

An Oracle White Paper
July 2011

Using SAP NetWeaver with the Oracle Exadata Database Machine

A Best Practices Guide

1	Preface.....	3
2	ASM Disk Group Recommendations for SAP Databases.....	5
3	Character Set Requirements for SAP Databases.....	6
4	Non-Unicode SAP Installations.....	6
5	SAP Database Administration with BR*Tools.....	6
6	Lifecycle Management for SAP Databases.....	7
	6.1 Installation of the OPatch and MOPatch Utilities.....	8
	6.2 Installation of the Oracle Exadata Bundle Patch.....	8
	6.3 Installation of the SAP Bundle Patch for Oracle Exadata. .	8
7	Migration of SAP Databases	9
	7.1 Migration Approach 1: Oracle-to-Oracle (O2O) ACS Service and Customer Self-Service.....	9
	7.2 Migration Approach 2: Oracle-to-Oracle (Triple-O) ACS Service only.....	10
	7.3 Migration Approach 3: Oracle RMAN and Transportable Tablespaces.....	11
8	Shared Filesystems in SAP Environments.....	12
9	Protection of SAP Central Services	13
	9.1 Installation Procedure for SAP Central Services	14
	9.2 Installation Preparation.....	14
	9.3 Setup of SAP Enqueue Replication Service	15
	9.4 Configuration of SAPCTL.....	17
10	Appendix 1:.....	18
	10.1 SAP Oracle_Home Naming Convention.....	18
	10.2 Default Oracle Environment Settings.....	18

Preface

This document explains all the necessary steps to setup an SAP system based on the SAP NetWeaver technology using the Oracle Exadata Database Machine. All SAP products based from SAP NetWeaver 7.0 on are certified to use the Oracle Exadata Database Machine.

The paper describes the required Oracle software environment settings on the database nodes (Appendix 1 lists a working example which should be followed for an Oracle Exadata deployment for SAP), SAP specific database requirements, information on how to install SAP required database patches to the database nodes, suggestions for the implementation of shared filesystems for SAP installations and how to install, configure, manage and control the SAP central services on the database nodes through Oracle Clusterware and its service program SAPCTL.

The Oracle Exadata Database Machine is used for storing the databases of the individual SAP systems. The Oracle Exadata Database Machine cannot be used to run SAP Instances. SAP Instances have to run on separate machines which use the Ethernet or InfiniBand network to exchange data with the database(s) on the Exadata Database Machine. In SAP terminology this is called a three tier architecture. This flexible three tier architecture allows for any combination of hardware and operating systems running the SAP instances to be used with the Oracle Exadata Database Machine. So for instance you can run SAP Application servers on AIX or HP-UX against the Oracle Exadata Database Machine. This flexibility allows an easy introduction of the Oracle Exadata Database Machine in existing SAP environments as it leaves the SAP layer unchanged. On the database nodes of the Oracle Exadata Database Machine you can choose between the Oracle Solaris 11 Express or the Oracle Linux 5 operating system. The only SAP components which are supported to run on the database nodes of the Oracle Exadata Database Machine are the SAP database administration tools (BR*Tools) and the SAP central services (SCS and ASCS). There is currently no support by SAP through SAPINST to install a new SAP system using the Oracle Exadata Database Machine as the database backend. Databases from already installed SAP systems have to be migrated from existing database servers to the Oracle Exadata Database Machine.

No changes to the standard database schema of the SAP database should be done when being migrated to the Oracle Exadata Database Machine. Changes should not be done to the table and/or index design, the partitioning concept or storage attributes of tables, indexes and partitions. The standard schema of the SAP database is very well designed, tested and proven with thousand's of customers. In addition many SAP administration, monitoring and upgrade tasks depend on the standard database schema layout. Any change to the standard SAP database schema therefore has to be discussed with SAP and an SAP support calls should be opened.

SAP Note 1590515 will be updated on a regular base to reflect any changes on using SAP applications with the Oracle Exadata Database Machine.

Overall this documentation complements the existing standard documentation on the Oracle Exadata Database Machine and therefore it is assumed that the reader is familiar with the standard Oracle Exadata documentation.

To understand the requirements and steps outlined in this document it is necessary that the reader is also familiar with the SAP specific support notes and white papers on Oracle RAC (“Configuration of SAP NetWeaver for Oracle Grid Infrastructure 11.2.0.2 and Oracle Real Application Clusters 11g Release 2: A Best Practices Guide”), Oracle ASM (“SAP Databases on Oracle Automatic Storage Management 11g Release 2: Configuration Guidelines for Unix and Linux Platforms”), Oracle Linux and SAPCTL (“Providing High Availability for SAP Resources with Oracle Clusterware 11g Release 2”). All these white papers are stored on the SAP Developer Network (SDN) under <https://www.sdn.sap.com/irj/sdn/ora>. The SAP Notes are available through the SAP Support Portal for authorized users.

ASM Disk Group Recommendations for SAP Databases

Although there are no specific requirements for ASM Disk Groups storing the SAP databases on the Oracle Exadata Database Machine, it is a best practice to use a redundancy level of high for a production SAP database to achieve the highest level of protection against any type of storage failure. Other SAP databases used for development, test and QA may use a normal ASM redundancy level.

As there is no need for storage based replication with the Oracle Exadata Database Machine the following ASM Disk Groups for each SAP database should be used. In line with the standard Oracle Exadata setup you should have at least one ASM Disk Group “+DATA” and another ASM Disk Group “+RECO”. The DATA Group should contain all data files, control files, online redo log files, spfiles, OCR and voting disks. The RECO Group should contain temporary files, archive logs, flashback files and backups. The DATA Group should use a redundancy level of high and the RECO Group a redundancy level of normal.

When storing more than one SAP database (for instance an SAP ERP database and an SAP BW database or an SAP ERP database with an SAP CRM database or multiple SAP ERP databases) on the Oracle Exadata Database Machine all files of each of these SAP databases should follow the above recommendation and therefore files should be stored in the DATA and RECO Group.

For performance and throughput reasons it is recommended to only have two control files and non-multiplexed online redo log files for each SAP database all stored in the DATA Group. As standard SAP installations use three control files in the database it is recommended to remove one control file from the spfile or init.ora. Standard SAP installations also use two members for each online redo log file. On the Oracle Exadata Database Machine it is therefore necessary to remove one member of each online redo log file for each redo thread. The source database will have multiple redo threads if it was a RAC database. Three control files and multiplexed online redo log files are not needed on the Oracle Exadata Database Machine as the control files and the online redo log files are stored in the DATA Group which already provides three way mirroring for each file at the Oracle ASM level due to the redundancy level of high.

Character Set Requirements for SAP Databases

New installations of SAP systems from NetWeaver 7.0 on are Unicode installations only. For an SAP Unicode installation it is required that both the character and the national character set in the database is set to UTF8. When deploying a new Oracle Exadata Database Machine for Unicode installations of SAP it is therefore mandatory to specify UTF8 for both the character and the national character set in the Exadata configuration worksheet. A working example for the Exadata configuration worksheet can be found in Appendix 1 of this paper.

So please make sure that the default database on the Oracle Exadata Database Machine is created with the SAP required UTF8 character and national character set or a new SAP database must be created on the Oracle Exadata Database Machine for instance through DBCA with a character and national character set of UTF8.

Non-Unicode SAP Installations

Existing Non-Unicode SAP installations can be used with the Oracle Exadata Database Machine. It is important for these Non-Unicode installations that the character and national character set of the migrated databases from existing systems to the Oracle Exadata Database Machine is kept the same as in the original system. The character set will either be WE8DEC or UTF8. The national character set should always be UTF8.

It is mandatory that the SAP application of such a Non-Unicode SAP installation runs on an operating system which supports the Non-Unicode runtime requirements of SAP. The Product Availability Matrix (PAM) of SAP (<http://www.service.sap.com/PAM>) should be checked for valid operating system support for Non-Unicode SAP installations.

In the case of Non-Unicode installations it is highly recommended to not change the hardware or operating system for the SAP layer. Only the existing database server and storage layer should be changed to the Oracle Exadata Database Machine.

Note: In the case of Non-Unicode SAP installations you cannot run the SAP Central Services on the database nodes of the Oracle Exadata Database Machine.

SAP Database Administration with BR*Tools

The installation, configuration and operation of BR*Tools on the database nodes of the Oracle Exadata Database Machine are documented in SAP Notes 1598594 and 1598868.

Lifecycle Management for SAP Databases

An Oracle Exadata Database Machine requires lifecycle management at several levels of its hardware and software stack:

- Exadata Storage Server
- Database server
 - Oracle Database Server
 - Operating system and firmware
- InfiniBand switch
- Additional components

This section focuses on the Oracle Database Server and describes how to install Oracle Database Server software patches into the Grid Infrastructure Oracle Home and the RAC Oracle Home of an SAP database. For more information on the other components mentioned above, see MOS note 1262380.1 (“Exadata Patching Overview and Patch Testing Guidelines”).

The Oracle Database Server of an SAP database requires two bundle patches for a complete update:

- The regular Oracle Exadata Database Machine Bundle Patch (which contains patches for Database, ASM and Clusterware), also referred to as "Oracle Exadata Bundle Patch" in the following, and
- The SAP Bundle Patch for Oracle Exadata, also referred to as "SAP Exadata Bundle Patch" in the following.

Oracle tests and certifies both bundle patches for SAP databases on a regular basis and makes them available for SAP customers on the SAP Service Marketplace. You can find up-to-date release information on both bundle patches and their download locations in SAP Note 1591389 ("Exadata 11.2.0: Patches for 11.2.0.2").

Since each SAP Exadata Bundle Patch requires a specific version of the Oracle Exadata Bundle Patch, you cannot use versions of the Oracle Exadata Bundle Patch for SAP databases which have not been specifically certified for that purpose.

Installation of the OPatch and MOPatch Utilities

The installation of both bundle patches requires an up-to-date version of the OPatch utility. The installation of the SAP Exadata Bundle Patch additionally requires an up-to-date version of the MOPatch utility.

Appropriate versions of both utilities are available in the SAP Exadata Bundle Patch. See section "OPatch and MOPatch Utility Information" in the Readme document of the SAP Exadata Bundle Patch for instructions on how to extract and install these utilities.

Installation of the Oracle Exadata Bundle Patch

There are a number of options to install the Oracle Exadata Bundle Patch, which are described in more detail in the Readme document of the patch.

All these options are available for SAP databases. It is recommended, however, that you use the rolling installation approach with command "opatch auto", which is described in section "Patch Installation and Deinstallation" of the Readme document.

After you have installed the Oracle Exadata Bundle Patch, for example, with "opatch auto", you should immediately install the SAP Exadata Bundle Patch. In particular, you should not execute the SQL statements mentioned in section "Patch Postinstallation" of the Readme document, since the catsxd.sql script provided by the SAP Exadata Bundle Patch runs also the SQL statements required by the Oracle Exadata Bundle Patch.

Installation of the SAP Bundle Patch for Oracle Exadata

You must install the SAP Exadata Bundle Patch after you have installed the Oracle Exadata Bundle Patch.

Note: The SAP Exadata Bundle Patch is not installed by the Oracle OneCommand Exadata installation utility.

The SAP Exadata Bundle Patch contains patches which must be installed in the Grid Infrastructure Oracle Home (GI Home) and patches which must be installed in the RAC Oracle Home (RAC Home) of an SAP database. Use the MOPatch utility as described in section "Patch Installation and Deinstallation" of the Readme document to install the patches into all Oracle Homes of the Oracle Exadata Database Machine.

Finally, follow the post-installation steps in section "Executing Post-Installation Instructions" of the Readme document to run all required SQL statements, update the database dictionary, and maintain the database initialization parameters.

Migration of SAP Databases

Although there are several possibilities to migrate an existing SAP database to the Oracle Exadata Database Machine it is recommended to choose one of the following approaches as these have been successfully tested.

Migration Approach 1: Oracle-to-Oracle (O2O) ACS Service and Customer Self-Service

This way of database migration exists for many years and is being used to migrate an SAP database between different systems. The method is also described in the SAP Note 1508271. O2O allows for homogenous and heterogenous system copies.

A homogenous system copy is a migration where the source and target operating system is the same. A heterogenous system copy is a migration where the source and target system have different operating systems. With an heterogenous system copy you can for instance migrate an existing SAP AIX database to an Exadata Linux database.

O2O supports all operating systems SAP products are certified on.

The advantage of this method is, that you can combine the operating system change with multiple options to get most out of the migration:

- As part of the database migration, the whole database is reorganized. This can free up a significant amount of space within tables and indexes.
- The tablespace layout can be changed to the new SAP standard or to a customer own customized one. It is also possible to move single tables and indexes to separate tablespaces or to merge them into existing or new ones. This allows you to unify the SAP landscape by using a default tablespace name like “PSAPSR3” in all SAP systems
- The SAP schema name can be changed for instance to “SAPSR3” to unify the SAP landscape.
- The number of data files and mount points can be significantly reduced, by either optimizing the tablespace layout or the size of the data files and file systems
- Tablespaces are created with LMTS and ASSM
- Data files will be converted from filesystems to Oracle ASM
- LOB or LONG data types can be converted to Secure files (11.2 only)
- You can compress the indexes on the target database by using Oracle index compression. The compression calculation is executed as described in SAP Note 1109743. The correct index compression for each index is determined automatically.
- You can compress tables on the target system. The compression will compress all SAP tables as recommended in SAP Note 1431296.
- With O2O it is possible to combine a platform migration with a release upgrade. The migration method supports every combination of Unix, Windows or Linux

on source and target system. So you can migrate an existing Oracle 10.2 database on HP-UX to an Exadata Solaris database.

- It is possible to upgrade directly to a higher database release. Currently with the O2O method direct database migrations are possible between different Oracle versions. So it is possible to upgrade directly from Oracle 9i to Oracle 11g by using O2O. You also do not need the most current patchset of the lower Oracle release to run the migration. A complete overview about the upgrade paths between different Oracle versions is given at the end of this chapter.

The downtime needed to migrate a database with the O2O method is depending on the database size, the included database objects (SAP cluster tables, partitioned tables) and the available hardware resources (CPU, Memory, Storage, Network). Up to 1TB/hour is possible.

Migration Approach 2: Oracle-to-Oracle (Triple-O) ACS Service only

If the downtime requirements cannot be fulfilled with the O2O offline method, you can use as an alternative the Triple-O method. The Triple-O method allows to migrate an Oracle database with a downtime for the database copy of not more than 15 minutes. The downtime is independent from the database size. The downtime is needed to stop the SAP system on the source, check the final synchronization of the migrated system and to restart SAP on the target system.

The Triple-O method is a further development of the O2O method. This method is certified by SAP to perform an online migration and documented in SAP Note 1508271. So all systems migrated with this method are fully supported by SAP after the migration.

The Triple-O method is using the O2O technology to perform the initial data load (copying the current contents of the source database to the target database) while the SAP application is up and running. Oracle GoldenGate is used to synchronize all changes applied on the source system, while the migration is running. The method fully supports DDL and DML operations. There are no limitations for the usage of the SAP system during the migration.

With the Triple-O method you can make use of all features listed for the O2O method. Also heterogenous database migrations and direct database upgrades are supported.

When using Triple-O:

- There is no downtime needed to start or stop the online migration
- No need for special database patches
- Migration can be started and stopped at any time without risk for the productive system
- Downtime needed for the database copy part of the migration is typically only 15 minutes. The downtime is independent from the database size.
- Triple-O fully supports DDL and DML commands
- Triple-O is SAP release independent

Prerequisites for Triple-O:

- At least Oracle version 9i on source
- 30% free CPU time in average on database server
- Additional filesystem space of 50% of the redo log amount in 24h on the source
- Additional filesystem space of 50% of the redo log amount in 48h (or the duration of the initial data load) on the target

Limitations of Triple-O:

- Oracle GoldenGate currently does not support compressed tables on the source system. These tables must be either decompressed before the migration, or must be migrated at the end in a downtime windows

Impact on productive system:

- Increased redo log volume, because supplemental logging must be activated on the source

Summary:

These well established methods offer a fast and reliable approach, not only to migrate a SAP database from any SAP supported OS platform from Oracle Database 9i Release 2 on to the Oracle Exadata Database Machine, but also to implement a couple of additional features and benefits into the migration process. This adds more value to the database migration and will increase the ROI for instance by reducing the storage costs with table/index compression.

Migration Approach 3: Oracle RMAN and Transportable Tablespaces

Through the DUPLICATE DATABASE command of Oracle RMAN it is possible to copy an existing Linux or Solaris database running on an x86_64 platform to the Oracle Exadata Database Machine without any downtime.

By using Transportable Tablespaces it is possible to copy an existing database from any UNIX or Windows platform to the Oracle Exadata Database Machine.

Both ways are explained in detail in the document “Moving your SAP Database to Oracle Automatic Storage Management 11g Release 2”. The document can be downloaded from the SAP Developer Network (SDN) <https://www.sdn.sap.com/irj/sdn/ora>.

Shared Filesystems in SAP Environments

In an SAP environment it is common that all SAP Application Servers have access to a shared filesystem (/sapmnt, /usr/sap/trans, ...) which store the SAP kernels, profiles, trace files and provide the global SAP transport directory. In typical SAP installations such a shared filesystem is implemented using a NAS appliance, a cluster filesystem or through an NFS exported filesystem from the database server. For high availability reasons a cluster filesystem is being used or the source of the NFS location is protected by special configurations such as HA-NFS to not be a single point of failure in an SAP environment.

If you already have an existing shared filesystem solution in your SAP environment not using an NFS exported filesystem from the database server it is recommended to continue to use this solution when moving to the Oracle Exadata Database Machine.

If the shared filesystem is NFS exported from the database server in your existing environment it is necessary to implement the shared filesystem on a different system other than the Oracle Exadata Database Machine as the Oracle Exadata Database Machine does not offer any HA-NFS or cluster filesystem capability. The preferred solution is to use a separate Sun ZFS Storage Appliance. Other NAS appliances, HA-NFS or cluster filesystem solutions can be used as well. The Sun ZFS Storage Appliance can also be used for very fast backups from the Oracle Exadata Database Machine by directly connecting the ZFS Storage Appliance to the Oracle Exadata InfiniBand fabric.

Protection of SAP Central Services

In an SAP environment certain components such as the Enqueue (SCS, ASCS) or Message Server or the Web Dispatcher need to be monitored and controlled by non-SAP high availability software to make the whole SAP system highly available. Almost every high availability software like IBM PowerHA, HP Serviceguard, Veritas Cluster Server, Oracle Solaris Cluster or Oracle Clusterware provide additional services to protect these critical SAP central services. In typical SAP environments this high availability software is running either on the clustered database server or outside the database server on a separate cluster of servers.

If a separate cluster other than the database cluster is already used for the SAP Central Services then it is recommended to continue to use this separate cluster for the SAP Central Services when deploying the Oracle Exadata Database Machine.

Should the SAP Central Services run on the clustered database server in your existing environment then you should consider to install two additional x86_64 systems running Oracle Solaris and Oracle Solaris Cluster on these two x86_64 systems to protect the SAP Central Services.

Note: Oracle Solaris Cluster on x86_64 hardware cannot be used for the SAP Central Services for Non-Unicode SAP installations. As an alternative for Non-Unicode SAP installations you should built an Oracle Solaris Cluster using SPARC hardware.

Another alternative for Unicode-only SAP installations is to use the Oracle Clusterware running on the database nodes of the Oracle Exadata Database Machine along with the Oracle Clusterware utility SAPCTL to protect the SAP Central Services . The next chapter describes in detail all the necessary steps to install the SAP Central Services on the database nodes of the Oracle Exadata Database Machine and how to protect them through Oracle Clusterware and SAPCTL.

Note: Please be aware that any Exadata Storage Software change (patch, patch bundle or upgrade) may affect the configuration and the operation of the SAP Central Services on the database nodes of the Oracle Exadata Database Machine. So please check the configuration and correct operation of the SAP Central Services on the database nodes after an Exadata Storage Software change was applied.

Installation Procedure for SAP Central Services

The SAP Central Services for ABAP and JAVA are typically provided by SAP instance type ASCS for ABAP and SCS for JAVA. The installation of these instances on Oracle Exadata Database Machine must be performed using SAP installation tool SAPINST. For a High-Availability Setup including Enqueue Replication Service (SAP instance type ERS), the initial installation must be performed on every database node on the Oracle Exadata Database Machine. Installation of instance type ASCS and SCS is fully supported by SAPINST, whereas the configuration and setup of the Enqueue Replication Service must be done manually. The following sections specify the steps required for the initial installation of SAP Central Services on Oracle Exadata Database Machine.

Installation Preparation

Please check the minimum requirements for your specific SAP software version or system type regarding OS parameters, user limits, etc. Consult SAP documentation for the recommended values and check if requirements are met on all the database nodes. If necessary, adjust at least to the minimum required value.

Download JCE (Java Cryptography Extension) policy file `jce_policy-1_4_2.zip` to all database nodes. Note that exactly this version is required during installation.

Assign a virtual hostname for the ASCS or the SCS instance. This virtual hostname will be used for the network name resolution and represents the IP address of the Oracle Clusterware VIP which will be used by SAPCTL to provide failover protection for the SAP Central Services.

We will refer to these virtual hostnames as **xsapdb_abap** and **xsapdb_java** throughout the following sections.

Add the virtual hostnames to the `/etc/hosts` file on all database nodes.

```
[root@xsapdb01] # vi /etc/hosts
. . .
10.165.110.180 xsapdb_abap xsapdb_abap.de.oracle.com
10.165.110.181 xsapdb_java xsapdb_java.de.oracle.com
```

As user root, change to the directory containing the SAP Installation Master CD/DVD.

Add the variable `JCE_POLICY_ZIP` and `SAPINST_USE_HOSTNAME` to the environment.

For the installation of the SAP Central Services Instance for ABAP set

```
[root@xsapdb01] # export JCE_POLICY_ZIP=/<full-path-to>/jce_policy-1_4_2.zip
[root@xsapdb01] # export SAPINST_USE_HOSTNAME=xsapdb_abap
```

For the installation of the SAP Central Services Instance for JAVA set

```
[root@xsapdb01] # export JCE_POLICY_ZIP=/<full-path-to>/jce_policy-1_4_2.zip
[root@xsapdb01] # export SAPINST_USE_HOSTNAME=xsapdb_java
```

On the Oracle Exadata Database Machine the X11 libraries for a graphical user interface are not installed. You must use another host providing a graphical user interface for software installation with SAPINST. Export the DISPLAY variable to the host providing the GUI.

```
[root@xsapdb01] # export DISPLAY=<host_with_gui>:1
```

Now start the SAP software installation tool.

```
[root@xsapdb01] # ./sapinst
```

Follow the installation steps.

Use identical user-id and group-id for user <sid>adm on all database nodes.

The instance number for ASCS instances must be identical on all database nodes.

The instance number for SCS instances must be identical on all database nodes.

Preferable, also the password for SAP System Administrator is identical on all nodes.

Complete the installation on all database nodes.

Setup of SAP Enqueue Replication Service

The configuration task to setup an SAP Enqueue replication server instance ERS for either SAP ASCS or SAP SCS is not supported by SAPINST and must be done manually.

Check the SAP documentation for further guidance. In a nutshell we will explain the necessary steps here to give a overview on the tasks to do.

Remember that these steps must be executed on all database nodes of the Oracle Exadata Database Machine.

As user <sid>adm, create directory `/usr/sap/<SID>/ERS<NR>` with sub-directories **data**, **exe**, **log**, **sec** and **work**.

In the profile directory `/usr/sap/<SID>/SYS/profile` create the Instance and START profile for the ERS instance. Check the SAP documentation of the SAP software release installed to determine what type of profile is required to start up the instance. This may differ depending on the SAP product or version used.

The following example shows a working configuration for the START and Instance profile which can be used for reference purpose.

START profile for ABAP replication service

```
SAPSYSTEMNAME = KCM
SAPSYSTEM = 02
INSTANCE_NAME = ERS02
DIR_CT_RUN = $(DIR_EXE_ROOT)/run
DIR_EXECUTABLE = $(DIR_INSTANCE)/exe
SAPLOCALHOST = xsapdb_abap
DIR_PROFILE = $(DIR_INSTALL)/profile
_PF = $(DIR_PROFILE)/KCM_ERS02_xsapdb_abap
SETENV_00 = LD_LIBRARY_PATH=$(DIR_LIBRARY):%(LD_LIBRARY_PATH)
SETENV_01 = SHLIB_PATH=$(DIR_LIBRARY):%(SHLIB_PATH)
SETENV_02 = LIBPATH=$(DIR_LIBRARY):%(LIBPATH)
#-----
```

```
# Copy SAP Executables
#-----
_CPARG0 = list:$(DIR_CT_RUN)/scs.lst
Execute_00 = immediate $(DIR_CT_RUN)/sapcpe$(FT_EXE) pf=$(PF) $_CPARG0
#-----
# Start SAP locking service replication
#-----
_ERS = enr.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
Execute_01 = local rm -f $_ERS
Execute_02 = local ln -s -f $(DIR_EXECUTABLE)/enrepsvr$(FT_EXE) $_ERS
Start_Program_01 = local $_ERS pf=$(PF)
```

Instance profile for ABAP replication instance

```
SAPSYSTEMNAME = KCM
SAPSYSTEM = 02
INSTANCE_NAME = ERS02
DIR_CT_RUN = $(DIR_EXE_ROOT)/run
DIR_EXECUTABLE = $(DIR_INSTANCE)/exe
SAPLOCALHOST = xsapdb_abap
#-----
# SAP Message Server parameters are set in the DEFAULT.PFL
#-----
ms/standalone = 1
ms/server_port_0 = PROT=HTTP,PORT=81$$
#-----
# SAP Enqueue Server
#-----
enqueue/table_size = 4096
rdisp/myname=ERS02_xsapdb_abap
rdisp/enqname = $(rdisp/myname)
enqueue/snapshot_pck_ids = 100
#-----
# SAPCTL HA Impl. with SAP Enq. Replication
#-----
rdisp/mshost = xsapdb_abap
enqueue/process_location = local
enqueue/server/internal_replication = true
enqueue/server/replication = true
enqueue/enrep/keepalive_count = 1
enqueue/server/threadcount = 1
```

Edit the file `/usr/sap/sapservices`. Add an entry for every replication server instance.

Example from a working configuration for reference purpose:

```
LD_LIBRARY_PATH=/usr/sap/KCM/ERS02/exe:$LD_LIBRARY_PATH;
export LD_LIBRARY_PATH;
/usr/sap/KCM/ERS02/exe/sapstartsrv
pf=/usr/sap/KCM/SYS/profile/START_ERS02_xsapdb_abap -D -u
kcmadm
```

Examples of the configuration files for the JAVA Central Services can be found in the SAPCTL white paper (“Providing High Availability for SAP Resources with Oracle Clusterware 11g Release 2”).

Please check also all available SAP documentation and support notes on additional configuration tasks which are required for an SAP HA installation with failover capabilities.

Configuration of SAPCTL

Please refer to the latest version of the SAPCTL white paper stored on the SAP Developer Network (SDN) under <https://www.sdn.sap.com/irj/sdn/ora>. Follow the instructions on installation and configuration described in the collateral documentation.

Few additional tasks on the Oracle Exadata Database Machine for running SAPCTL must be executed.

After the installation and setup of SAPCTL on the first database node, copy the files **sapctl**, **saptctl.pl**, **crssapactions.pl**, **sapwrap** from directory `/usr/sap/<SID>/SYS/exe/run` to the same directory on all remaining nodes.

We recommend to copy the perl subdirectory from the GRID software installation recursively to directory `/usr/sap/sapctl` and change the ownership to `<sid>adm:sapsys` to avoid access permission problems.

```
[root@xsapdb01] # dcli -g <group> cp -r /u01/app/11.2.0/grid/perl
/usr/sap/<SID>/sapctl
[root@xsapdb01] # dcli -g <group> chown -R <sid>adm:sapsys
/usr/sap/<SID>/sapctl
```

Set variable `PERL_HOME` in file `sapctl` to point to this directory.

Note that there is no need to duplicate the `START` and instance profiles as described in the SAPCTL white paper if the installation of SAP Central Services for ABAP or JAVA as well as the setup of SAP Enqueue ERS replication services was performed using virtual hostnames as described in the section “Installation Procedure for SAP Central Services”.

Appendix 1:

SAP Oracle_Home Naming Convention

The correct installation and operation of any SAP utility such as SAPIInst or BR*Tools on the database nodes of the Oracle Exadata Database Machine requires some preparation for the correct setting of the Oracle_Home environment variable in the SAP environment. The SAP environment requires the Oracle_Home environment variable to be set to /oracle/<Database name>/<release>.

Therefore the following directories and symbolic link (according to the values listed in the configuration sheet below) must be manually created on each database node by the OS user who owns the Oracle software (In the configuration sheet below the name oracle is being used for the OS user):

```
root> mkdir /oracle
root> chown oracle:oinstall /oracle
root> su - oracle
oracle> mkdir -p /oracle/X11
oracle> ln -s /u01/app/oracle/product/11.2.0/dbhome_1 /oracle/X11/112
```

Default Oracle Environment Settings

The following configuration sheet for the Oracle Exadata Database Machine is an example for a successful SAP installation with Oracle Exadata. The sheet below shows in red the values required by SAP.

Note: Please be aware that for SAP installations only Standard OS Authentication is supported therefore Table 2 is not included in this paper.

Table 1 lists the default settings when selecting Standard OS Authentication for the OS owner used during installation to create the Oracle software environment. These default settings are in addition to the information in the configuration worksheets.

Table 1 Oracle Environment Default Settings when Using Standard OS Authentication

Oracle Database Item	Default Setting
Oracle Inventory group name	oinstall
Oracle Inventory group identifier	1001
DBA group name	dba
DBA group identifier	1002
Oracle software owner user name	oracle
Oracle software owner user identifier	1000

Oracle software owner default password	welcome
Oracle base directory (ORACLE_BASE)	/u01/app/oracle
Oracle inventory directory	/u01/app/oraInventory
Grid infrastructure home directory	/u01/app/11.2.0/grid
Database name	X11 (must be three characters)
Database character set	UTF8 (WE8DEC for Non-Unicode SAP)
Database national character set	UTF8
Database block size	8192
ASM disk groups	DATA for the default data file location RECO for the fast recovery area Note: Default DATA and RECO disk group size depends on the type of system, the type of disk drives, and the type of backup.
Exadata Smart Flash Cache	All flash disks are configured as flash cache
Starting IP address for InfiniBand private network	192.168.10.1
Subnet mask for InfiniBand network	255.255.252.0

In the general configuration worksheet (Table 3) you should always specify a workload type of **OLTP** for any SAP database. Even for an SAP BW database you should use the OLTP workload type OLTP as SAP BW uses serial DML (Insert, Update, Delete) and very few full table scans against the database. Most of the data access in an SAP BW database occur through indexes (Bitmap and B-trees).



Using SAP NetWeaver with the
Oracle Exadata Database
Machine
July 2011
Authors: Kurt Brög, Stephan
Bühne, Jan Klokkers, Christoph
Kurucz, Martin Sautter, Jens
Schmidt

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065



Oracle is committed to developing practices and products that help p

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

0711