctl config database
TEST
```

Display configuration for our RAC database

```
srvctl config database -d TEST
node1 TEST1 /opt/oracle/product/10.2/db_1
node2 TEST2 /opt/oracle/product/10.2/db_1
```

Display all services for the specified cluster database

```
srvctl config service -d TEST
orcl_taf PREF: orcl2 orcl1 AVAIL:
```

Display the configuration for node applications - (VIP, GSD, ONS, Listener)

```
srvctl config nodeapps -n node2 -a -g -s -l
VIP exists.: /node2-vip/10.1.1.214/255.255.255.0/eth0
GSD exists.
ONS daemon exists.
Listener exists.
```

References

- **Setup Guide: Creating a RAC Logical Standby Database for a RAC Primary Database**

- "Switchover and Failover Best Practices: Oracle Data Guard 10g Release 2" http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_FastStartFailoverBestPractices.pdf

- "Fast-Start Failover Best Practices: Oracle Data Guard 10g Release 2" http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_FastStartFailoverBestPractices.pdf

- Oracle Maximum Availability Architecture http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

- Oracle Database High Availability Best Practices (Part #B25159) http://otn.oracle.com/pls/db102/db102.to_toc?partno=b25159

- Workload Management with Oracle Real Application Clusters **http://www.oracle.com/technology/products/database/clustering/pdf/twpracwkldmgmt.pdf**

- **MAA 10g Setup Guide: Creating a RAC Physical Standby Database for a RAC Primary Database**

- **Configuration of SAP NetWeaver for Oracle Database 10g Real Application Clusters** http://www.oracle.com/newsletters/sap/products/rac/rac4sap_howto.html

- Oracle Data Guard http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardOverview.html

- Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide (Part #14197) http://download-west.oracle.com/docs/cd/B19306_01/rac.102/b14197/toc.htm

- Oracle Data Guard Broker (Part #B14230) http://otn.oracle.com/pls/db102/db102.to_toc?partno=b14230

Please also refer various **sapnotes** from www.sap.com

Appendix A – Basic Scripts

Use the following scripts to automate many of the installation steps needed in this documentation.

This script automates a number of RPM and other OS-level installations and checking.

```
#! /bin/bash

export ORACLE_STAGE=/opt/stage

unzip linux.x64_10gr2_clusterware.zip
unzip linux.x64_10gr2_database.zip
unzip linux.x64_10g2_examples.zip
unzip rda*

rpm -Uv elfutils-libelf-devel-0.97.1-5.x86_64.rpm
rpm -Uv elfutils-devel-0.97.1-5.x86_64.rpm

rpm -qa --qf '%{NAME}-%{VERSION}-%{RELEASE}  (%{ARCH})\n'|grep compat-libstdc++
rpm -qa --qf '%{NAME}-%{VERSION}-%{RELEASE}  (%{ARCH})\n'|grep ocfs

rpm -q --qf '%{NAME}-%{VERSION}-%{RELEASE}  (%{ARCH})\n' \
  binutils \
  compat-db \
  control-center \
  elfutils \
  elfutils-libelf \
  elfutils-libelf-devel \
  gcc \
  gcc-c++ \
  glibc \
  glibc-common \
  glibc-devel \
  gnome-libs \
  libaio \
  libaio-devel \
  libgcc \
  libstdc++ \
  libstdc++-devel \
  make \
  pdksh \
  sysstat \
  xscreensaver
```

setupBasicOracle.sh

```
#! /bin/bash

# This first part sets up the basic Oracle stuff needed, regardless
# of version of Red Hat.

# Create the default Oracle group
groupadd -g 501 dba

# Create the Oracle account
useradd -g 501 -u 501 -m -c 'Oracle User Account' oracle
usermod -G dba,wheel oracle
passwd oracle

# Create the ORACLE_BASE and CRS_BASE directories
mkdir /oracle
mkdir /oracleCRS
chown oracle:dba /oracle
chown oracle:dba /oracleCRS

# set the needed kernel parameters
cat >> /etc/sysctl.conf <<EOF
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.file-max = 76800
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default=4194304
net.core.wmem_default=262144
net.core.rmem_max=4194304
net.core.wmem_max=262144
EOF

# Set security limits
cat >> /etc/pam.d/login <<EOF
session required pam_limits.so
EOF

# Set process limits
cat >> /etc/security/limits.conf <<EOF
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
EOF

# Create default profile settings for ulimits
cat >> /etc/profile <<EOF
if [ \$USER = "oracle" ]; then
 if [ \$SHELL = "/bin/ksh" ]; then
 ulimit -p 16384
```

```
 ulimit -n 65536
 else
 ulimit -u 16384 -n 65536
 fi
 umask 022
fi
EOF

# Install hangcheck at reboot
cat >> /etc/rc.d/rc.local << EOF
modprobe hangcheck-timer hangcheck_tick=30 hangcheck_margin=180
EOF
```

1.30 Starting Oracle User bash Shell Profile

This is a good starting point for .bash_profile for the oracle user, copy this snippet to the end of the existing oracle user $HOME/.nash_profile file.

```
export ORACLE_BASE=/opt/oracle
export CRS_BASE=/opt/oracleCRS
export ORACLE_HOME=$ORACLE_BASE/product/10.2/db_1
export AGENT_HOME=$ORACLE_BASE/product/10.2/agent10g
export CRS_HOME=$CRS_BASE/product/10.2/crs_1
export ORACLE_TERM=xterm
export ORACLE_SID=TEST
export ORACLE_LISTENERS=''
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/use/local/lib::/lib:/usr/lib:/usr/local/lib
export
CLASSPATH=$ORACLE_HOME/JR:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network.jl
ib
export TEMP=/tmp
export TMPDIR=/tmp
PATH=$ORACLE_HOME/bin:$CRS_HOME/bin:/usr/local/bin:$PATH
export PATH
#
# Basic Oracle Stuff
#
alias oh='cd $ORACLE_HOME'
alias ob='cd $ORACLE_BASE'
alias sql='sqlplus "/ as sysdba"'
alias lstat='lsnrctl status ${ORACLE_LISTENER}'
alias lstart='lsnrctl start ${ORACLE_LISTENER}'
alias lstop='lsnrctl stop ${ORACLE_LISTENER}'
#
# CRS Related
#
alias crs='${CRS_HOME}/bin/crs_stat -t'
#
# Agent Related
#
alias astatus='${AGENT_HOME}/bin/emctl status'
alias astart='${AGENT_HOME}/bin/emctl start agent'
alias astop='${AGENT_HOME}/bin/emctl stop agent'
EOF
```

Appendix B : SAP J2EE Server Configuration for RAC Database

This section is valid for Java Add-In installations (JAVA + ABAP stack) as well as for Java Standalone installations of SAP NetWeaver.

The JVM (Java Virtual Machine) by SAP uses the thin JDBC driver from Oracle to connect to the database. Therefore the Oracle network configuration for the Java stack is different from the ABAP stack. With the thin JDBC driver the network configuration files tnsnames.ora, sqlnet.ora and listener.ora are not used. The URL which defines the connect description is stored in an encrypted flat file in the filesystem. The content of this file is also stored in the database. The content of this configuration file is valid for all Java Server Engines in an SAP system. The configuration file is shared via the /sapmnt access path.

The path to the configuration tool in a mixed ABAP+JAVA installation is /usr/sap/ZLS/CI/j2ee/configtool.

As user **ZLSADM**, go to this directory and start the configuration tool configtool.sh from there:

sapappsrvr01:zlsadm 7> pwd
/sap/usersap/ZLS/DVEBMGS00/j2ee/configtool

sapappsrvr01:zlsadm 10> ./configtool.sh

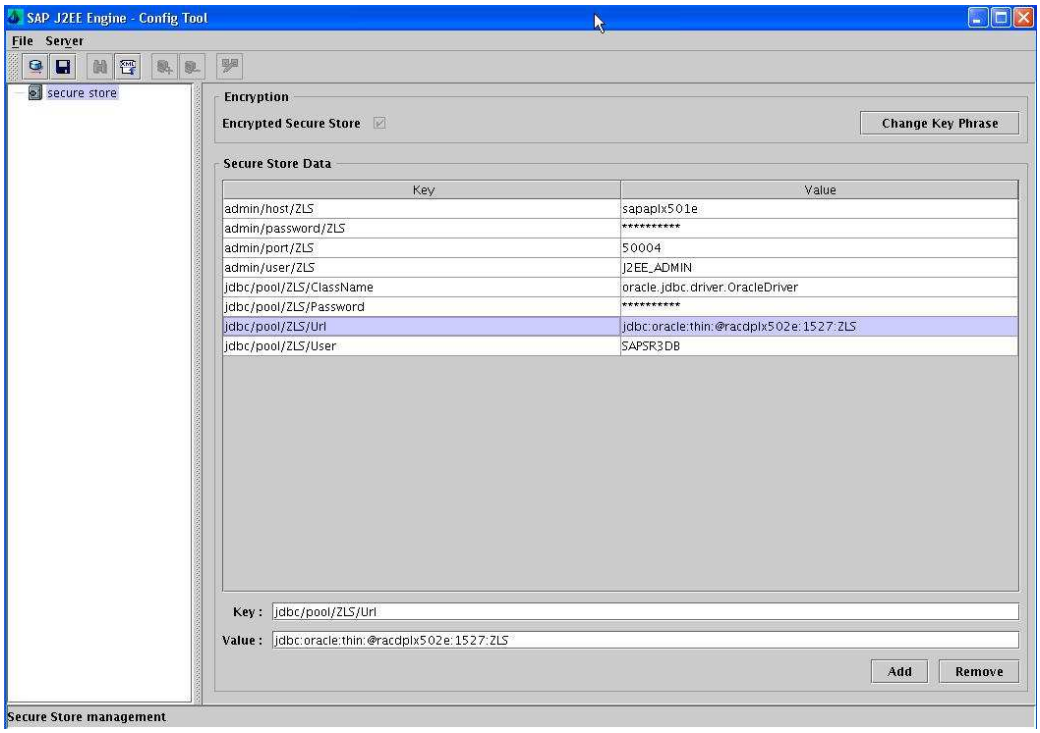

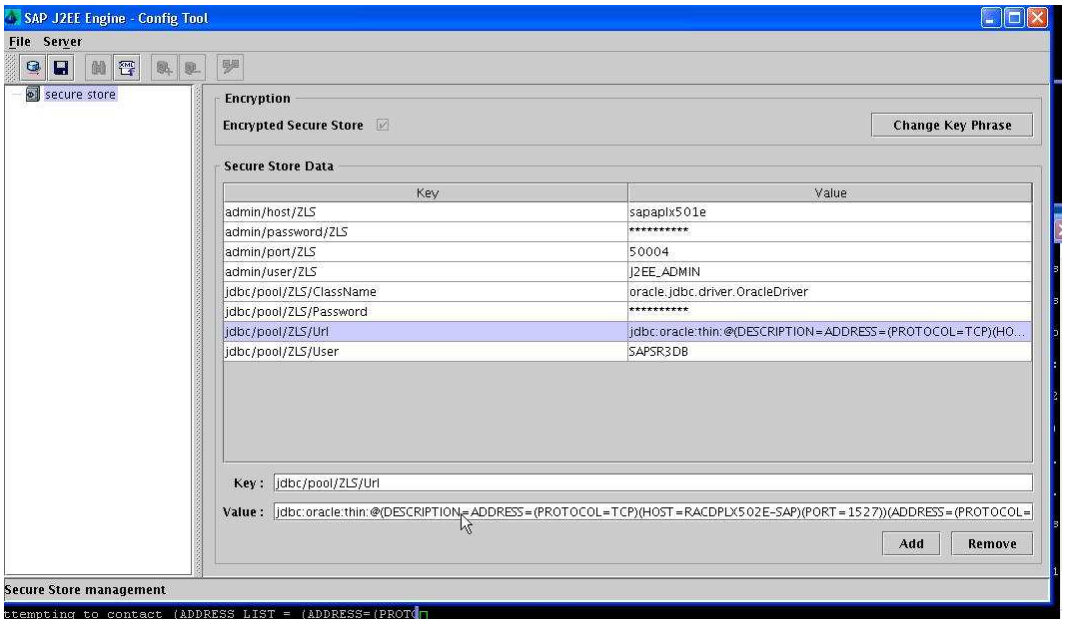

Modify the Key – ***/jdbc/pool/ZLS/URL*** : Cut and paste the TNSNAMES DESCRIPTION from the tnsnames file created for ABAP server to address all the RAC servers.

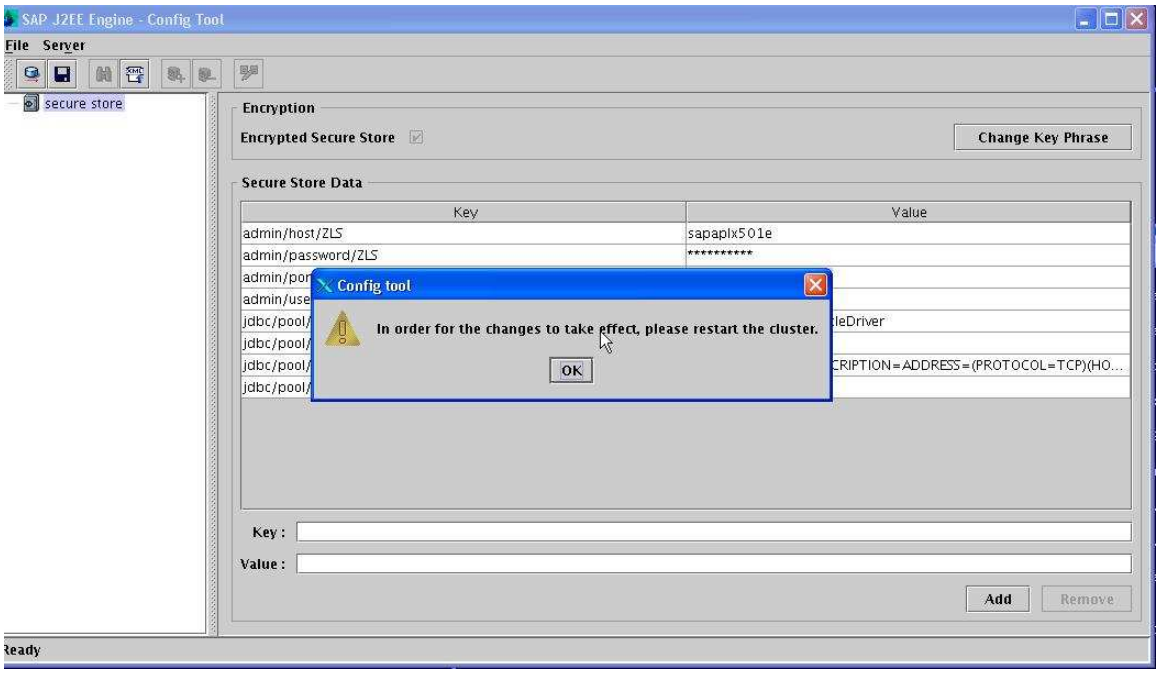
SAP J2EE Engine - Config Tool
File Server
secure store
Encryption
Encrypted Secure Store
Change Key Phrase
Secure Store Data
Key
Value
admin/host/ZLS
sapaplx501e
admin/password/ZLS
****
Config tool
In order for the changes to take effect, please restart the cluster.
OK
Key :
Value :
Add
Remove
Ready

Appendix C : SAP ABAP Server Configuration for RAC Database

Linking tnsnames.ora File to TNS_ADMIN Location for Instant Client

From Oracle Release 10g on, SAP software for the ABAP stack uses the instant client libraries for database connections. An instant client implementation does not require a $ORACLE_HOME directory. The location of the network configuration file tnsnames.ora being used by the client is obtained from the environment variable TNS_ADMIN.

SAP uses environment variable TNS_ADMIN to specify the location of tnsnames.ora. The value of TNS_ADMIN may differ on SAP application server nodes outside of the cluster. Copy the file tnsnames.ora to all SAP application servers to the location given in the environment variable TNS_ADMIN for user **ZLSADM** on those servers.

For SAP application instances running within the cluster, you can build a softlink of the value of $TNS_ADMIN to the shared Oracle Home directory holding the network configuration.

#cd /sapmnt/ZLS/profile/oracle

sapappsrvr01:zlsadm 1> echo $TNS_ADMIN
/sapmnt/ZLS/profile/oracle

The specific about the tnsnames.ora file and listener.ora are taken from the DB server. (Attachment)

Appendix D : Database Install Using the SAP Solution Manager

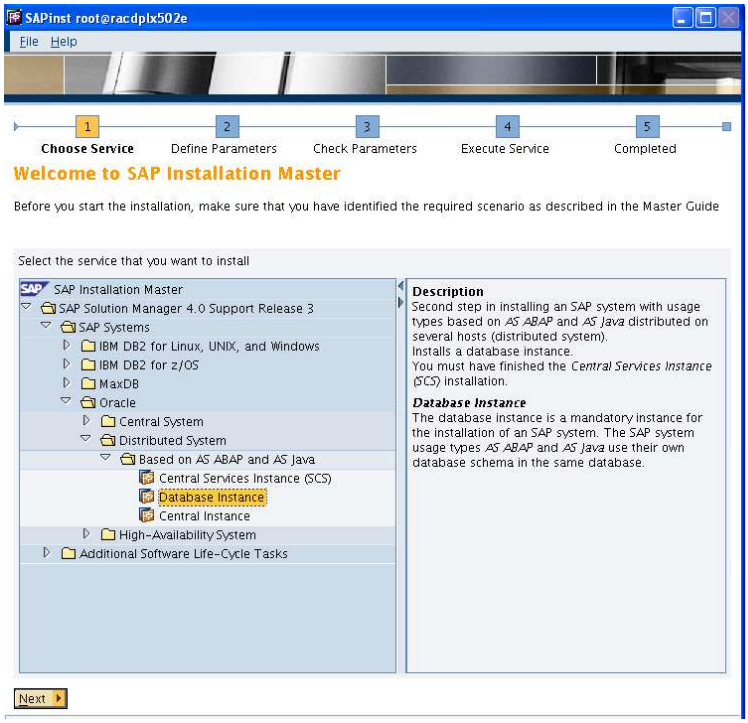

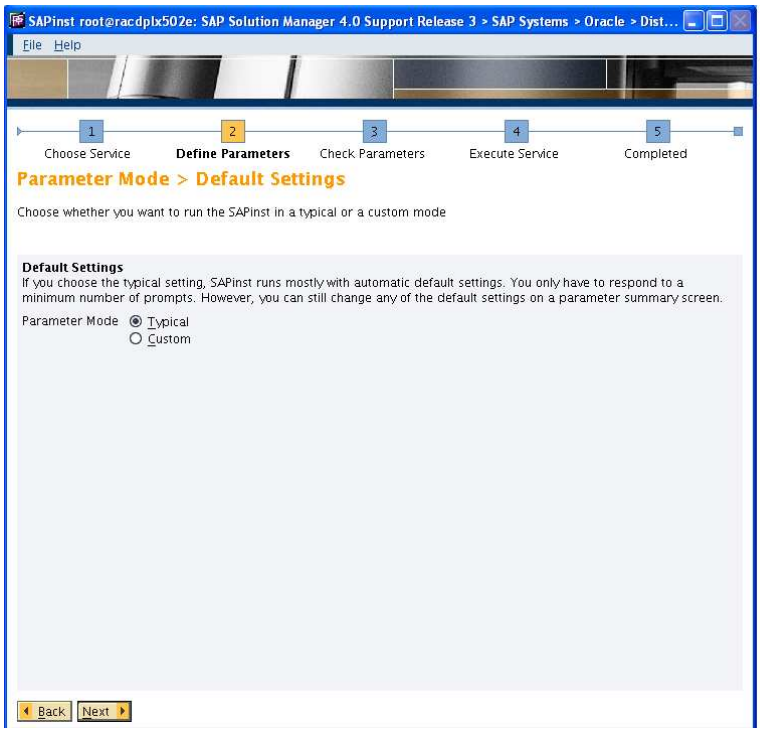
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Parameter Mode > Default Settings
Choose whether you want to run the SAPinst in a typical or a custom mode
Default Settings
If you choose the typical setting, SAPinst runs mostly with automatic default settings. You only have to respond to a minimum number of prompts. However, you can still change any of the default settings on a parameter summary screen.
Parameter Mode
Typical
Custom
Back
Next

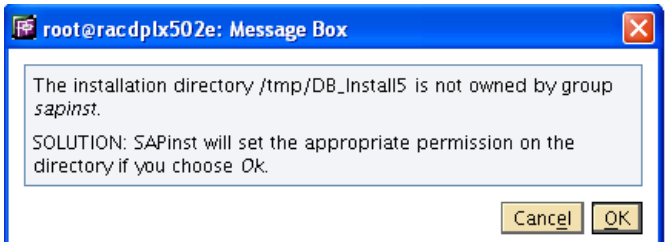
root@racdplx502e: Message Box
The installation directory /tmp/DB_Install5 is not owned by group sapinst.
SOLUTION: SAPinst will set the appropriate permission on the directory if you choose Ok.
Cancel
OK

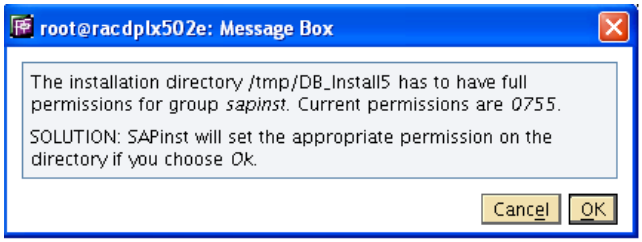
root@racdplx502e: Message Box
The installation directory /tmp/DB_Install5 has to have full permissions for group sapinst. Current permissions are 0755.
SOLUTION: SAPinst will set the appropriate permission on the directory if you choose Ok.
Cancel
OK


SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Media Browser > Software Package Request
Enter the location of the required software packages
Software Package(s)
Media Name
Package Location
Copy Package To
Java Component SOLMAN40SR3
/usr/sapsoftware/SOLMAN_JAVA_COMP
Browse...
Additional Information
SAPinst will detect the required software packages on the media and check the corresponding package identification files LABEL.ASC.
If you want to copy the media to your local disk, enter the target location in the Copy Package To column.
Cancel
OK

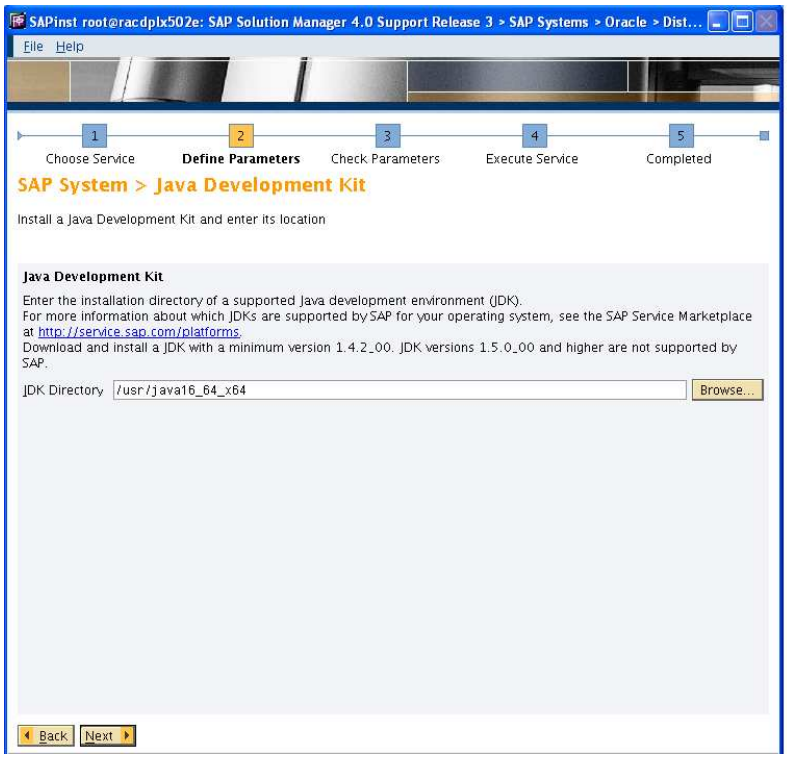
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
SAP System > Java Development Kit
Install a Java Development Kit and enter its location
Java Development Kit
Enter the installation directory of a supported Java development environment (JDK).
For more information about which JDKs are supported by SAP for your operating system, see the SAP Service Marketplace at http://service.sap.com/platforms.
Download and install a JDK with a minimum version 1.4.2_00. JDK versions 1.5.0_00 and higher are not supported by SAP.
JDK Directory
/usr/java16_64_x64
Browse...
Back
Next

Cancel

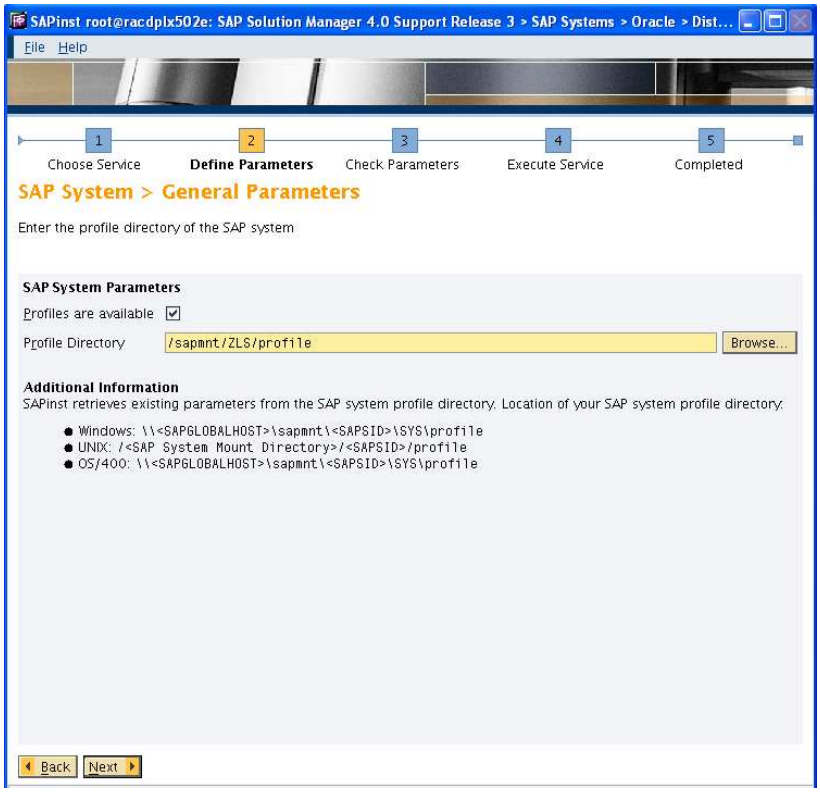
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
SAP System > General Parameters
Enter the profile directory of the SAP system
SAP System Parameters
Profiles are available
Profile Directory
/sapmnt/ZLS/profile
Browse...
Additional Information
SAPinst retrieves existing parameters from the SAP system profile directory. Location of your SAP system profile directory:
Windows: \\<SAPGLOBALHOST>\sapmnt\<SAPSID>\SYS\profile
UNIX: /<SAP System Mount Directory>/<SAPSID>/profile
OS/400: \\<SAPGLOBALHOST>\sapmnt\<SAPSID>\SYS\profile
Back
Next

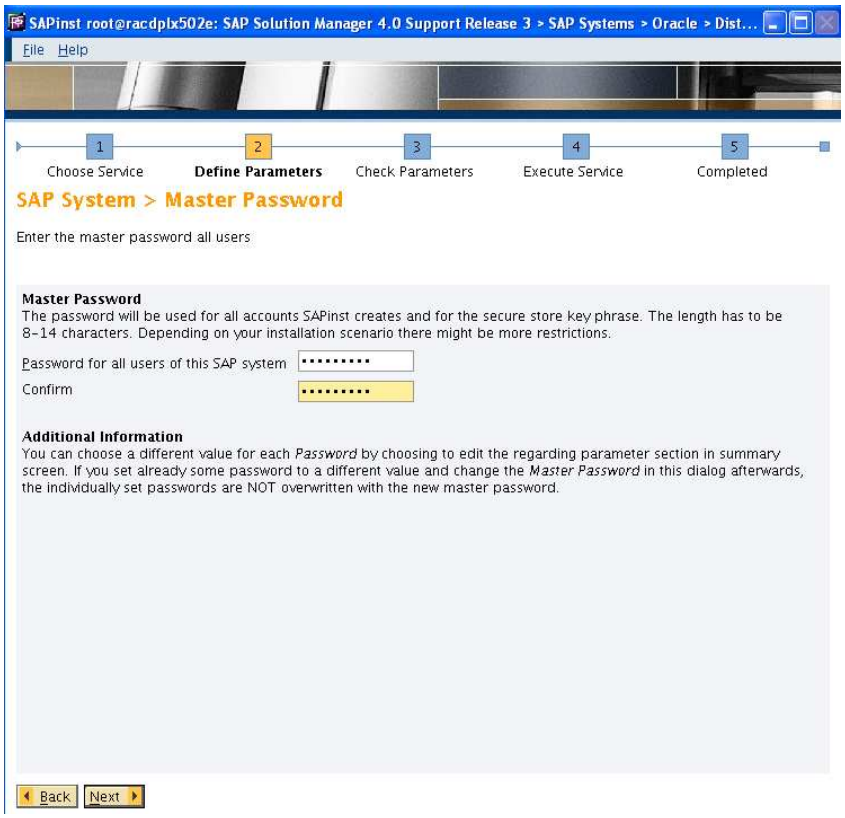
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
SAP System > Master Password
Enter the master password all users
Master Password
The password will be used for all accounts SAPinst creates and for the secure store key phrase. The length has to be 8-14 characters. Depending on your installation scenario there might be more restrictions.
Password for all users of this SAP system
Confirm
Additional Information
You can choose a different value for each Password by choosing to edit the regarding parameter section in summary screen. If you set already some password to a different value and change the Master Password in this dialog afterwards, the individually set passwords are NOT overwritten with the new master password.
Back
Next

sap@dell

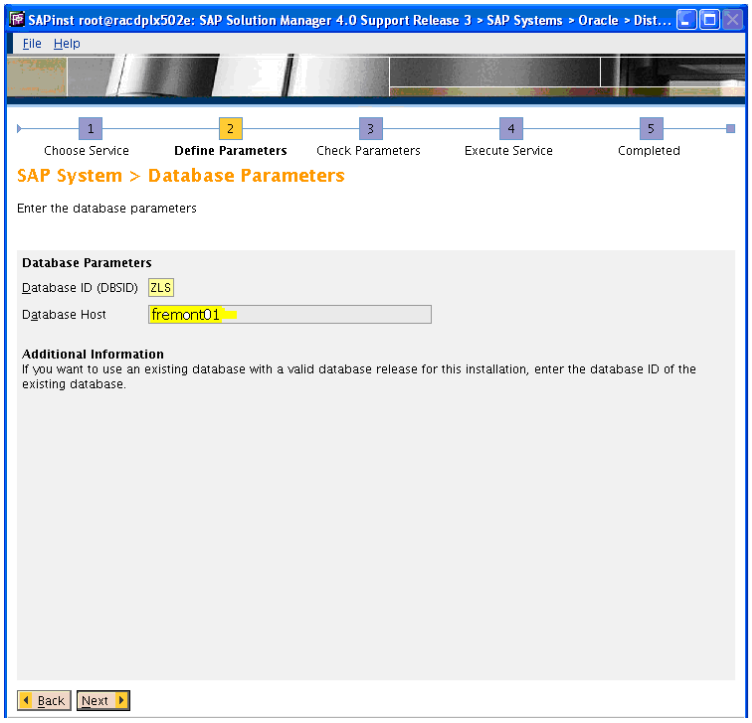
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
SAP System > Database Parameters
Enter the database parameters
Database Parameters
Database ID (DBSID) ZLS
Database Host fremont01
Additional Information
If you want to use an existing database with a valid database release for this installation, enter the database ID of the existing database.
Back
Next

Cancel

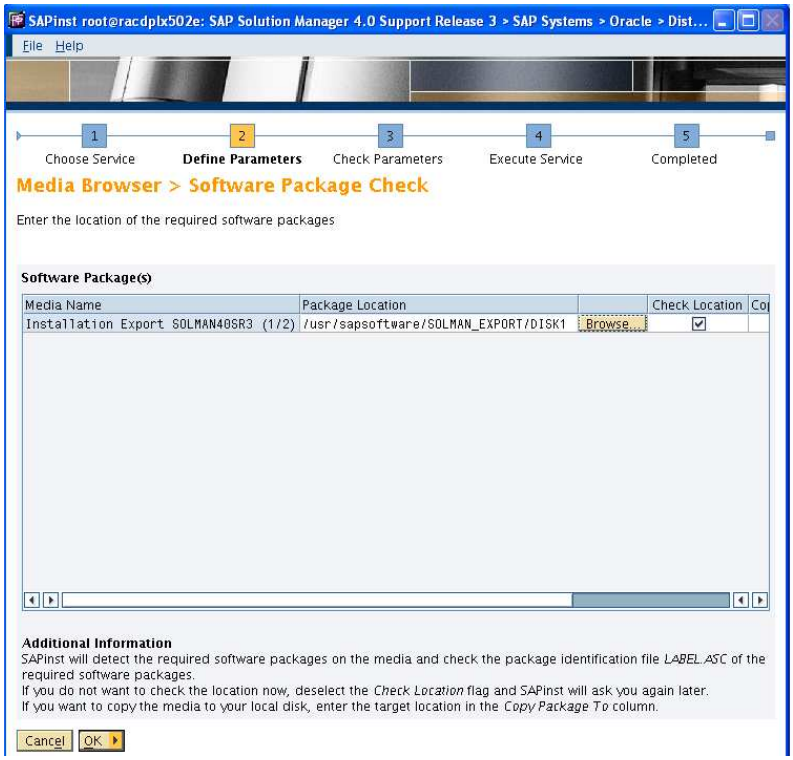
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Media Browser > Software Package Check
Enter the location of the required software packages
Software Package(s)
Media Name
Package Location
Check Location
Installation Export SOLMAN40SR3 (1/2)
/usr/sapsoftware/SOLMAN_EXPORT/DISK1
Browse...
Additional Information
SAPinst will detect the required software packages on the media and check the package identification file LABEL.ASC of the required software packages.
If you do not want to check the location now, deselect the Check Location flag and SAPinst will ask you again later.
If you want to copy the media to your local disk, enter the target location in the Copy Package To column.
Cancel
OK

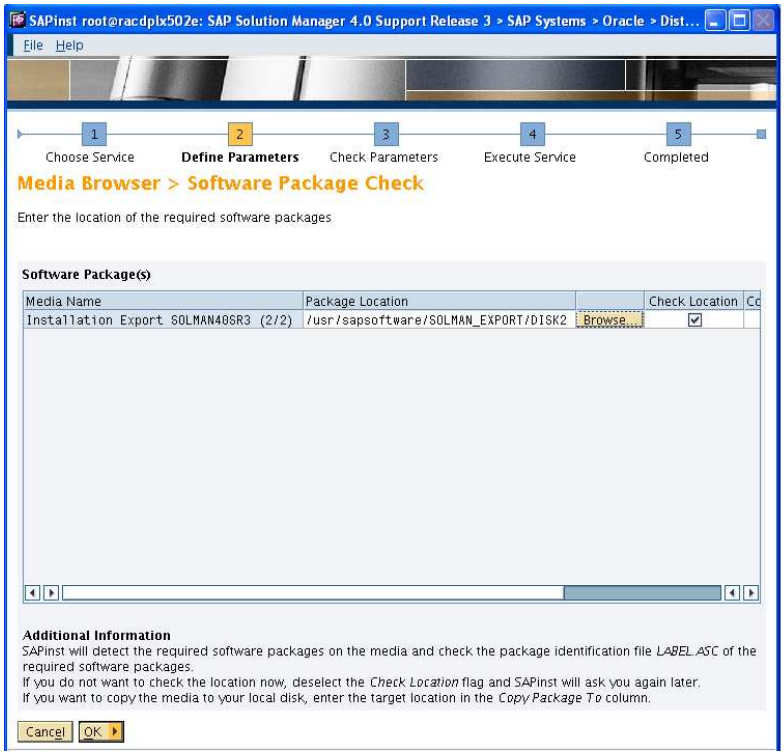
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Media Browser > Software Package Check
Enter the location of the required software packages
Software Package(s)
Media Name
Package Location
Check Location
Installation Export SOLMAN40SR3 (2/2)
/usr/sapsoftware/SOLMAN_EXPORT/DISK2
Browse...
Additional Information
SAPinst will detect the required software packages on the media and check the package identification file LABEL.ASC of the required software packages.
If you do not want to check the location now, deselect the Check Location flag and SAPinst will ask you again later.
If you want to copy the media to your local disk, enter the target location in the Copy Package To column.
Cancel
OK

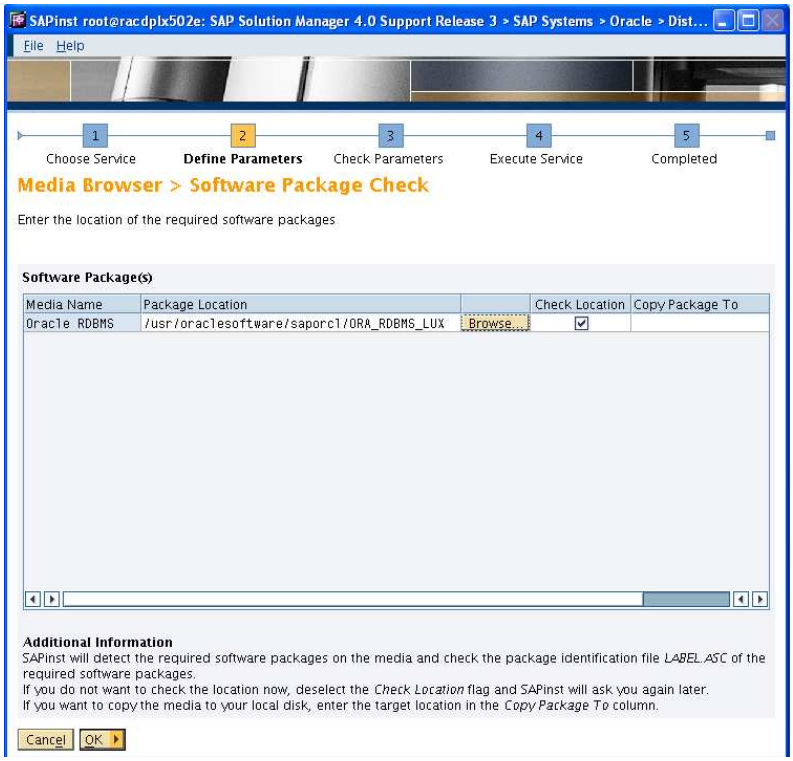
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Media Browser > Software Package Check
Enter the location of the required software packages
Software Package(s)
Media Name
Package Location
Check Location
Copy Package To
Oracle RDBMS
/usr/oraclesoftware/saporcl/ORA_RDBMS_LUX
Browse...
Additional Information
SAPinst will detect the required software packages on the media and check the package identification file LABEL.ASC of the required software packages.
If you do not want to check the location now, deselect the Check Location flag and SAPinst will ask you again later.
If you want to copy the media to your local disk, enter the target location in the Copy Package To column.
Cancel
OK

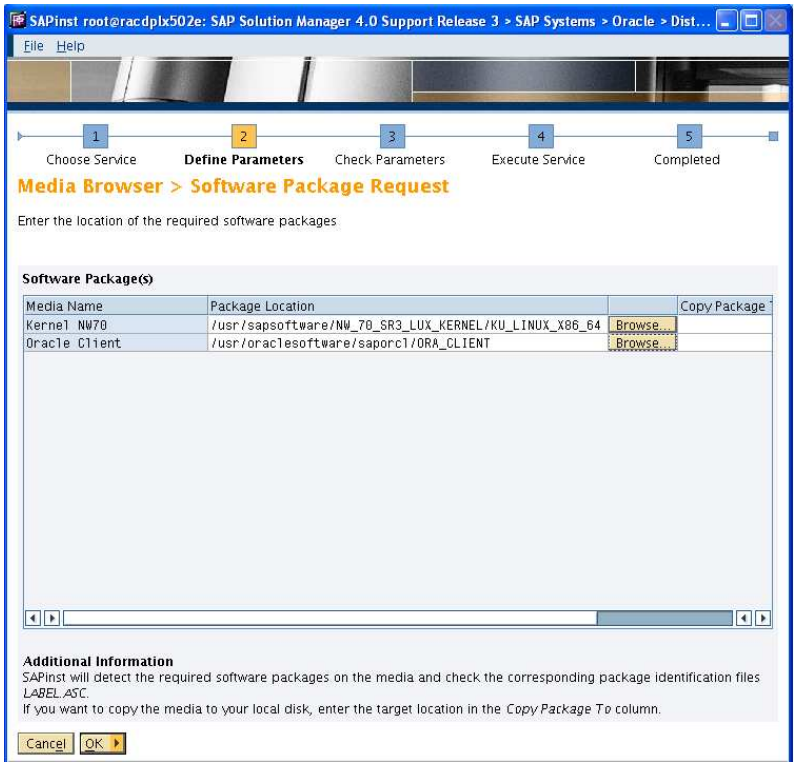
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Media Browser > Software Package Request
Enter the location of the required software packages
Software Package(s)
Media Name
Package Location
Copy Package
Kernel NW70
/usr/sapsoftware/NW_70_SR3_LUX_KERNEL/KU_LINUX_X86_64
Browse...
Oracle Client
/usr/oraclesoftware/saporcl/ORA_CLIENT
Browse...
Additional Information
SAPinst will detect the required software packages on the media and check the corresponding package identification files LABEL.ASC.
If you want to copy the media to your local disk, enter the target location in the Copy Package To column.
Cancel
OK

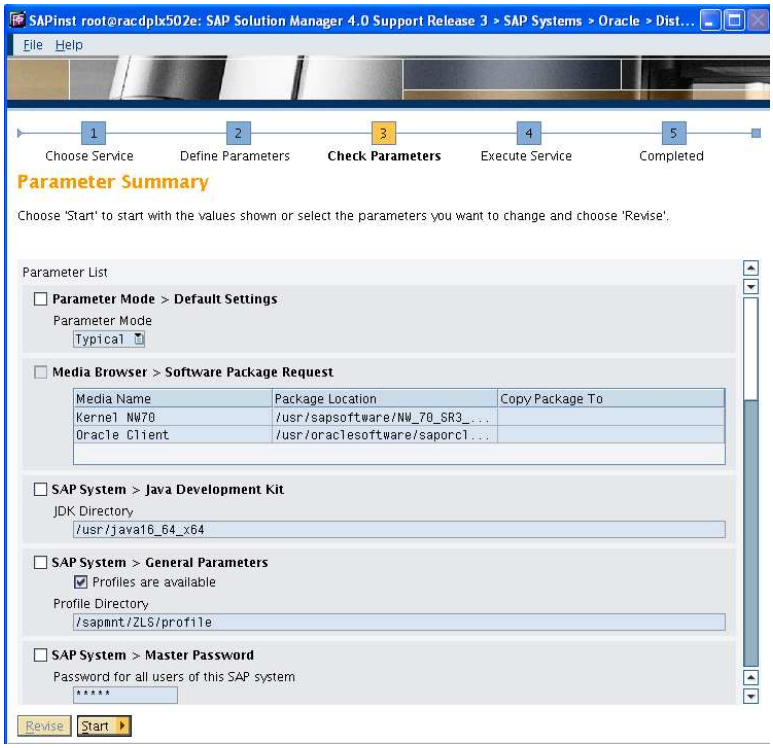
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Parameter Summary
Choose 'Start' to start with the values shown or select the parameters you want to change and choose 'Revise'.
Parameter List
Parameter Mode > Default Settings
Parameter Mode
Typical
Media Browser > Software Package Request
Media Name
Package Location
Copy Package To
Kernel NW70
/usr/sapsoftware/NW_70_SR3_...
Oracle Client
/usr/oraclesoftware/saporcl...
SAP System > Java Development Kit
JDK Directory
/usr/java16_64_x64
SAP System > General Parameters
Profiles are available
Profile Directory
/sapmnt/ZLS/profile
SAP System > Master Password
Password for all users of this SAP system
*****
Revise
Start

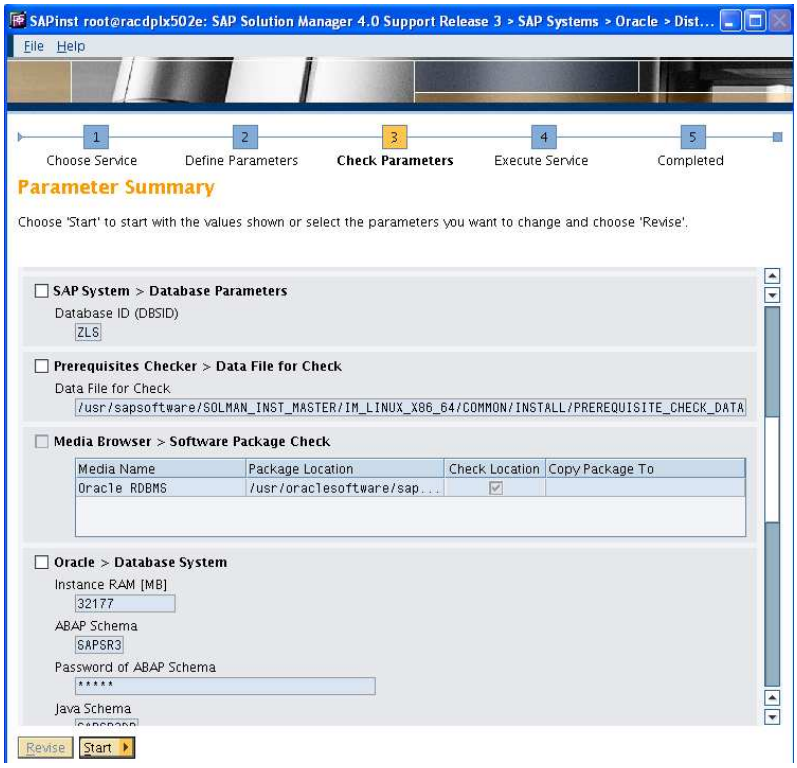
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Parameter Summary
Choose 'Start' to start with the values shown or select the parameters you want to change and choose 'Revise'.
SAP System > Database Parameters
Database ID (DBSID)
ZLS
Prerequisites Checker > Data File for Check
Data File for Check
/usr/sapsoftware/SOLMAN_INST_MASTER/IM_LINUX_X86_64/COMMON/INSTALL/PREREQUISITE_CHECK_DATA
Media Browser > Software Package Check
Media Name
Package Location
Check Location
Copy Package To
Oracle RDBMS
/usr/oraclesoftware/sap...
Oracle > Database System
Instance RAM [MB]
32177
ABAP Schema
SAPSR3
Password of ABAP Schema
*****
Java Schema
Revise
Start

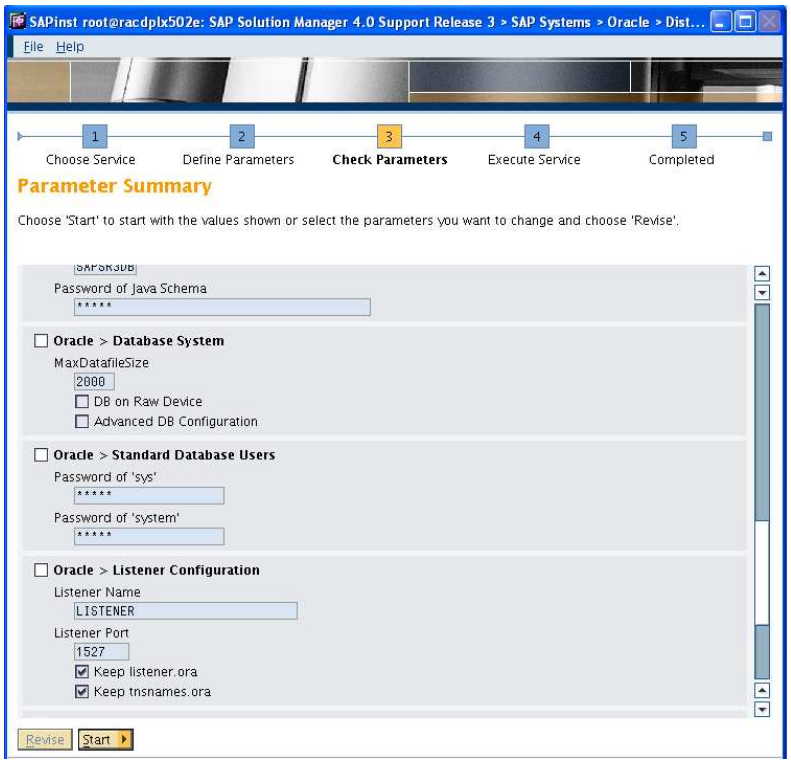
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Parameter Summary
Choose 'Start' to start with the values shown or select the parameters you want to change and choose 'Revise'.
Password of Java Schema
*****
Oracle > Database System
MaxDatafileSize
2000
DB on Raw Device
Advanced DB Configuration
Oracle > Standard Database Users
Password of 'sys'
*****
Password of 'system'
*****
Oracle > Listener Configuration
Listener Name
LISTENER
Listener Port
1527
Keep listener.ora
Keep tnsnames.ora
Revise
Start

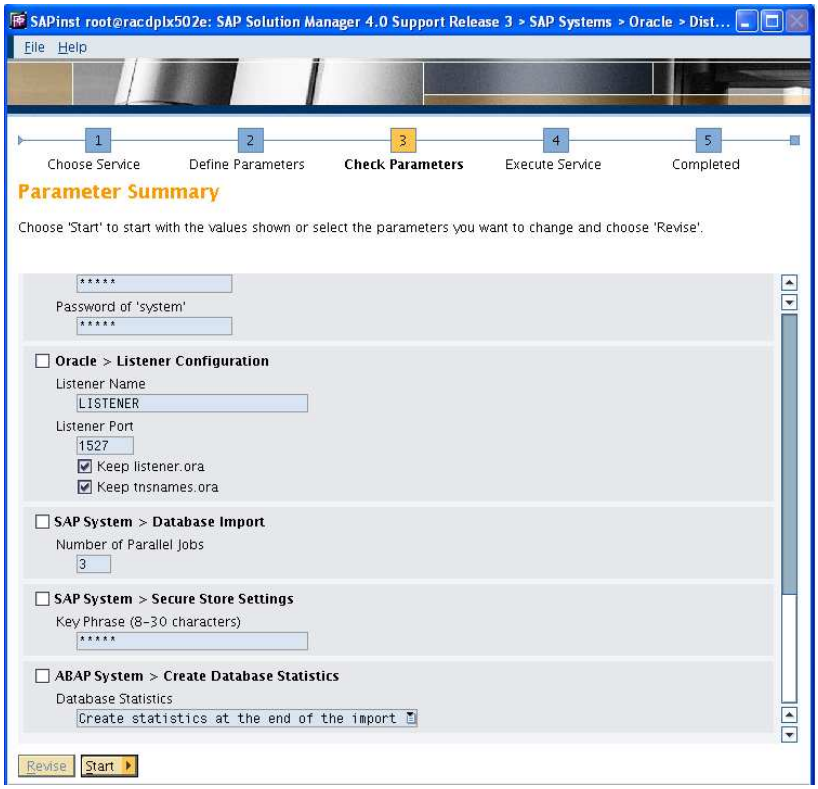
SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Parameter Summary
Choose 'Start' to start with the values shown or select the parameters you want to change and choose 'Revise'.
Password of 'system'
Oracle > Listener Configuration
Listener Name
LISTENER
Listener Port
1527
Keep listener.ora
Keep tnsnames.ora
SAP System > Database Import
Number of Parallel Jobs
3
SAP System > Secure Store Settings
Key Phrase (8-30 characters)
ABAP System > Create Database Statistics
Database Statistics
Create statistics at the end of the import
Revise
Start


SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Task Progress
Task is running
Phase 1 of 22
Phase List
Create users for SAP system
Install common system files
Unpack SAP archives
Create and load database
Install/Check database
Evaluate database parameters
Install database server software
Install Oracle client software
Configure Oracle server network
Prepare database server configuration
Create database
Create/Check tablespaces
Check Oracle NLS_CHARACTERSET
Create database schema
Check Oracle patch
Install DBA tools
Import ABAP
Create secure store
Import Java dump
The phase is in progress.

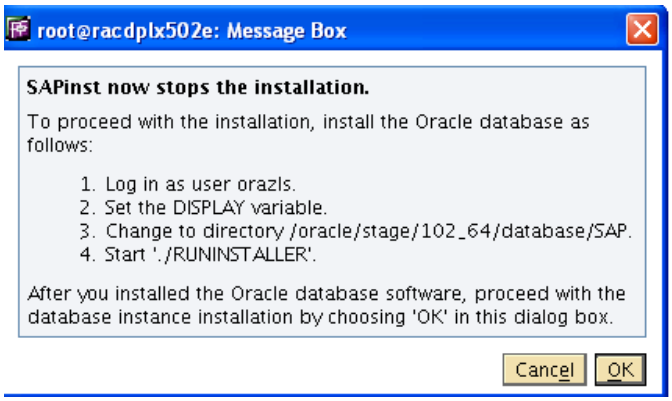
root@racdplx502e: Message Box
SAPinst now stops the installation.
To proceed with the installation, install the Oracle database as follows:
1. Log in as user orazls.
2. Set the DISPLAY variable.
3. Change to directory /oracle/stage/102_64/database/SAP.
4. Start './RUNINSTALLER'.
After you installed the Oracle database software, proceed with the database instance installation by choosing 'OK' in this dialog box.
Cancel
OK


SAPinst root@racdplx502e: SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Dist...
File Help
1 Choose Service
2 Define Parameters
3 Check Parameters
4 Execute Service
5 Completed
Task Progress
The task has been completed successfully.
All phases completed
Phase List
Install/Check database
Evaluate database parameters
Install database server software
Install Oracle client software
Configure Oracle server network
Prepare database server configuration
Create database
Create/Check tablespaces
Check Oracle NLS_CHARACTERSET
Create database schema
Check Oracle patch
Install DBA tools
Import ABAP
Create secure store
Import Java dump
Post load activities
Perform post-load activities
Create database statistics
Execution of Service SAP Solution Manager 4.0 Support Release 3 > SAP Systems > Oracle > Distributed System > Based on AS ABAP and AS Java > Database Instance has been completed successfully
OK
Execution of Service has been completed successfully

DONE!!!

Authors

SAP installation using MAA Architecture: SAP & Oracle Database 10g Release 2

January 2009

Author: Mahesh Pakala

Mahesh Pakala is a highly experienced professional in Information Technology and Media Industry in the Silicon Valley. He has worked in various senior management roles in the industry and is on advisory panel of companies as a technology and business consultant.

He works in GICS (Global Infrastructure Consulting Services) Group of Dell Inc., assisting large enterprise customers with HA architectures and solutions. He has extensive work experience in areas of – engineering, media and technology with companies, such as Oracle Corporation (System Performance Group & RDBMS Kernel Escalations), Ingres (Computer Associates), Fujitsu and startups like eLance and Grand Central Communications. He has been a presenter in the areas of High Availability and Disaster Recovery Solution at previous Oracle Open World.

Contact for more Information – Mahesh_Pakala@dell.com

Reviewed By: Joseph Meeks, Lawrence To, Michael Smith, Lyn Pratt, Shari Yamaguchi, Ron Piwetz