WebSphere software

WebSphere Process Server Version 6.2.0 Configuring The Process Choreographer And Other BPM Components In A Clustered Environment

A step by step guide based on Linux and Oracle 11g

including Business Space and IBM HTTP Server

August 2009

This edition applies to Version 6.2 of WebSphere Process Server including the Business Space and the IBM HTTP Server.

Disclaimer

This document is subject to change without notification and will not comprehensively cover the issues encountered in any customer situation.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS.

For updates or newer releases please contact the service team.

The Team of Authors

This document is produced by the Business Process Choreographer team in Böblingen.

Michael Mann
IT Senior Specialist
IBM Software Group, Application and Integration Middleware Software,
WSS Business Process Solutions Test

Eduard Welte
IT Architect
IBM Global Buisness Services, Enterprise Application Integration

Marco Lezajic
IT Specialist
IBM Software Group, Application and Integration Middleware Software,
WSS Business Process Solutions Test

Ekkehard Voesch
Business Process Management Test Architect
IBM Software Group, Application and Integration Middleware Software,
WSS Business Process Solutions Test

Thanks for the contributions of the Business Process Choreographer test team and the WSS BPM competence center.

Preface

Scope

This document contains an overview of the tasks required for setting up a production environment of a Business Process Management solution with a main focus on WebSphere Process Server and in particular the Business Process Choreographer. Additional Business Process Management components or products will be taken into consideration, especially if interactions with the Process Choreographer are of interest.

This document introduces the topology of the chosen production environment and describes the different tasks with a series of click-by-click instructions including the installation, configuration and a verification for each main step.

The information provided in this document is related to a production environment with several hosts and clusters. Please notice that a single server setup will be significant less complex in terms of the WebSphere topology. Also the overall setup of a remote database system is much more complex in production environment than in a single server setup where a local database system might be used.

WebSphere Business Space or WebSphere Business Monitor might be added as optional components to this production environment. They are part of the same cell and represented in different, additional clusters. IBM HTTP Server is required for request distribution into cluster members. The related chapters can be skipped if these components or products are not needed.

The installation and configuration of additional Business Process Management components or products is not part of this documentation, but might be added in a later version of this document.

The general concepts for building cells and clusters apply to the different releases, but several details might have been added or changed between the different releases. To make reading and understanding easier this document is related to the following explicit releases:

- WebSphere Process Server V 6.2 (including Business Space)
- WebSphere Business Monitor V 6.2
- Oracle 11g Release 1 (11.1.0.6.0)
- IBM Tivoli Directory Server V 6.1

The used operating system for this documented setup is Red Hat Enterprise Linux. In general other operating systems especially UNIX based systems will operate in a similar way. From a configuration point of view there will be only minor differences, for example the default path names might be different.

Additional document sources being of interest

IBM WebSphere Business Process Management Version 6.2 information center

```
http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r2mx/index.jsp
```

Especially the sections for the following products are of interest

- IBM WebSphere Process Server 6.2

IBM Redbooks

```
http://www.redbooks.ibm.com/
```

There are numerous publications available for WebSphere Business Process Management from which the following are of interest

- WebSphere Application Server V6 System Management and Configuration Handbook (SG24-6451)
- WebSphere Application Server Network Deployment V6: High Availability Solutions (SG24-6688)
- WebSphere Business Process Management V6.1.2 Production Topologies (SG24-7665)
- WebSphere Application Server V6 Problem Determination for Distributed Platforms (SG24-6798)

Summary of changes

Changes in the August 2009 version.		
Type	**Chapter**	**Change**
Content	7	Changed the user IBMBUSSP to WPS_BSPACE.
	7.2.1	User IBMBUSSP in no longer mandatory for the BusinessSpace configuration. The way of configuring the BusinessSpace has changed.
	14.5	Changed image of installed applications
	15	Changed the way how to configure the BusinessSpace.
	16 - 26	Added description how to install WebSphere Business Monitor
	27	Improved the Claims Handling application installation with additional imformation.
	28	Added description how to view WBM details in Business Space
	29	Reworked the Appendix. Added new Chapters in the Oracle Appendix
Wording	7	Changed all terms "rights" to "privileges"

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A current list of IBM trademarks is available on the Web at

http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
AlphaBlox®
alphaWorks®
DataPower®
DB2 Universal Database™
DB2®
developerWorks®
HACMP™
IBM®
IMS™
Lotus®
Redbooks®
Redbooks (logo) ®
Tivoli®
WebSphere®
Workplace™
Workplace Messaging®
z/OS®

Trademarks

The following terms are trademarks of other companies:

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

SAP NetWeaver, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

EJB, Enterprise JavaBeans, J2EE, Java, JavaBeans, JavaServer, JDBC, JMX, JSP, JVM, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.

How to read this document

The intention of this document is to provide information which might be used to facilitate the setup of a production environment. To get a better understanding further documentation might be referenced.

This document contains numerous illustrations being formatted as follows.

- ***Graphical interactions***
  Typical interactions with the system are performed on screen (or panels) and a description how to interact is much easier if pictures are used as outlined here

Title of the screen

1. Special related instructions – maybe several items are available.
   They are marked with correspondend green bulleted numbers.
   In addition sometimes related areas are marked with green rounded rectangle.
2. The final instruction, e.g. press the Unpack button, is typically not highlighted.

- ***Interaction via the console***
  Mostly two types are of interest, one for the input requested on a console window sometimes mentioned as command line input and a second for the output provided on the console.

```
Console input (sometimes mentioned as command line input)
```

```
Console output
This might be several rows
In case of large lists important aspects are highlighted in this way
```

- ***Listings***
  Whenever a list of parameters is necessary to be discussed they will look like this. Also file content will be formatted in this way. In case of large lists highlighting might be added to put the focus on the major aspects.

```
Parameter1 = value1
Parameter2 = value2
Parameter3 = string1
etc
```

- ***Hyperlinks***
  For making it easier to find references hyperlinks are used and formatted like the following link which leads to the IBM Redbooks homepage

```
http://www.redbooks.ibm.com/
```

- ***Notices***
  To emphasize information two types of formatted notices are used

Standard notice – typically used

Importand notice – used in special cases

Table of Contents

Part I Introduction

Chapter 1 Introduction

This part of the document provides planning information. It is highly recommended to plan the setup before starting with the installation and configuration. Careful planning can avoid both the need to reset the involved systems or a complete restart of the overall setup process.

The first part of this chapter (Planning the cell and cluster setup primer) provides a brief overview of the prerequisite actions that need to be completed before starting with the actual setup.

In the second part of this chapter (Planning the cell and cluster setup) a brief description of the major components for building a cell is provided followed by an overview of how the components are assigned to the systems.

The term system is used in multiple different facets depending on the author, industry, functional area etc. In the IT world itself various interpretations also exist. A common understanding does not really exist, but often the term system is a synonym for a machine or a set of machines, e.g. a server or personal computer. Furthermore the term system is used to specify a combination of machine and inclusive software used to build a solution. From a usage perspective such solutions (systems including software) are often referenced as service. An application server, especially the WebSphere Process Server, might be understood as a system providing services.

Depending on the requirements related to performance and scalability the service can be implemented on a single machine or it might be necessary to use a set of machines where the term machine in this context is a server. On large machines, e.g. IBM Power Systems, the overall power might be distributed in partitions where each partition will be referenced as an independent system. Another option might be to configure the software multiple times on a single machine where it is usual to reference each configuration as a system.

When multiple configurations of the WebSphere Process Server are reside on the same machine they are sometimes referenced as instance.

In this document the term system is used to define a machine where the WebSphere Process Server is installed and configured.

1.1 Planing the cell and cluster setup primer

The following list of tasks/subtasks should be completed in order to enable a straightforward cell setup including clusters (A checklist can be found in Appendix):

1.1.1 Assign WebSphere components to the systems

This task comprises multiple subtasks, including:

- Topology selection (performance, scalability)
- WebSphere Process Server software level selection (e.g. version 6.2.0.1)
- Database system selection, database tuning, and database distribution (e.g. DB2 version 9.5 FP1, Oracle 11g Release 1)
- Number of systems, hardware and operating system
- Establish new systems (hardware) or reuse/adapt existing ones
- Provide hardware and operating system requirements
- Check conflicts when using existing systems (e.g. ports, applications, firewalls)
- Create a topology chart
- Plan hostnames (used during WebSphere Process Server configuration)

For further more detailed information related to concepts, architecture and background aspects please check the following documentation.

WebSphere Process Server V6.2 Information Center

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r2mx/index.jsp?topic=/com.ibm.websphere.wps.620.doc/welcome_wps.html

WebSphere Business Process Management V6.1.2 Production Topologies

http://www.redbooks.ibm.com/abstracts/sg247665.html?Open

1.1.2 Install the operating system and the prerequisites

Depending on the customers business preferences one of the supported operating systems may be used. In this document the systems used will be Linux systems. For a detailed list of hardware and software prerequisites refer to the WebSphere Process Server V6.2 Information Center.

> A list of the minimum software and hardware required for WebSphere Process Server Version 6.2 can be found here:
>
> http://www.ibm.com/software/integration/wps/sysreqs/

Another important requirement is to ensure all involved systems are using a synchronized time. Typically this will be achieved by a timeserver providing system time for all systems.

1.1.3 Select the type of user directory for the cell

General options are: Federated repository, IBM Tivoli Directory Server (LDAP) or one of the other supported custom registries. The federated repository enables the support of multiple repositories which can be file-based, LDAP, database, and custom registries.

In the various documents an additional option (local operating system) is mentioned, but in a cell or clustered environment (which is the scope of this document) it is not supported.
Select what fits your needs best. As most environments already maintain a user directory which can be attached easily.

> Note: If you intend to use the Business Space component it is mandatory to use the federated repository.

1.1.4 Get the software packages

Obtain the required software packages of the desired product version and fix packs in a format that fits your needs best (DVD, CD, zip-pkg …)

1.1.5 Summarize the values used during installation

The following values are needed during the installation and configuration steps. You need to plan them carefully before starting.

- Operating system user IDs and passwords
- Database user IDs and passwords
- Messaging user IDs and passwords
- WebSphere Process Server user IDs and passwords
- Directory locations for:
  - Install sources
  - Database system
  - Database drivers
  - WebSphere Process Server install directory
  - Profile directories for deployment manager and custom nodes
- Hostnames of the involved systems
- Naming of cluster components
  - Cell name
  - Node names
  - Profile names
  - Database names

"Naming considerations for profiles, nodes, hosts, and cells":, InfoCenter:

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r2mx/index.jsp?topic=/com.ibm.websphere.wps.620.doc/doc/cins_naming.html

1.2 Planning the Cell and Cluster Setup

With a cell setup planning is highly recommended since a cell setup is much more complex compared to a single server setup. A single server contains all required elements in one unit so that no communication between the different components across systems is necessary. However with the scaling limitations of a single server, and with more complex requirements a single server may not longer be sufficient. In this case a cell with multiple server instances should be used.
Just adding more single servers will not fit all aspects. Especially synchronous administration as well as controlled communication are not covered. This can be achieved by using the WebSphere Application Server Network Deployment solution which provides several additional capabilities. To understand this concept several basic built-in components will be outlined.

1.2.1 Major Components of a Cell

Various elements are building a cell also referred as network deployment topology.

1.2.1.1 Cell

A WebSphere cell is a logical unit containing all elements being required for operation. In other words it is a management domain for a distributed environment of elements that are centrally managed and have access to shared resources in the cell. The key elements are briefly described in the following paragraphs. In physical terms a WebSphere cell typically involves multiple systems.

1.2.1.2 Node

A WebSphere node is an element within a cell which typically contains one or more application servers. A managed node consists of a node agent managing the the node and all application servers hosted on that node. An unmanaged node does not have a node agent and will be used to administer the IBM HTTP server (also referred as WebServer). A WebSphere node resides on a single host, but on a host it is possible to configure multiple nodes having then multiple node agents.

1.2.1.3 Node Agent

A WebSphere node agent is an architectural component that enables the deployment manager for the cell to remotely manage the node, its application servers, and their applications. The node agent represents the node in the management cell and keeps the configuration up-to-date.

1.2.1.4 Deployment Manager

A WebSphere deployment manager is a special application server whose only task is the management and configuration of all elements in the cell in which it exists. The deployment manager requires no node agent. The deployment manager runs a single application, a Web-based configuration front-end known as the Integrated Solutions Console or Administrative Console.
When J2EE™ applications are deployed within a cell, they are deployed to a server, cluster or multiple instances of both being referred to as the *deployment target*.

1.2.1.5 Application Server

A WebSphere application server hosts zero or more applications. An application server instance can be configured as follows:

- Stand-alone application server
  A stand-alone application server does not belong to a cell and runs its own administrative console.
- Managed application server
  A managed application server is part of a cell and is managed by a deployment manager residing on a separate node. A managed application server is not part of a cluster.
- Application server as member of a cluster
  An application server that is a cluster member resides on a node belonging to a cell, and is managed by a deployment manager residing on a separate node. The application server is part of a cluster.

1.2.1.6 Cluster

A WebSphere cluster is a logical collection of application servers configured to perform the same task as a team. The members of a cluster can be distributed across one or more nodes in any configuration. Clusters are typically used for scalability, workload and high availability aspects.

1.2.1.7 Service Integration Bus

A service integration bus supports applications using message-based and service-oriented architectures. A bus is a group of interconnected servers and clusters that have been added as members of the bus. Applications connect to a bus at one of the messaging engines associated with its bus members.
A messaging engine is a server component that provides core messaging functionality of a service integration bus. A messaging engine manages bus resources and provides a connection point for applications.
Each messaging engine is associated with a server or a server cluster that has been added as a member of a bus. When you add an application server or a server cluster as a bus member, a messaging engine is automatically created for this new member. If you add the same server as a member of multiple buses, the server is associated with multiple messaging engines (one messaging engine for each bus). You can create additional messaging engines for use with server clusters that are bus members, for availability and scalability reasons. However, in its simplest form a single engine can realize a bus.

1.2.1.8 Database

In a complex environment typically various data exist which need to be stored persistently. Within WebSphere accessing a database is transparent. It will be accessed via a JDBC implementation related to the used database system. All required database variables like name, location, user, password, etc are collected in the data source and are administered via the deployment manager. Within WebSphere Process Server several components persist data. Depending on the overall database convention used this results in several independent databases or in several database schema managed all in one database.

1.2.1.9 Profile

Each profile defines a separate run-time environment, with separate command files, configuration files, log files, and so on. Profiles enable you to have more than one runtime environment on a system, without having to install multiple copies of the system files. In WebSphere there are three kinds of profiles:

- The stand-alone server profile
- The deployment manager profile
- The custom profile.

The profile management tool will be used to built one of these profile types. In a cell the profile for the deployment manager is the first one followed by the custom profiles. After a custom profile has been created it needs to become a known element in the cell. This is achieved by federating the node into the cell. Please note that typically profiles have a one to one relation ship to nodes.
When additional products like the WebSphere Process Server or the WebSphere Business Monitor are installed, the profiles need to be augmented to be able to serve the additional functions.

1.2.2 WebSphere Process Server

The WebSphere Process Server is an SCA-compliant runtime element that provides a fully converged, standards-based process engine that is underpinned by WebSphere Application Server. Along with WebSphere Enterprise Service Bus, it is a strategic product for integration and modernization of IT assets, including core systems using SOA. Following the principles of SCA, there is a single invocation model, a single data model, and a component-based framework.

The WebSphere Process Server can consist of a single server running on a single machine, several independent servers running on multiple machines administered in a cell (referred as managed servers), or it may consist of several servers combined to a cluster running on several machines. Within a cell multiple clusters and multiple managed servers are allowed and are referred as deployment targets.

1.2.3 WebSphere Business Monitor

The WebSphere Business Monitor is a CEI (Common Event Infrastructure) based runtime element that provides business specific monitoring capabilities. Common Base Events which can be emitted by various applications (i.e BPEL Processes running on WebSphere Process Server) are consumed by monitor models. Based on defined metrics Key Performance Indicators (KPIs) and Visual Diagrams can be generated, viewed and monitored within Business Space.

Like WebSphere Process Server WebSphere Business Monitor can consist of a single server running on a single machine, several independent servers running on multiple machines administered in a cell (referred as managed servers), or it may consist of several servers combined to a cluster running on several machines. Within a cell multiple clusters and multiple managed servers are allowed and are referred as

deployment targets.

1.2.4 Assign the cell and cluster components to the systems

In general numerous topologies are possible and a decision on the best solution depends on the appropriate customer needs. In this document a topology is chosen which has an average complexity and is suitable for showing the various configuration steps.

The selected topology for the environment is a mixed WebSphere Process Server (WPS) and WebSphere Business Monitor (WBM) cell. The cell contains several clusters.

WPS clusters:

- **Application Cluster ("BPEL Cluster")**
  Contains the business process applications (I.e BPEL Processes).

- **A separate Message Engine Cluster ("MECluster")**
  Contains the messaging engines and the SI buses.

- **Support Cluster**
  Contains the Common Event Infrastructure (CEI), the Business Process Explorer/Observer and and the Business Process Event Collector.

WBM clusters:

- **Monitor Support Cluster**
  Contains the monitor action services and the monitor data movement services application.

- **Monitor Application Cluster**
  Contains the moderator and the logic part of the monitor model.

- **Web Dashboard Cluster**
  Contains Business Space and Alphablox.

The cluster environment is established on 5 dedicated physical machines, the database on a separate sixth machine. As security is also required, an existing user directory (IBM Tivoli Directory Server – in this document typical mentioned as LDAP) is attached to the cluster:

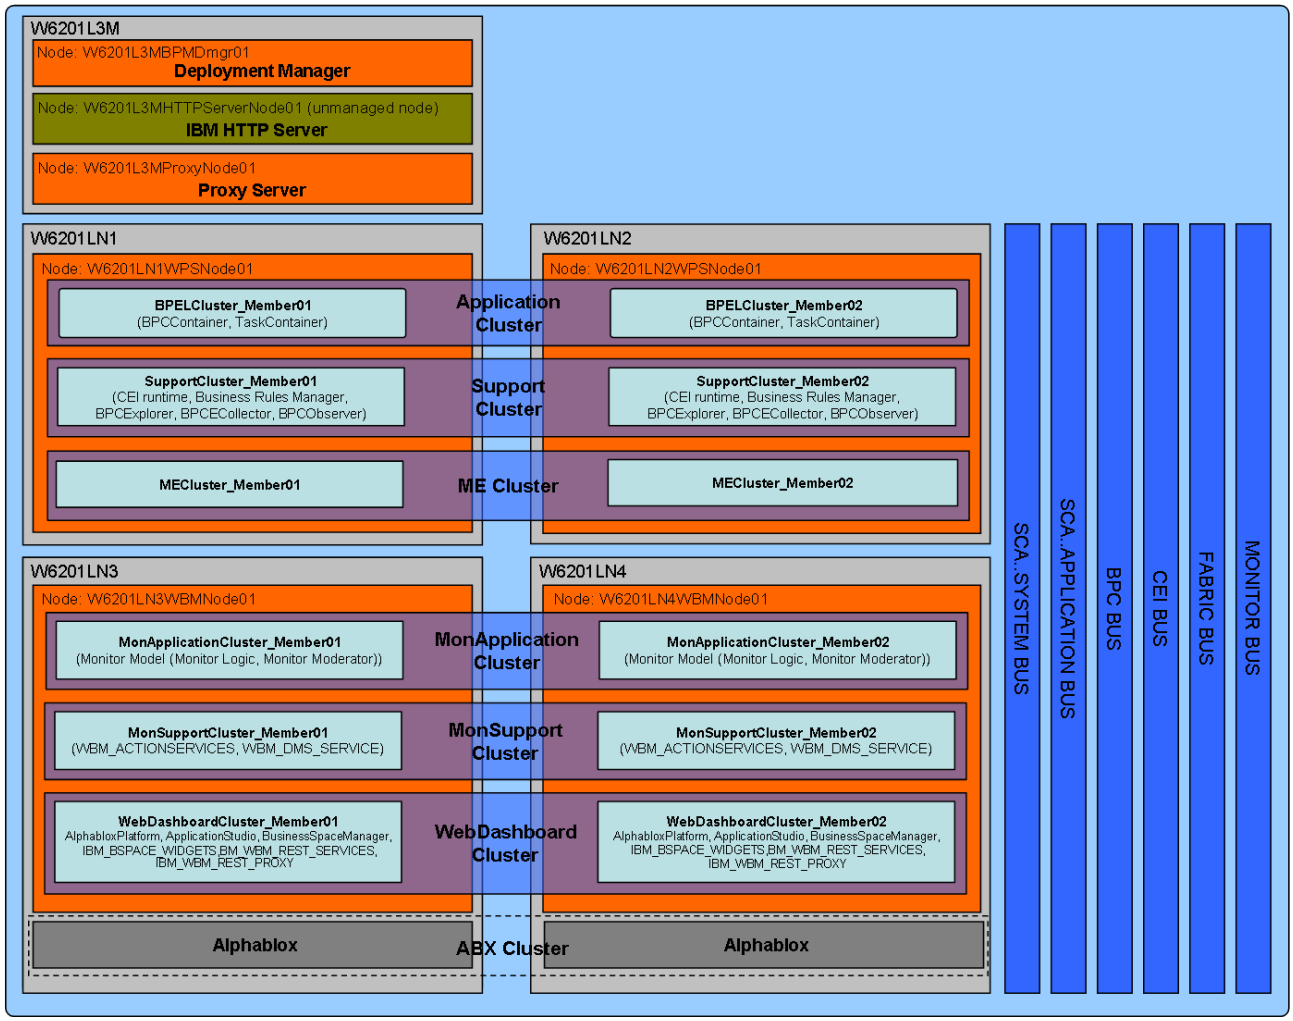

The setup is built on six dedicated machines, one hosting the database installation, the other hosting the deployment manager, the WebSphere Process Server nodes, the WebSphere Business Monitor nodes and the Alphablox insallation building up the cell. In this topology the messaging engines are separated from the applications. Therefore the messaging will be done in a separate cluster, the “MECluster”. Since the messaging functionality does not need any WebSphere Process Server specific functionality the “MECluster” might be based on simple WebSphere Application Server nodes, but this is not discussed in this document. On the separate cluster called “BPELCluster”, the BPE Container and the Task Container will be installed to host all process applications in the future. Last but not least the supporting functions provided by WebSphere Process Server (like the Common Event Infrastructure, the BPC Explorer, BPC Explorer Reporting Function and the Business Space) are installed in the third cluster called “SupportCluster”. Both, the “BPELCluster” and the “SupportCluster” will use the “MECluster” for all messaging purposes.

The Monitor Application Cluster hosts the moderator and the logic part of the monitor model. The moderator part receives Common Base Events and forwards them to the logic part of the monitor model. The logic part receives common base events from the moderator part and populates the monitor model (assigns values to metrics, calculates KPIs, creates visual diagrams, etc.)
The Monitor Support Cluster hosts monitor specific support applications like the monitor action services application and the monitor data services scheduler application. The applications and their purpose are described in detail when they are installed and configured. The Web Dashboard Cluster hosts the Business Space, Monitor REST services and the Alphablox application artifacts.

This document describes how to install and configure the outlined WebSphere Process Server and WebSphere Business Monitor clusters. There is only one active messaging engine per bus in the cluster. The other messaging engines are in stand-by mode, ready to become active if the currently active messaging engine goes down. The following picture illustrates the Messaging Cluster with the different messaging engines in active and standby state (1-of-N policy).

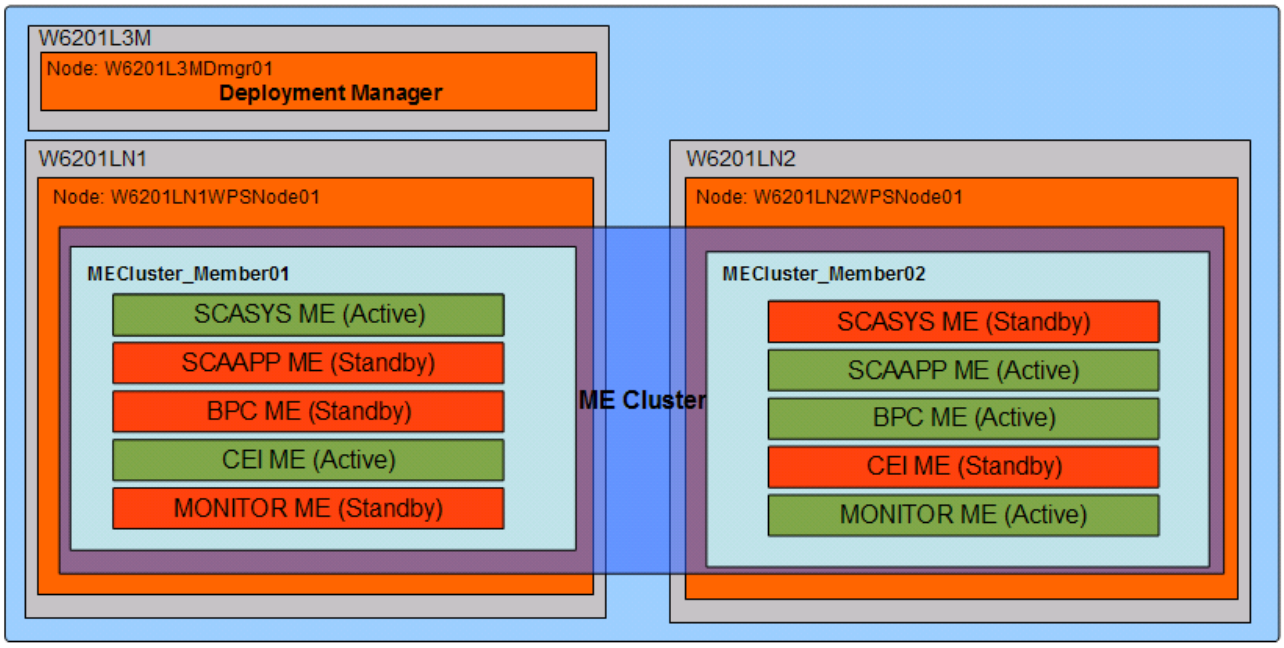

PART II Preparation

Chapter 2 Prerequisites and Operating system installation

2.1 Install operating system

In this document it is assumed that a pre-installed Red Hat Enterprise Linux system can be used. Therefore the installation steps of the operating system are skipped. Thus continue directly to check/provide the prerequisites:

Verify operating system level. Please note that this is only a sample and the output on the actual used system might be different.

```
cat /etc/redhat-release
```

```
Red Hat Enterprise Linux Server release 5.3 (Tikanga)
```

```
cat /proc/version
```

```
Linux version 2.6.18-128.el5xen (mockbuild@hs20-bc1-7.build.redhat.com)
(gcc version 4.1.2 20080704 (Red Hat 4.1.2-44)) #1 SMP Wed Dec 17 12:01:40
EST 2008
```

Check on all machines, that the WPS 6.2.0 required prerequisites are met. Do this using the information provided here:

"Preparing Linux systems for installation", InfoCenter:

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r2mx/index.jsp?topic=/com.ibm.btools.help.monitor.install.doc/plan/prep_os_lin.html

2.2 Time server setup

To avoid conflicts in a production environment caused by date/time mismatch in a cell, all machines running WebSphere components as well as the system hosting the database system are set up to receive their system date and time from a time server:

1. Logon as root
   (if you are not using the native system console ensure x11 forwarding is enabled)
2. Start the graphical Configuration Tool /usr/bin/system-config-time
3. Select Tab Network Time Protocol and set
4. Servers (primary, secondary, tertiary): 0.rhel.pool.ntp.org, 1.rhel.pool.ntp.org, 2.rhel.pool.ntp.org
5. Enable Network Time Protocol: selected
6. Click Ok.

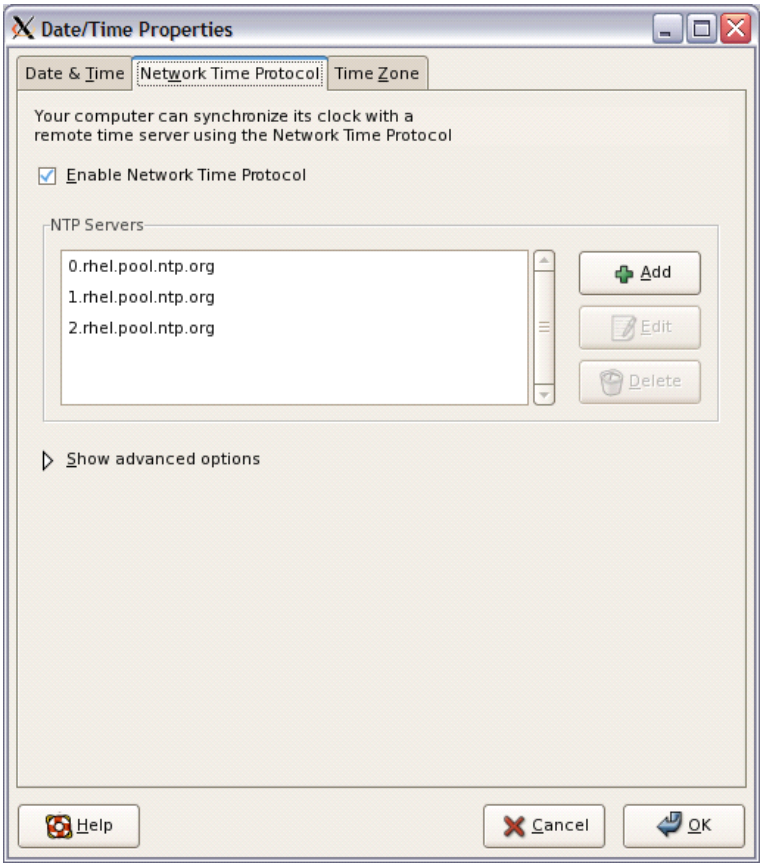

2.3 Select the type of user directory

To demonstrate how to use ldap an existing user directory server IBM Tivoli Directory Server 6.1. will be used. To keep administration simple, security will be set up with a single user ID "**vmmuser**" of group "**FVTStaff**". Additional LDAP credentials are:

Property	Value
Directory type	Tivoli Directory Server Version 6.1
LDAP host	ldap.boeblingen.de.ibm.com
Port	389
Server user ID	vmmuser
Server user password	********
Server user group	ou=StaffTest,cn=FVTStaff,o=ibm,c=us
Base distinguished name (DN)	o=ibm, c=us
Bind distinguished name (User DN)	cn=ldapuser,o=ibm,c=us
Bind password	
Anonymous bind	Allowed

A useful utility to explore a LDAP directory is:

LDAP Explorer Tool on Sourceforge (available for Windows and Linux)

http://sourceforge.net/projects/ldaptool/

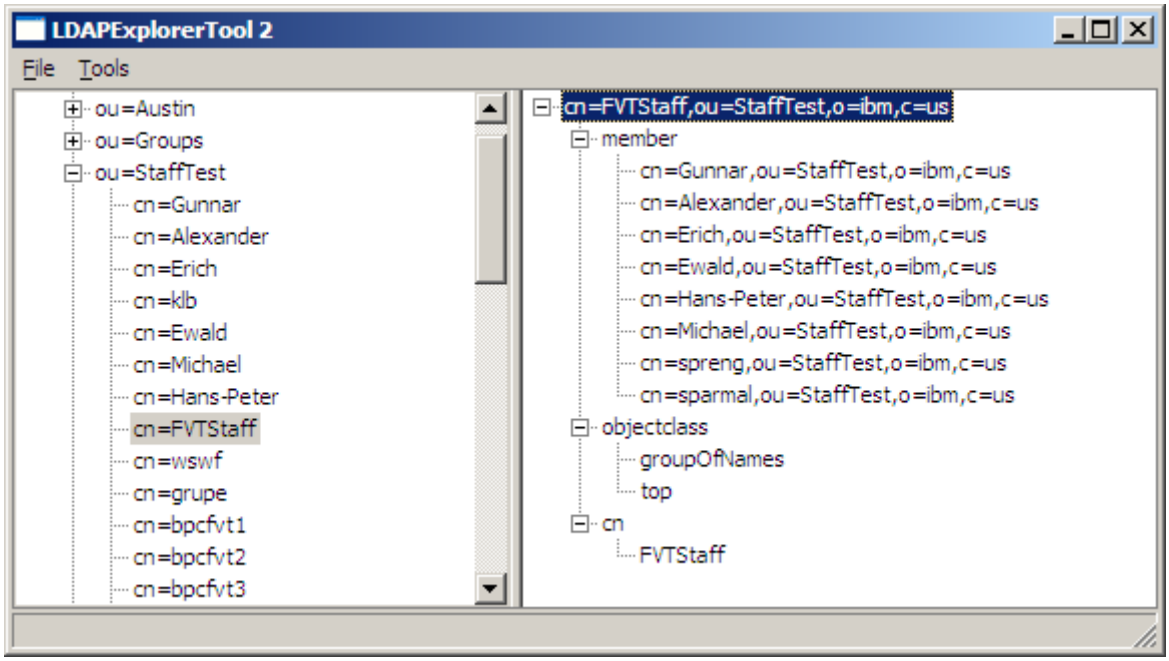

2.4 Software packages

WebSphere Process Server V6.2.0.0 (GA CDs)

Use the installation CDs or download the following file from Passport Advantage:

Xtreme Leverage > Software Downloads> Find by Part Number

WebSphere Process Server V6.2 Linux X86 32Bit
Multilingual (C1M4XML)

- C1M4XML.tar.gz

WebSphere Business Monitor V6.2.0.0 (GA CDs)

Use the installation CDs or download the following file from Passport Advantage:

Xtreme Leverage > Software Downloads> Find by Part Number

WebSphere Business Monitor V6.2 Linux X86 32Bit
Multilingual (C1LK5ML)

- C1LK5ML.tar.gz

WPS, WBM and WAS Fixes (6.2.0.1 / 6.1.0.23)

WPS and WBM Fixpack 6.2.0.1 is applied. This includes WAS 6.1.0.23

Oracle 11g Release 1 (11.1.0.6.0)

Get the Oracle database installation media for example:

- linux.x64_11gR1_database_1013.zip

For the installation of the products the files have been copied to the local directory /root/SWINST on each machine.

2.5 Installation properties

2.5.1 User IDs

To keep the setup simple, the same user IDs will be used for dedicated functions.

UserType	Username	Role
Linux User	root	Linux system user and WebSphere Process Server functional user ID
Linux DB User	oracle	Linux system user for the database installation
WPS User	vmmuser	Central user ID for WebSphere components, also SCA / JMS userid and BPE/HTM Administration role userid
LDAP User	ldapuser	Central user ID for the ldap directory server in our case the server ldap.boeblingen.de.ibm.com

2.5.2 Directory locations

Directory	Brief description
/root/SWINST	Location of install sources (WebSphere Process Server, Oracle, DB2, scripts, etc)
/opt/oracle	Oracle Product install directory
/opt/oracle/driver	Oracle JDBC driver directory
/WPS62	WebSphere Product install directory on W6201L3M, W6201LN1, W6201LN2
/WPS62	WebSphere Business Monitor install directory on W6201L3M
/WBM62	WebSphere Business Monitor install directory on W6201LN3, W6201LN4
/WPS62/profiles/W6201L3MBPMDmgr	Profile directory deployment manager
/WPS62/profiles/W6201LN1WPSCustom01	Profile directory WPS custom profile on W6201LN1
/WPS62/profiles/W6201LN2WPSCustom01	Profile directory WPS custom profile on W6201LN2
/WBM62/profiles/W6201LN3WBMCustom01	Profile directory WBM custom profile on W6201LN3
/WBM62/profiles/W6201LN4WBMCustom01	Profile directory WBM custom profile

	on W6201LN4
/opt/Alphablox	Alphablox install directory on W6201LN3, W6201LN4
/opt/ihs7	IBM HTTP Server install directory on W6201L3M

Note: Depending on your local needs, these directory names might be different.

2.5.3 Hostnames of the involved systems

Hostname	Brief description
W6201L3M	Deployment manager
W6201L3M	IBM HTTP Server
W6201L3M	Proxy Server
W6201LN1	WPS Node 1 in the cell
W6201LN2	WPS Node 2 in the cell
W6201LN3	WBM Node 1 in the cell
W6201LN4	WBM Node 2 in the cell
W6201L3O	Oracle Database System
LDAP	IBM Tivoli Directory Server

Note: Depending on your local needs, the amount of used hosts might be different.

2.5.4 Naming of cluster components

Name	Brief description
W6201L3MBPMDmgr	Profile name deployment manager (WPS-DMGR)
W6201LN1WPSCustom01	Profile name of WPS custom node 01 (W6201LN1WPSNode01)
W6201LN2WPSCustom01	Profile name of WPS custom node 02 (W6201LN2WPSNode01)
W6201LN3WBMCustom01	Profile name of WBM custom node 01 (W6201LN3WBMNode01)
W6201LN4WBMCustom01	Profile name of WBM custom node 02 (W6201LN4WBMNode01)
Cell01	Cell name
W6201L3MBPMDmgr	Node name Deployment Manager
W6201LN1WPSNode01	WPS Node name node 1
W6201LN2WPSNode01	WPS Node name node 2
W6201LN3WBMNode01	WBM Node name node 1
W6201LN4WBMNode01	WBM Node name node 2
BPELCluster	Cluster name BPEL cluster
SupportCluster	Cluster name Support cluster

MECluster	Cluster name Messaging Engine cluster
MonSupportCluster	Cluster name Monitor Support cluster
MonApplicationCluster	Cluster name Monitor Moderator cluster
WebDashboardCluster	Cluster name Web Dashboard cluster
BPELCluster_Member0x	Member name BPEL cluster
SupportCluster_Member0x	Member name Support cluster
MECluster_Member0x	Member name Messaging Engine cluster
MonSupportCluster_Member0x	Member name Monitor Support cluster
MonModeratorCluster_Member0x	Member name Monitor Moderator cluster
MonLogicCluster_Member0x	Member name Monitor Logic cluster
WebDashboardCluster_Member0x	Member name Web Dashboard cluster
w6201l3m.boeblingen.de.ibm.com	Hostname deployment manager, http server, proxy server
w6201ln1.boeblingen.de.ibm.com	Hostname WPS custom profile 1
w6201ln2.boeblingen.de.ibm.com	Hostname WPS custom profile 2
w6201ln3.boeblingen.de.ibm.com	Hostname WBM custom profile 1
w6201ln4.boeblingen.de.ibm.com	Hostname WBM custom profile 2
w62l3ora.boeblingen.de.ibm.com	Hostname for the Oracle database

Chapter 3 WebSphere Process Server Cluster installation and configuration

3.1 WebSphere Process Server Cluster installation and configuration primer

On a high level, the following steps are necessary to create a clustered WebSphere Process Server (WPS). They are described in detail in the following chapters (A checklist can be found in Appendix):

- Product installation
  - Installation of the WPS binaries (including available fix packs and all iFixes (for WAS ND 6.1 and WPS 6.2.0)
- Installation of the database system and database creation
  - Create WPS server deployment manager profile
  - Create the WPS tablespaces. users and schema objects
  - Create WPS custom profiles
  - Federate the WPS custom nodes
  - Enable security in the cell
- Cluster configuration
  - Create the messaging engine cluster (MECluster)
  - Configure SCA on the MECluster
  - Create the SupportCluster
  - Configure Common Event Infrastructure and Business Rules Manager on SupportCluster
  - Create the BPEL cluster
  - Configure BPEL cluster for BPC and HTM
  - Configure HTTP and Proxy Server
  - Configure BPC Explorer on SupportCluster
  - Configure Reporting Function on SupportCluster
  - Configure Business Space endpoint XML files
- Optional: Install and configure the IBM HTTP Server

3.2 Installing WebSphere Process Server binaries

NOTE: The steps shown here have to be executed on each node that is supposed to contain WPS functionality; In this case this are the machines W6201L3M, W6201LN1 and W6201LN2. Make sure to execute this step as user root.

If you are installing from the product DVD, mount the DVD and change to the mount_point directory.

If you are installing from a downloaded image from Passport Advantage extract the image and change to the extract directory. There are two options on how to install the WPS binaries:

1. graphical wizard
2. silent mode

> Note: Installing the WPS binaries in silent mode is not described within this document

To start the installation with the graphical wizard, enter the following (if you are not locally working on the machine, make sure X11 tunneling is activated and an X server is running on your machine):

```
cd /<WPS62_EXTRACT_ROOT>/WBI
./install
```

Now the graphical wizard starts...

The "WebSphere Process Server 6.2 Welcome" panel is displayed:

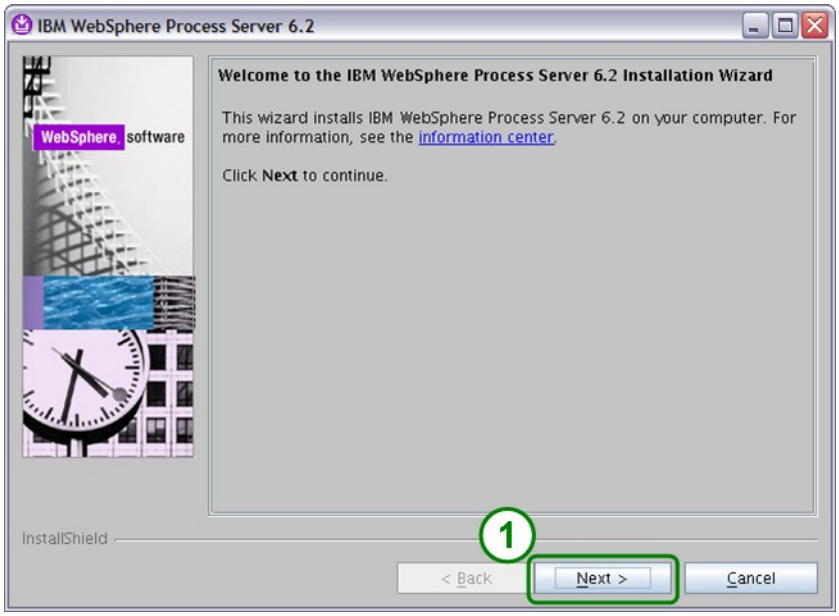

1. Press **Next**.

The "Software License Agreement" panel is displayed:

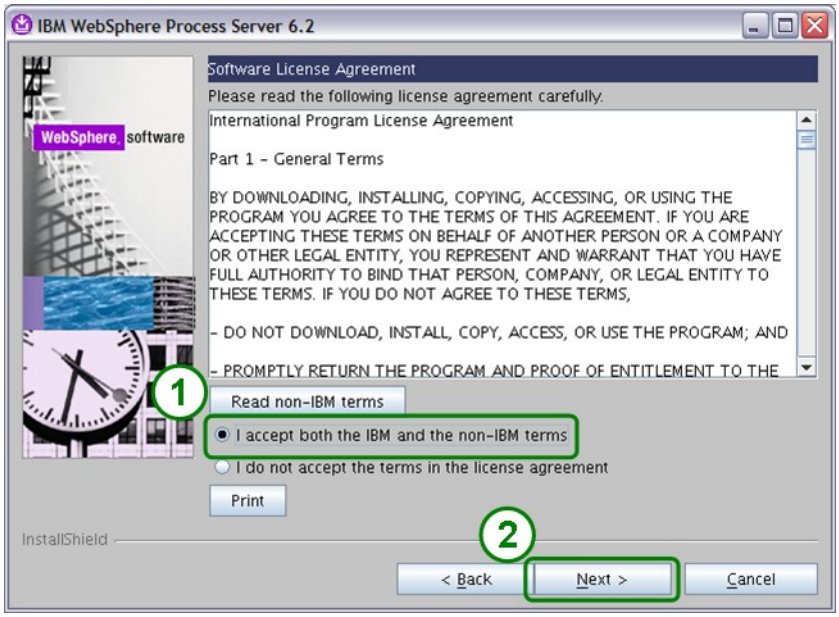

1. Select "I accept both the IBM and the non IBM terms" to accept the license agreement.

2. Press **Next**.

The "System Prerequisites Check" panel is displayed:

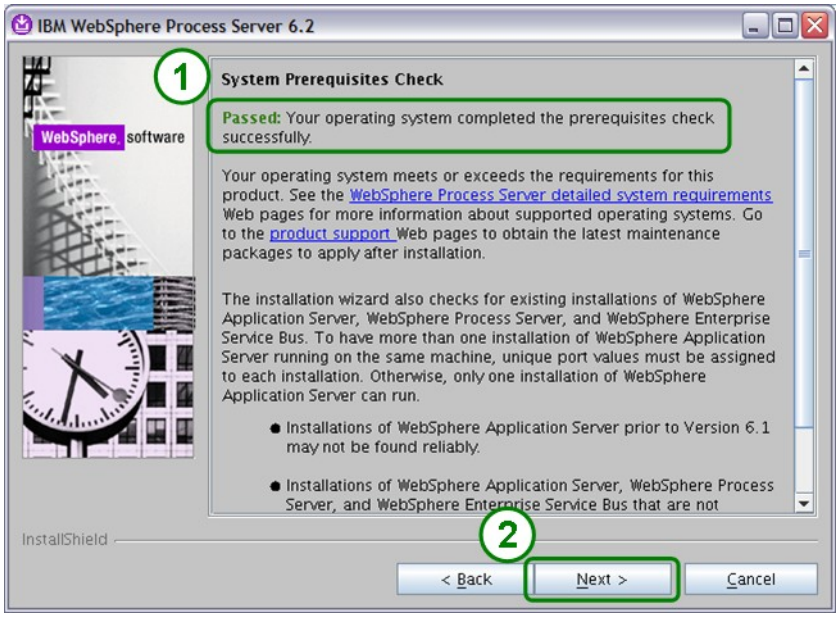

1. Verify that the System Prerequisites Check has passed.

2. Press **Next**.

The "Installation Type Selection" panel is displayed:

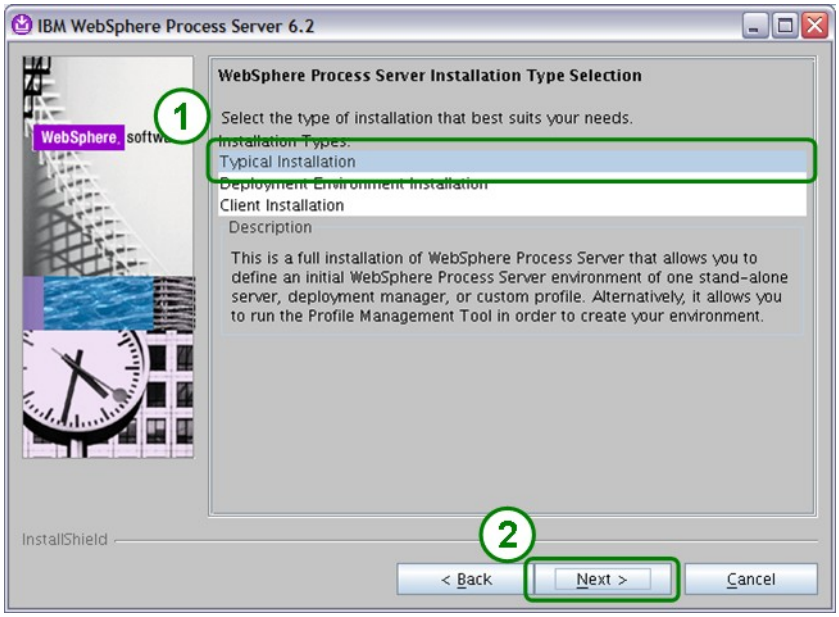

1. Select “Typical Installation” in the Installation Types selection box.

2. Press **Next**.

The "Features Selection" panel is displayed:

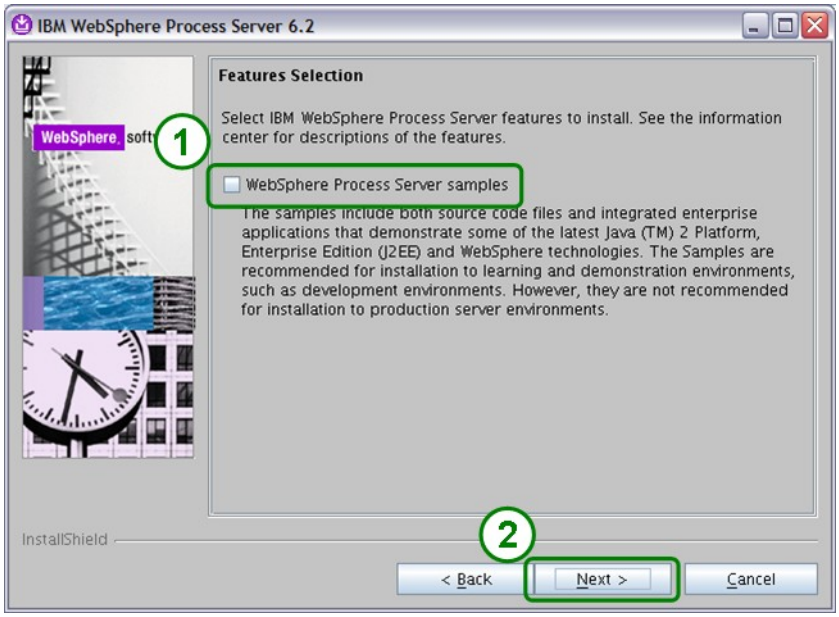

1. "Webshere Process Server samples" is not used so make sure it is NOT selected.

2. Press **Next**.

The "Installation Root Directory" panel is displayed:

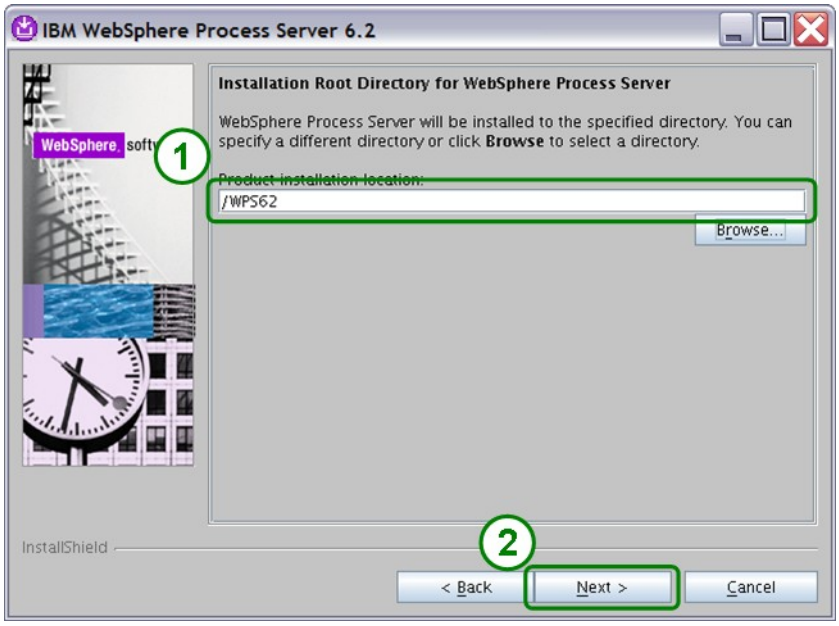

1. Specify the "Product installation directory".

2. Press **Next**.

The "Environments" panel is displayed:

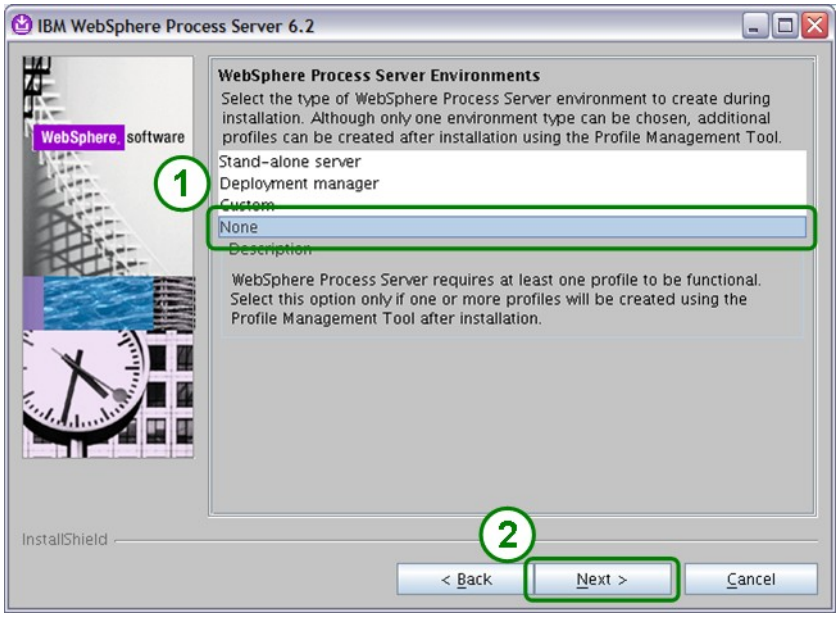

1. Select “None” for the type of WebSphere Process Server Environment. The profiles will be created later with the pmt tool.

2. Press **Next**.

A Warning panel pops up:

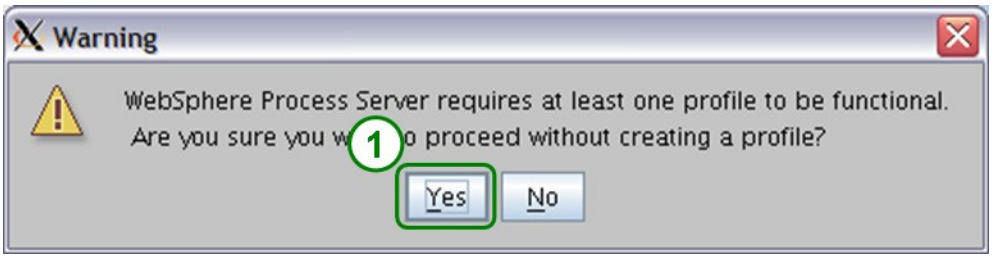

1. Press the **Yes** button to proceed without creating a profile.

The "Installation Summary" panel is displayed:

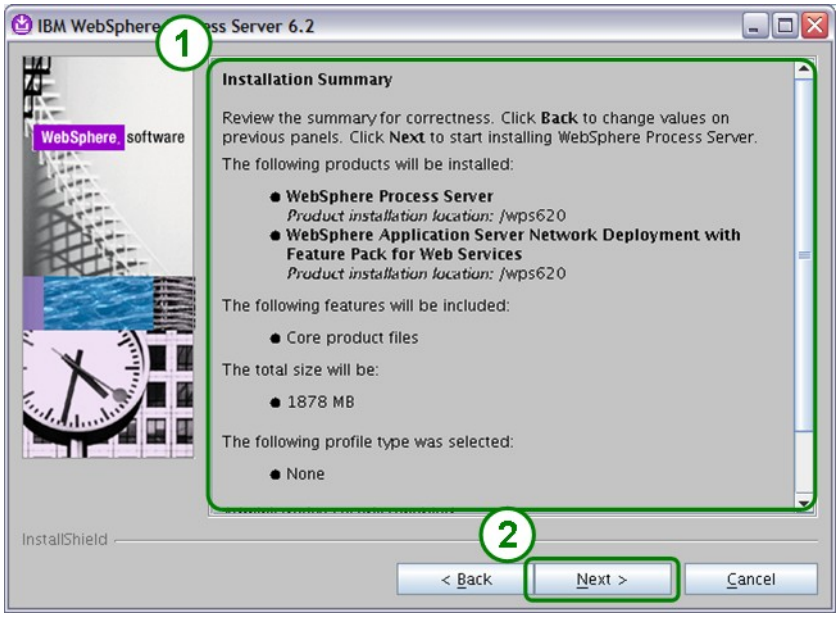

1. Verify the summary.

2. Press the **Next** button to start the installation.

The "installation" starts ...

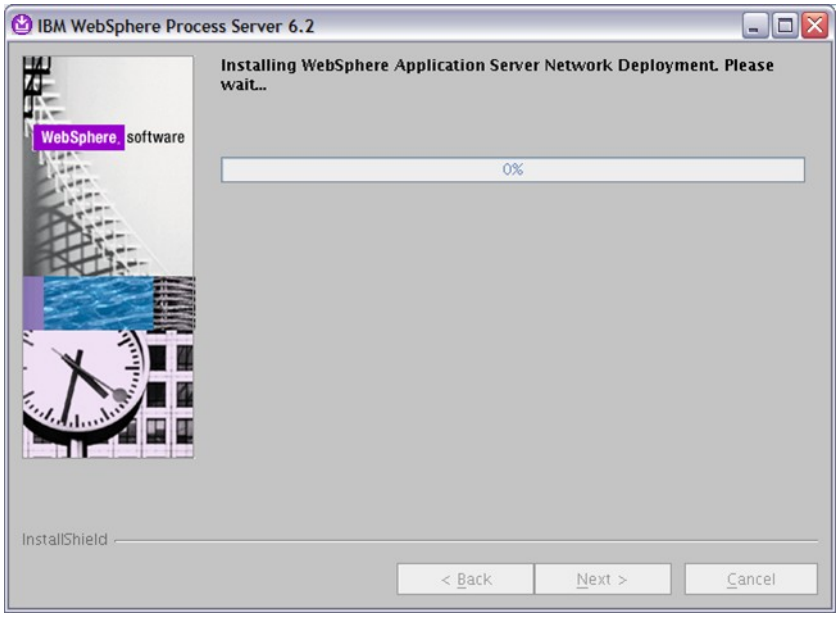

... wait until the installation finishes.

The "Installation Results" panel is displayed:

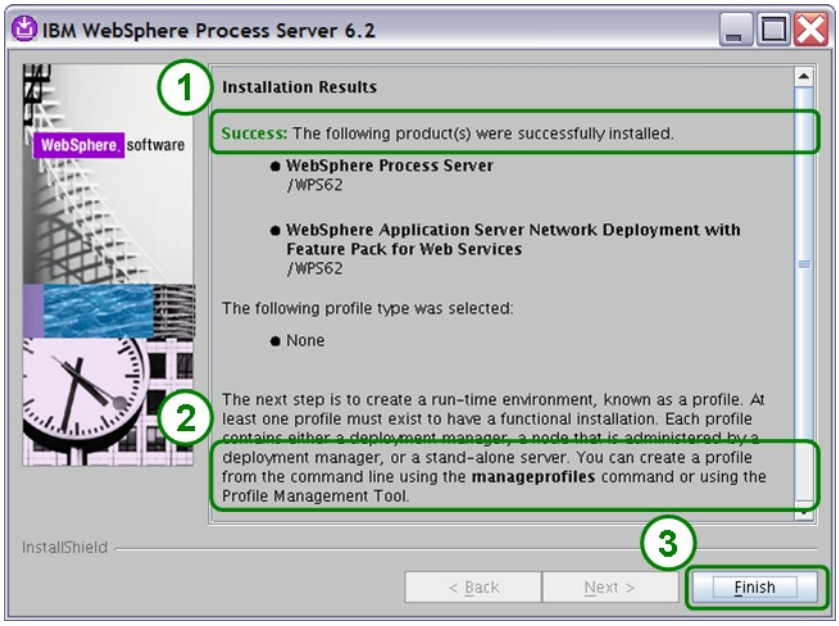

1. Verify that the installation was successful.

2. Scroll down and unselect the “Create a new WebSphere Process Server profile”. The profiles will be created later with the pmt tool.

3. Press **Finish**.

3.3 Installing the Update Installer

Get the latest Update Installer from:

http://www.ibm.com/support/docview.wss?rs=180&uid=swg24020446

and extract it.

Note: The current version at the time of writing this document is v7.0.0.3

There are two options on how to install the Update Installer:

1. graphical wizard
2. silent mode

Note: Installing the Update Installer in silent mode is not described within this document

Start the Installation Wizard for the Update Installer with the following command:

```
cd /UPDI_INSTALLER_EXTRACT_ROOT/UpdateInstaller
./install
```

The "Installation Wizard Welcome" panel is displayed:

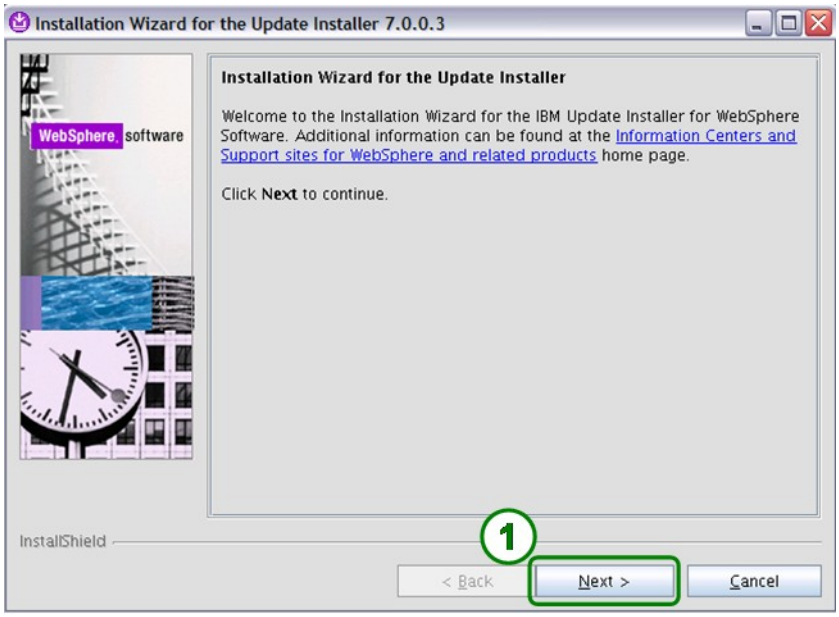

1. Click **Next**.

The "Software License Agreement" panel is displayed:

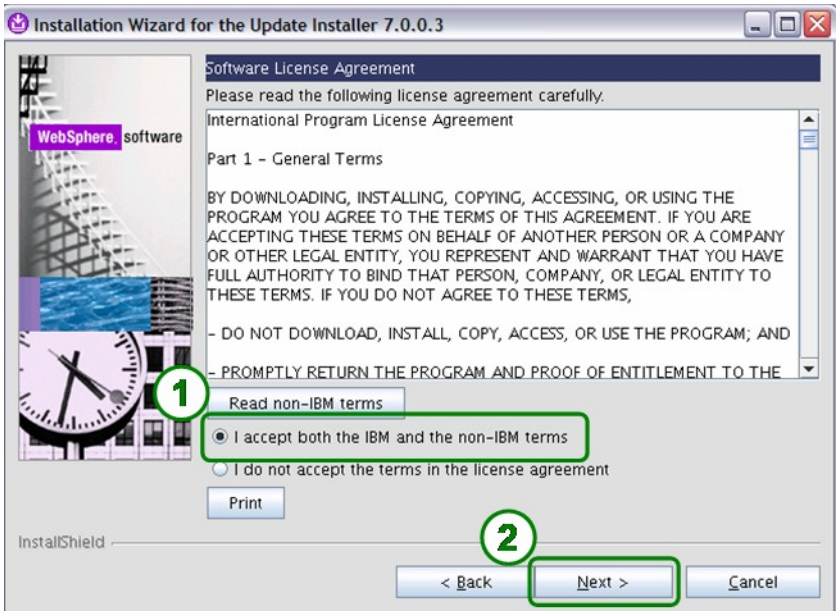

1. Select "I accept both the IBM and the non-IBM terms"

2. Press **Next**.

The "System Prerequisite Check" panel is displayed:

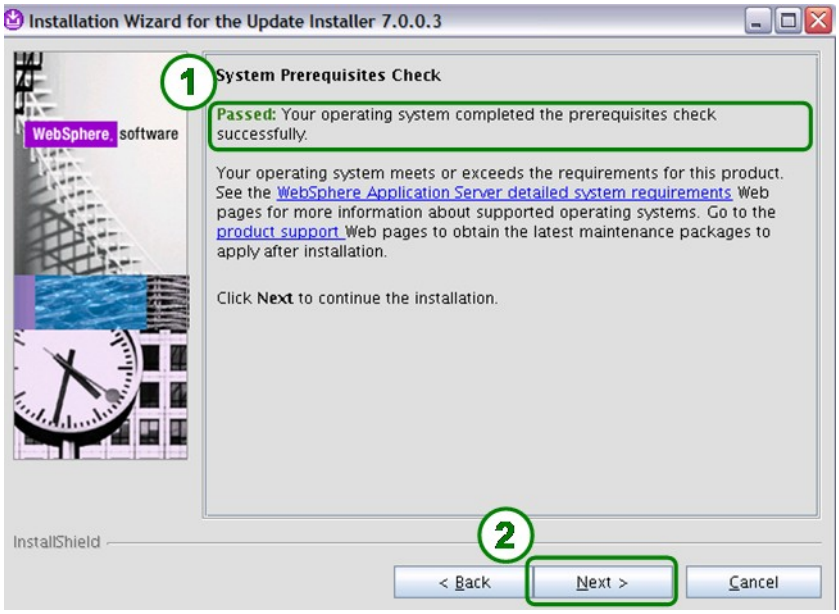

1. Verify that the Status of the System Prerequisites Check is passed.

2. Press **Next**.

The "Installation directory" panel is displayed:

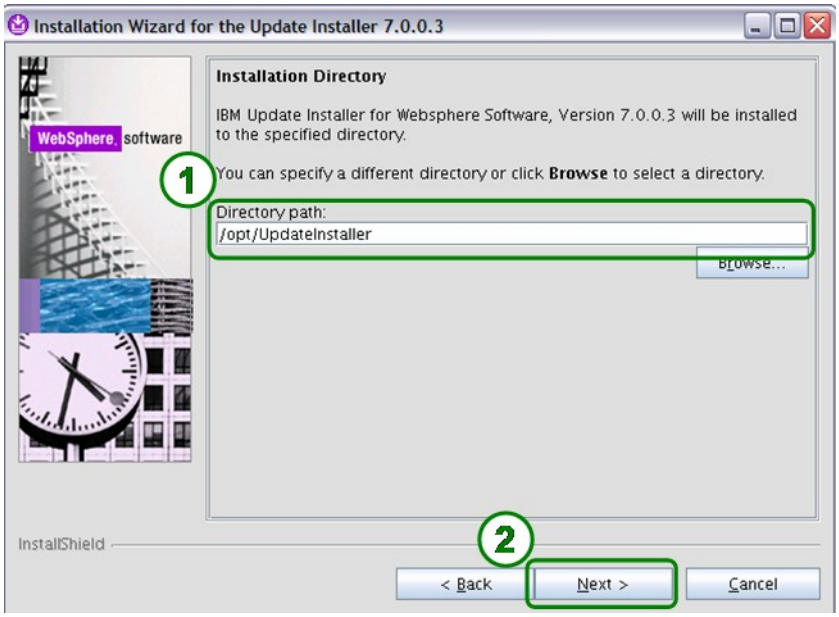

1. Enter the installation directory path.

2.Press **Next**.

The "Installation Summary" panel is displayed:

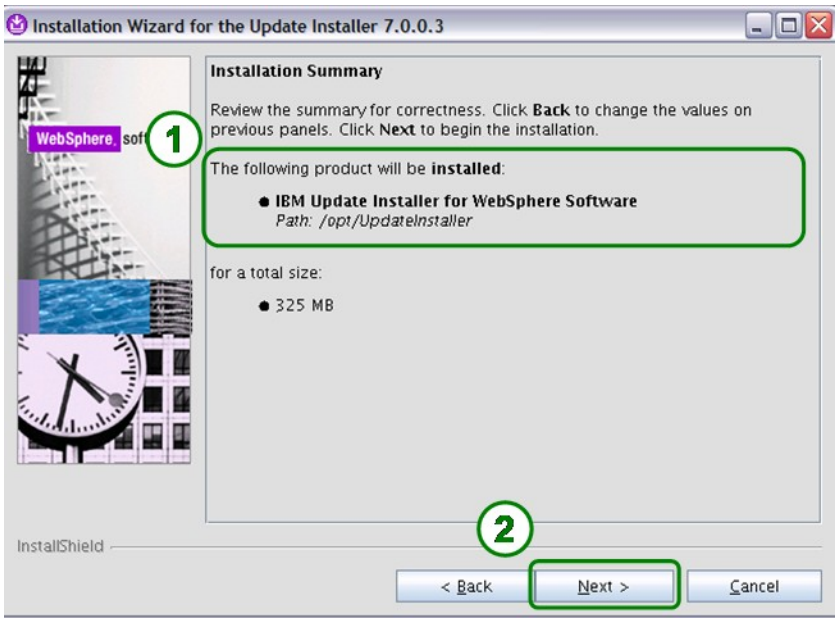

1. Verify that the Update Installer gets installed in the proper directory.

2.Press **Next**.

The "Installation Complete" panel is displayed after the installation has finished:

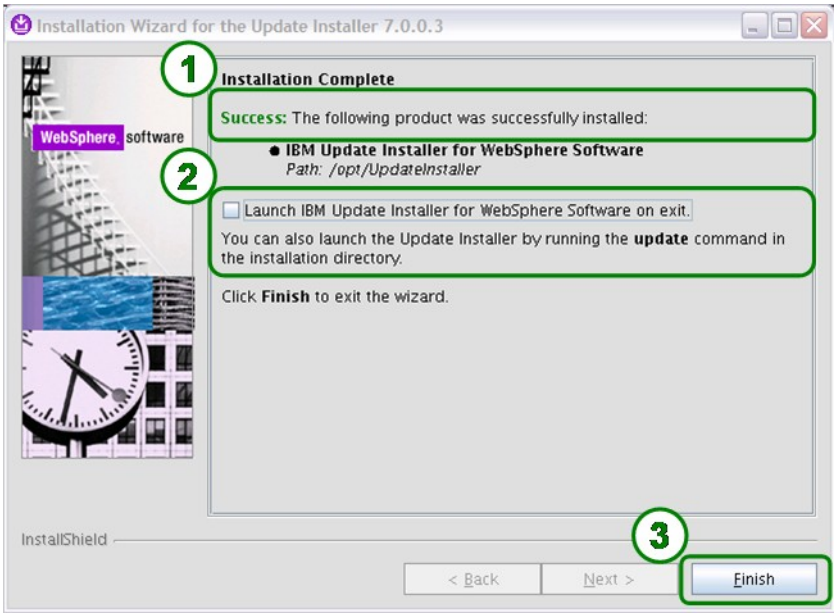

1. Check whether the installation was successful.

2. De-Select "Launch IBM Update Installer on exit".

3. Press **Finish**.

3.4 Installing mandatory fixes

Get the latest mandatory fixpack (Fixpack 1) from:

```
http://www.ibm.com/support/docview.wss?
rs=2307&context=SSQH9M&dc=DB600&uid=swg21383890&loc=en_US&cs=UTF-8&lang=en
```

and and extract it into the Update Installer maintenance directory {UpdateInstaller_Home}/maintenance.

> Note: The current fixpack at the time of writing this document is Fixpack 1 (v6.2.0.1)

There are two options on how to apply the fixpack:

1. graphical wizard
2. silent mode

> Note: Applying the fixpack using the silent mode is not described within this document.

Start the Update Installer with the following command:

```
cd /opt/UpdateInstaller/
./update.sh
```

The "Welcome to the IBM Update Installer for WebSphere Software wizard" panel is displayed:

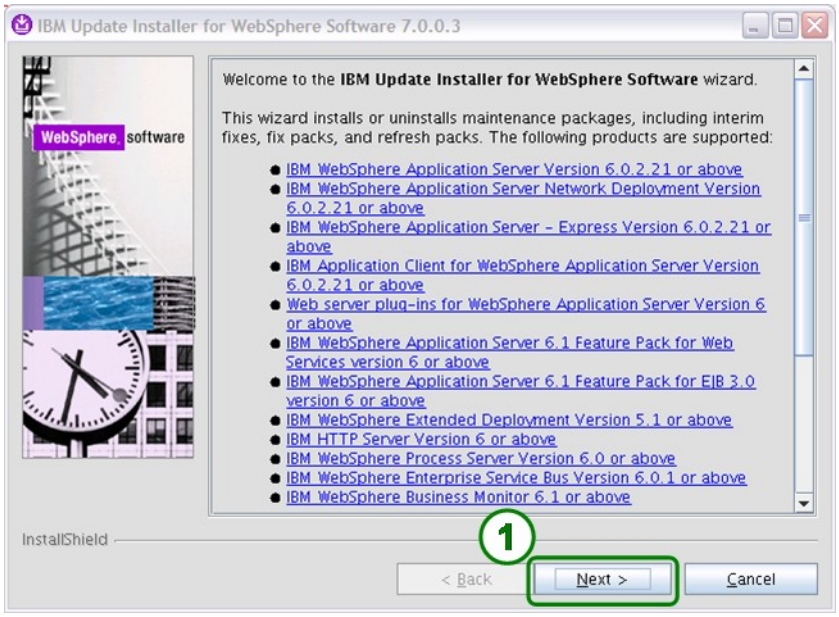

1.Press **Next**.

The "Product Selection" panel is displayed:

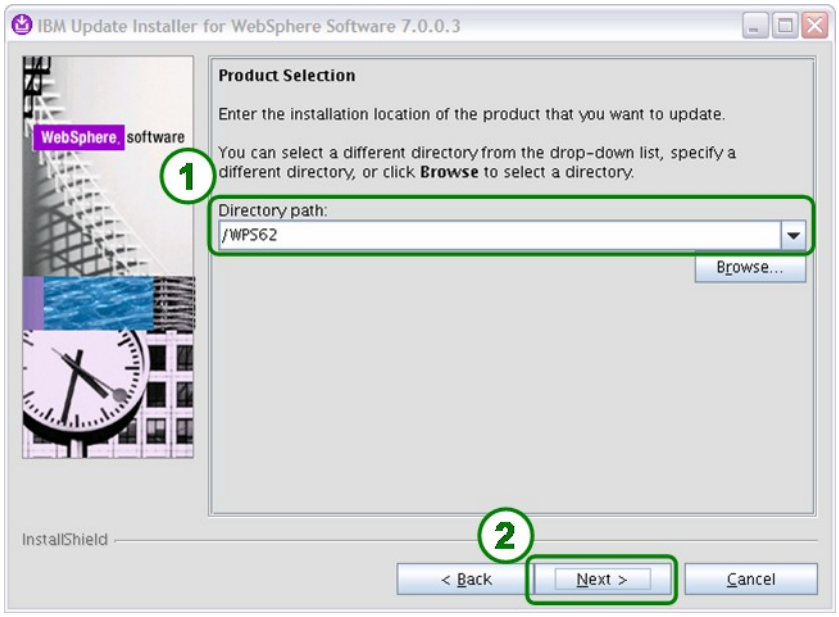

1. Enter the appropriate directory path.

2. Click **Next** .

The "Maintenance Operation Selection" panel is displayed:

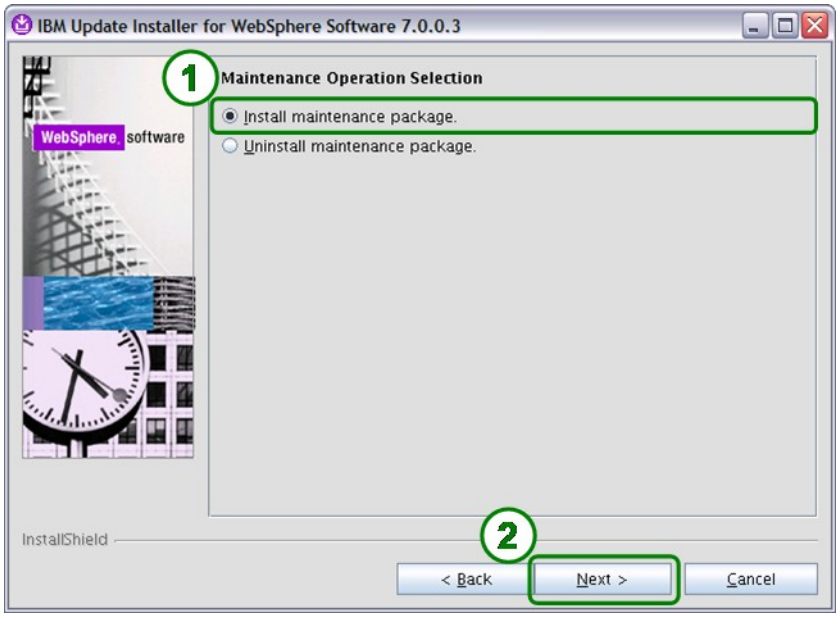

1. Check that Radio-button "Install maintenance package" is selected

2. Click **Next**.

The "Maintenance Package Directory Selection" panel is displayed:

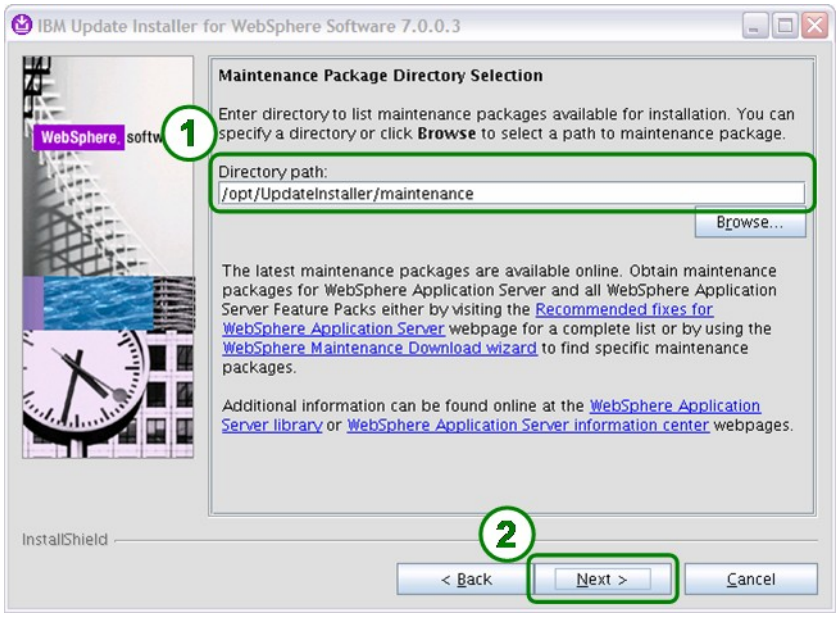

1. Check the Directory path

2. Click **Next**.

The "Available Maintenance Packages to Install" panel is displayed:

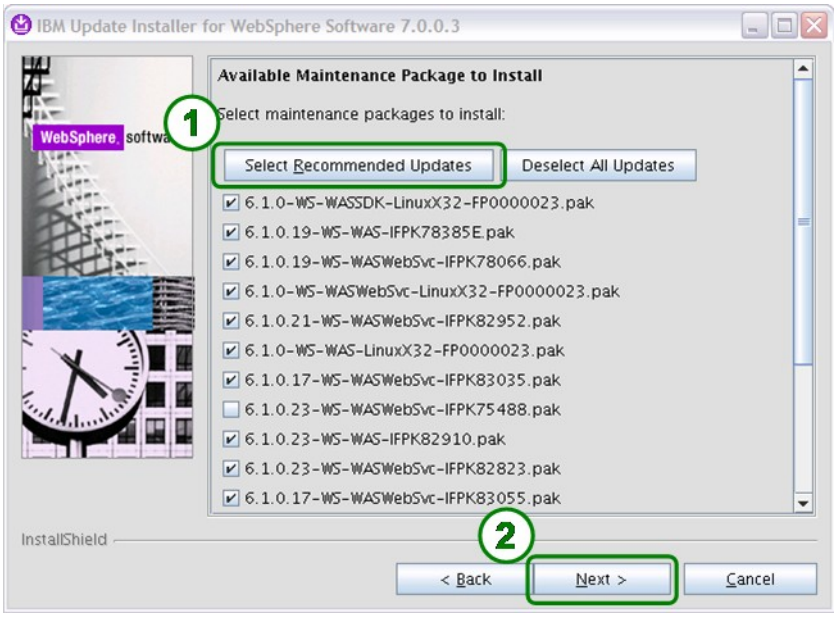

1. Press "Select Recommended Updates"

2. Click **Next**.

The "Warning" panel is displayed:

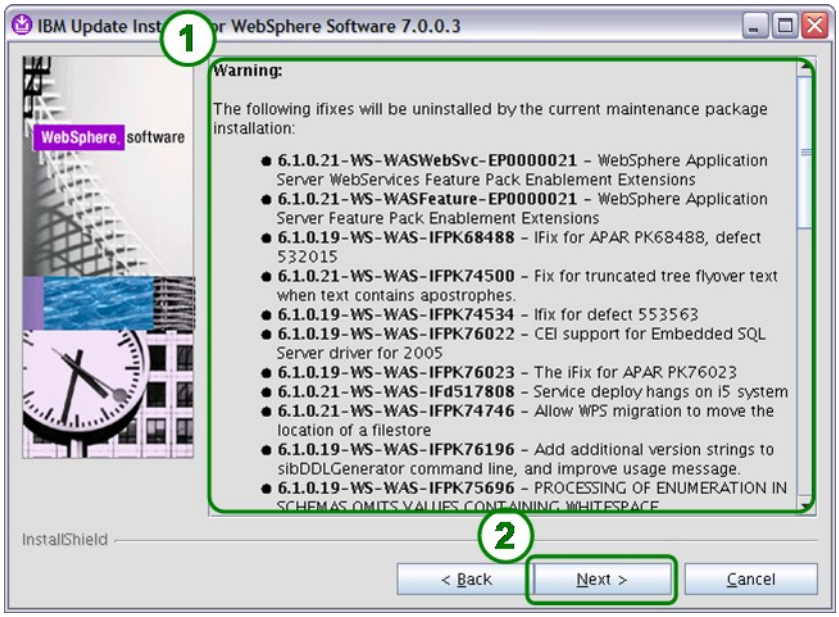

1. Notice that several ifixes will be uninstalled by the current maintenance package installation

2. Click **Next**.

The "Additional Maintenance Required for Installed WebSphere Software" panel is displayed:

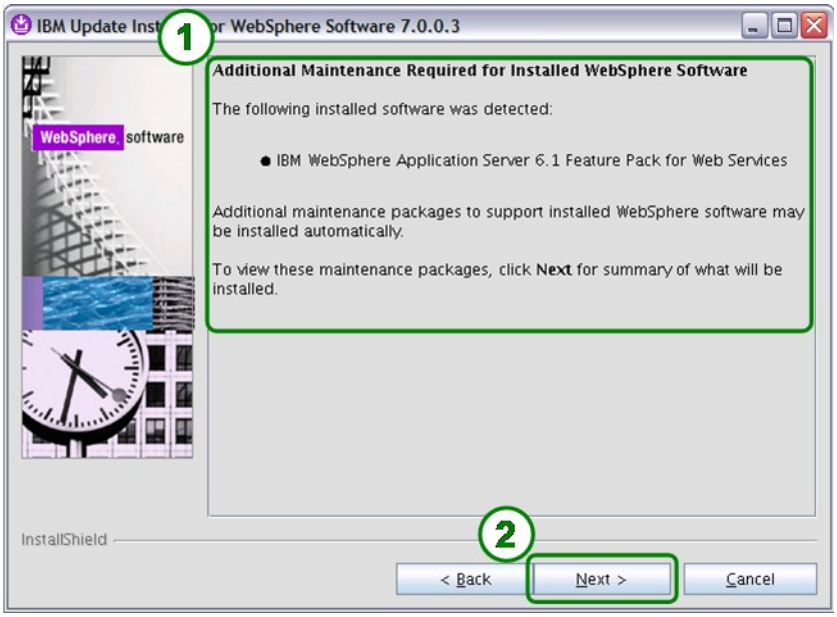

1. Notice that additional maintenance packages may be installed automatically"

2. Click **Next**.

The "Installation Summary" panel is displayed:

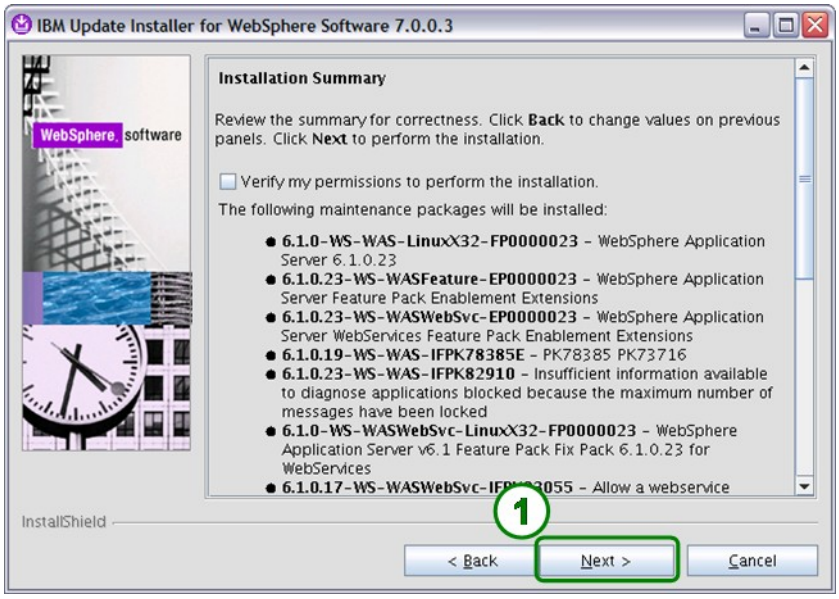

1. Click **Next**. This will start the installation of the fixpack.

The "Installation Complete" panel is displayed:

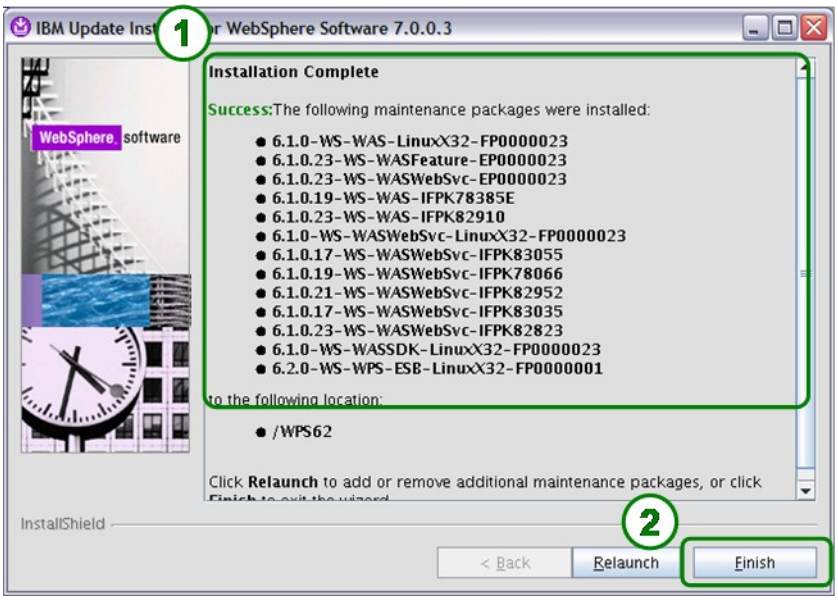

1. Verify that the ifixes have been installed successfully

2. Click **Finish**.

3.5 Verify WebSphere Process Server binary installation

Verify the success of the binary installation by examining the WebSphere Process Server log files. If the last line of the file contains the word **INSTCONFSUCCESS**, the selected WebSphere Process Server features were installed successfully.

The log file is located as follows:

```
/WPS62/logs/install/log.txt
```

The log file of each fix can be found under in directory:

```
/WPS62/logs/update/install/updatelog.txt
```

You can also use the IVT (Installation Verification Tool) to check if the binaries have been installed correctly. See the infos on how to do that here:

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r2mx/index.jsp?topic=/com.ibm.websphere.wps.620.doc/doc/tins_ivt.html

Part III Database System

Chapter 4 Oracle binary installation

This chapter covers the Oracle binary installation and its prerequisites on an Red Hat Enterprise Linux Version 5 system.

4.1 Oracle installation prerequisites

This chapter describes all Linux OS and user prerequisites for the Oracle installation on a Red Hat Enterprise Linux Version 5 system.
All information in the following sections were gathered from the **Oracle® Database Installation Guide 11g Release 1 (11.1) for Linux** Part Number B32002-06:

http://download.oracle.com/docs/cd/B28359_01/install.111/b32002/toc.htm

Before Oracle can be installed on a Linux operating system all requirements must be met and all preparation tasks must be finished successfully.

4.1.1 Oracle operating system requirements

This section describes the operating system requirements for Oracle 11g. These requirements will be checked during the Oracle installation process and marked as failed or succeeded. If a prerequisite fails in the check go back to this chapter and check the failed prerequisites. Continuing the installation with a failed prerequisite can lead to a unrecoverable installation failure.

4.1.1.1 Required Linux packages for Oracle

An Oracle 11g installation requires several Linux software packages (RPMs) to be installed on the RHEL5. Most of these packages have automatically been installed during the actual Linux installation. However, to avoid potential problems at a later point in time, it is recommended to check that the RPMs listed below are installed. Note that these packages must be at least at the level of the listed version.

RPMs required for Oracle 11g on a RHEL5 64-bit system - note that on the 64bit system, several software packages are required in both, the 32-bit and 64-bit versions:

```
binutils-2.17.50.0.6
compat-libstdc++-33-3.2.3
compat-libstdc++-33-3.2.3 (32 bit)
elfutils-libelf-0.125
elfutils-libelf-devel-0.125
gcc-4.1.1
gcc-c++-4.1.1
glibc-2.5-12
glibc-2.5-12 (32 bit)
glibc-common-2.5
glibc-devel-2.5
glibc-devel-2.5-12 (32 bit)
libaio-0.3.106
libaio-0.3.106 (32 bit)
libaio-devel-0.3.106
libgcc-4.1.1
libgcc-4.1.1 (32 bit)
libstdc++-4.1.1
libstdc++-4.1.1 (32 bit)
libstdc++-devel 4.1.1
make-3.81
sysstat-7.0.0
```

Note: The listed packages without the additional (32 bit) mark are 64 bit versions.

RPMs required for Oracle 11g on a RHEL5 32-bit system:

```
binutils-2.17.50.0.6
compat-libstdc++-33-3.2.3-61
elfutils-libelf-0.125
elfutils-libelf-devel-0.125
glibc-2.5-12
glibc-common-2.5-12
glibc-devel-2.5-12
glibc-headers-2.5-12
gcc-4.1.1-52
gcc-c++-4.1.1-52
libaio-0.3.106
libaio-devel-0.3.106
libgcc-4.1.1-52
libstdc++-4.1.1
libstdc++-devel-4.1.1-52
make-3.81-1.1
sysstat-7.0.0
unixODBC-2.2.11
unixODBC-devel-2.2.11
```

To verify that the required software packages are installed, enter the following command as user root:

```
rpm -qa |grep <package_name>
```

For example:

```
rpm -qa |grep binutils
```

To install the required software packages , enter the following command as user root:

```
rpm -ivh <package_name_version>
```

For example:

```
rpm -ivh binutils-2.17.50.0.6
```

4.1.1.2 Linux kernel requirements

Oracle 11g requires a Linux kernel version of at least 2.6.18.

To verify the installed Linux kernel version on your Oracle DB host maschine execute the following command as user root:

```
uname -r
2.6.18-92.1.18.el5
```

4.1.1.3 Linux kernel parameters

The Linux kernel parameters listed below show the values required for the Oracle 11g database. It is recommended to tune theses parameters for production environments.

For further information on tuning please refer to: **Oracle® Database Performance Tuning Guide 11*g* Release 1 (11.1)** Part Number B28274-02:

http://download.oracle.com/docs/cd/B28359_01/server.111/b28274/toc.htm

The following parameters and values are required as a minimum for the Oracle 11g installation:

Kernel parameter	Value	Description
kernel.shmmax	0.5 x RAM size or 4GB – 1byte (depends on which is lower)	shmmax is the maximum size of one shared memory segment.
kernel.shmmni	4096	shmmni is the maximum number of shared memory segments
kernel.shmall	2097152	shmall specifies the maximum number of shared memory pages

(continued)

Kernel parameter		Value	Description
kernel.sem	**semmsl**	250	semmsl is the maximum number of semaphores per semaphore set
	semmns	32000	semmns is the the total number of semaphores
	semopm	100	semopm is the maximum number of semaphore operations that can be performed per semop(2) system call.
	semmni	128	semmni is the maximum number of semaphore sets in the entire Linux system
fs.file-max		102696	file-max is the maximum of file handles that the Linux kernel will allocate
net.ipv4.ip_local_port_range		1024 65000	ip_local_port_range defines the local port range used by TCP und UDP traffic to choose the local port
net.core.rmem_default		4194304	rmem_default is the receive socket buffer's default size
net.core.rmem_max		4194304	rmem_max is the receive socket buffer's maximum size
net.core.wmem_default		262144	wmem_default is the send socket buffer's default size
net.core.wmem_max		262144	wmem_max is the send socket buffer's maximum size

Important: If the current system values are higher than the ones listed above, it is recommended to leave them as they are.

This is not valid for the SHMMAX parameter which needs to be calculated.

To check the values issue the following command as user root:

```
/sbin/sysctl -a | grep <kernel_parameter>
```

For example:

```
/sbin/sysctl -a | grep kernel.shmmax
```

```
kernel.shmmax = 1892380672
```

To change the values, edit /etc/sysctl.conf as user root and modify this file as shown below:

```
# --------------------------------------------------------------------
#  Linux kernel parameter for Oracle 11g
# --------------------------------------------------------------------
kernel.shmmax = <add value here>
kernel.shmmni = 4096
kernel.shmall = 2097152
kernel.sem = 250 32000 100 128
fs.file-max = 102696
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 4194304
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 262144
```

After completing the changes, activate them either by issuing the following command as user root:

```
/sbin/sysctl -p
```

```
...
kernel.shmmax = <added value>
kernel.shmall = 2097152
kernel.sem = 250 32000 100 128
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 4194304
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 262144
fs.file-max = 102696
...
```

or reboot the operating system.

4.1.1.4 Memory requirements

Oracle 11g requires a minimum of 1GB RAM memory.

4.1.1.5 Disk space requirements

The following are the disk space requirements for installing Oracle Database 11g Release 1.

4.1.1.5.1 Disk requirements in the /tmp directory

Between 150 and 200 MB of disk space in the /tmp directory is needed.
To determine the amount of disk space available in the /tmp directory, enter the following command:

```
df -k /tmp
```

If there is less than 400 MB of free disk space available in the /tmp directory, then complete one of the following steps:

- Delete unnecessary files from the /tmp directory to meet the disk space requirement.
- Set the TMP and TMPDIR environment variables when setting the oracle user's environment (described in Linux operating system user for Oracle).
- Extend the file system that contains the /tmp directory. If necessary, contact the system administrator for information about extending file systems.

4.1.1.5.2 Disk space requirements on the system

To determine the amount of free disk space on the system, enter the following command:

```
df -k
```

Between 3.5 GB and 5 GB of disk space for the Oracle software, depending on the installation type are needed. The following table describes the disk space requirements for software files for each installation type:

Installation Type	Requirement for Software Files (GB)
Enterprise Edition	4.35
Standard Edition	3.73
Custom (maximum)	4.54

The following table describes the disk space requirements for each installation type:

Installation Type	Requirement for Datafiles Files (GB)
Enterprise Edition	1.68
Standard Edition	1.48
Custom (maximum)	2.14

Additional disk space, either on a file system or in an Automatic Storage Management disk group, is required for the flash recovery area if you choose to configure automated backups.

4.1.1.5.3 Disk space requirements for WebSphere Process Server

The Oracle database for WebSphere Process Server (WPS) needs between 4.0 GB and 5,5 GB additional disc space for the Oracle system data files and redo logs.
This disk space requirement can raise depending on the amount of Processes and running process instances.
The following table describes the additional disc space requirements for the Oracle system data files and redo logs:

Installation Type	Additional Requirement for Datafiles Files (GB)
WPS stand-alone server	~ 4.00
WPS Cluster	~ 6.00

4.1.1.6 Swap Space requirements

Oracle 11g requires the following swap space size:

RAM size	Swap space size
1GB – 2GB RAM	1.5 x the RAM size
2GB – 8GB RAM	1 x the RAM size
8GB RAM and higher	0.75 x the RAM size

To display the size of the configured swap space execute the following command as user root:

```
fdisk -l | grep swap
/dev/sda         33         554     4192965   82  Linux swap / Solaris
```

4.1.2 The Linux operating system user for Oracle

The Oracle 11g installation is usually performed using a Linux operating system user. This user needs certain environment variable settings to make sure that the Oracle database is set up and functions correctly. This section describes how to create such a user for Oracle, and how to configure this user's environment.

4.1.2.1 Creating the Linux operating system user for Oracle

The following local operating system groups and users are required to install an Oracle database:

- The Oracle Inventory group (`oinstall`)
- The operating system database administrator group (`dba`)
- The Oracle software owner/user (`oracle`)

> Note: The group and user names are a recommendation from the Oracle Installation Guide for Linux. We will follow the recommendation in this case.

To determine if these groups and users already exist, and if necessary, to create them, follow the steps below.
To determine if the group `oinstall` exists, enter the following command as user `root`:

```
grep oinstall /etc/group
```

If the output of this command contains the group name `oinstall`, then the group already exists.

To determine if the group `dba` exists, enter the following command as user `root`:

```
grep dba /etc/group
```

If the output from this command contains the group name `dba`, then the group already exists.

If you need to create these groups, enter the following commands as user `root`:

```
groupadd oinstall
groupadd dba
```

To determine if the user `oracle` exists and belongs to the correct groups, enter the following command as user `root`:

```
id oracle
```

If the user `oracle` exists, this command will display information about the groups to which the user belongs. The output should be similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

This indicates that oinstall is the primary group and dba is a secondary group. Some hints and tips if you need to adapt your settings are described in the following.

If the oracle user exists, but its primary group is not oinstall or it is not a member of the group dba, enter the following command as user root:

```
usermod -g oinstall -G dba oracle
```

If the user `oracle` does not exist, enter the following command as user root to create it:

```
useradd -g oinstall -G dba oracle
```

This command creates the `oracle` user and specifies `oinstall` as the primary group and `dba` as the secondary group.

Enter the following command to set the password of the oracle user:

```
passwd oracle
```

4.1.2.2 Modifying the profile of user oracle

After creating the operating system user `oracle`, the user's profile needs to be changed. To do this, log in as user `oracle` and change or add the following entries in `/home/oracle/.bash_profile`:

```
export ORACLE_BASE=/opt/oracle
export ORACLE_HOME=$ORACLE_BASE/11g
export ORACLE_PATH=$ORACLE_BASE/common/oracle/sql:.:$ORACLE_HOME/rdbms/admin
export ORACLE_SID=ORCL

export PATH=.:${PATH}:$HOME/bin:$ORACLE_HOME/bin
export PATH=${PATH}:/usr/bin:/bin:/usr/bin/X11:/usr/local/bin
export PATH=${PATH}:$ORACLE_BASE/common/oracle/bin
export ORACLE_TERM=xterm
export TNS_ADMIN=$ORACLE_HOME/network/admin
export ORA_NLS10=$ORACLE_HOME/nls/data
export LD_LIBRARY_PATH=$ORACLE_HOME/lib
export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:$ORACLE_HOME/oracm/lib
export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:/lib:/usr/lib:/usr/local/lib
export CLASSPATH=$ORACLE_HOME/JRE
export CLASSPATH=${CLASSPATH}:$ORACLE_HOME/jlib
export CLASSPATH=${CLASSPATH}:$ORACLE_HOME/rdbms/jlib
export CLASSPATH=${CLASSPATH}:$ORACLE_HOME/network/jlib
export THREADS_FLAG=native
export TEMP=/tmp
export TMPDIR=/tmp
```

The environment variables **ORACLE_BASE** and **ORACLE_HOME** can be adapted in order to match a different Oracle installation directory.

> **Note:** In this document the defaults are used for the environment variables ORACLE_BASE and ORACLE_HOME. If other values are used the screens shown in 4.2 Installing the Oracle binaries will reflect those values instead of the defaults.

The environment variable **ORACLE_SID** defines the Oracle System ID for the particular database on the system. Set this system variable to match the value that will be used later in the database creation process.

4.1.3 Creating the directories for Oracle

The Oracle user needs a base directory for the software binaries and the database files. The user must have read and write rights in this directory.
Create this directory to match the path defined in the previous chapter in the system variable **ORACLE_BASE** (e.g. /opt/oracle).

To create the Oracle base directory execute the following commands as user `root`:

```
mkdir -p /opt/oracle
chown -R oracle:oinstall /opt/oracle
chmod -R 775 /opt/oracle
```

4.2 Installing the Oracle binaries

This chapter describes the Oracle 11g binary installation on a Red Hat Enterprise Linux 5 (RHEL5) operating system. The Oracle binary installation will be performed in silent mode using an response file. General information about how to install Oracle using a response file can be found in the **Oracle® Database Installation Guide 11g Release 1 (11.1) for Linux** Part Number B32002-06.

Before installing the binaries, make sure that the binaries for Oracle 11g have been unzipped in the `/tmp/db11106` directory or the Software CD/DVD is mounted in the system.
In this example the unzipped variant will be used. Please refer to the Linux man-page how to mount the CD/DVD and how to access the file on it.

Note: Make sure that you use Oracle 32-bit binaries on a 32-bit Linux system, and Oracle 64-bit binaries on a 64-bit system.

The following are the general steps to install and configure Oracle products using Oracle Universal Installer in silent:

1. Create the oraInst.loc file.
2. Prepare a response file.
3. Run Oracle Universal Installer in silent mode (Software-only).
4. Run Oracle Patch Set in silent mode.
5. Run Net Configuration Assistant in silent mode
6. Run Database Configuration Assistant in silent mode.

These steps are described in the following sections.

Note: The steps in the following sections imply that the user `oracle` is logged in on the database host.

4.2.1 Creating the oraInst.loc File

This file specifies the location of the Oracle Inventory directory where Oracle Universal Installer creates the inventory of Oracle products installed on the system.

Note: If Oracle software has been installed previously on the system, the oraInst.loc file might already exist. If the file does exist, you do not need to create a file.

To create the oraInst.loc file, follow these steps:

1. Switch user to `root`:

```
su - root
```

2. Change directory to /etc:

```
cd /etc/
```

3. Use a text editor to create the `oraInst.loc` file, containing the following lines:

```
vi oraInst.loc
```

```
inventory_loc=$ORACLE_BASE/oraInventory
inst_group=oinstall
```

In this example, $ORACLE_BASE is the path of the Oracle base directory, for example, /opt/oracle. Type in a absolute path for ORACLE_BASE, do not use variables.

4. Enter the following commands to set the appropriate owner, group, and permissions on the oraInst.loc file:

```
chown oracle:oinstall oraInst.loc
chmod 664 oraInst.loc
exit
```

4.2.2 Editing the Response File Template

Oracle provides response file templates for each product and installation type, and for each configuration tool. These files are located at the `{db_install_dir}/response` directory. In this example the response files are located in `/tmp/db11106/response` and it will use the `enterprise.rsp` for the silent installation.

To modify the enterprise.rsp response file finish the following steps:

1. Open the response file in a text editor:

```
vi /tmp/db11106/response/enterprise.rsp
```

In addition to editing settings specific to the Oracle Database installation, check that the FROM_LOCATION path is correct and points to the products.xml file in the stage directory in the installation media. You may want to set this variable to point to an absolute path, for example:

```
FROM_LOCATION="/directory_path/stage/products.xml"
```

Remember that sensitive information, such as passwords, can be specified at the command line rather than within the response file. "How Response Files Work?" explains this method. **See Also:** Oracle Universal Installer and OPatch User's Guide for detailed information on creating response files

2. Follow the instructions in the file to edit it.

Note: Oracle Universal Installer or configuration assistant fails if the response file is not configured correctly. Refer to "Silent-Mode Response File Error Handling" section for more information about troubleshooting a failed silent-mode installation.

Change the variables in the `enterprise.rsp` response file to the following values:

```
UNIX_GROUP_NAME="oinstall"
FROM_LOCATION="/tmp/db11106/stage/products.xml"
ORACLE_BASE="/opt/oracle"
ORACLE_HOME="/opt/oracle/11g"
ORACLE_HOME_NAME="OraDbHome1"
RESTART_SYSTEM=false
RESTART_REMOTE_SYSTEM=false
ORACLE_HOSTNAME="w620l13o.boeblingen.de.ibm.com"
INSTALL_TYPE="EE"
s_nameForDBAGrp="dba"
```

```
s_nameForOPERGrp="dba"
s_nameForASMGrp="oinstall"
n_configurationOption=3
```

4.2.3 Running Oracle Universal Installer using a response file

The Oracle Universal Installer executable, runInstaller, provides several options. For help information on the full set of these options, run the runInstaller command with the -help option:

```
/tmp/db11106/runInstaller -help
```

The help information appears in a window after some time.

To run Oracle Universal Installer using a response file:

1. Complete the preinstallation tasks listed in 4.2.2 Editing the Response File Template.

2. Log in as the Oracle software owner user (in this case oracle).

3. To start Oracle Universal Installer in silent mode, enter the following command:

```
/tmp/db11106/runInstaller -silent -noconfig -responseFile /tmp/db11106/response/enterprise.rsp
```

Note: Do not specify a relative path to the response file. If you specify a relative path, then Oracle Universal Installer fails.

Wait until the installation is finished.
If any errors occur, refer to
Oracle® Database Installation Guide 11g Release 1 (11.1) for Linux
Part Number B32002-06:

http://download.oracle.com/docs/cd/B28359_01/install.111/b32002/toc.htm

4. When the installation completes, log in as the root user and run the root.sh script:

```
su - root
password:
```

```
/opt/oracle/11g/root.sh
```

4.2.4 Installing the Oracle 11.1.0.7.0 patch set

This chapter describes the Oracle 11g patch set installation on a Red Hat Enterprise Linux 5 (RHEL5) operating system. The Oracle patch set installation will be performed in silent mode using an response file. In this case the Oracle patch set 11.1.0.7.0 will be installed.

> **Note:** The Oracle 11g (11.1.0.6.0) can be used for WebSphere Process Server 6.2.0. This step can be skipped if the Oracle Patch Set is not needed.

Finish the following steps to install the Oracle 11g Patch Set 11.1.0.7.0:

1. Download the Oracle 11g Patch Set 11.1.0.7.0 for Linux from the Oracle Support Homepage (Metalink).

```
https://metalink.oracle.com/CSP/ui/index.html
```

Copy the downloaded patch set file to the database host. Make sure, that the oracle user has

2. Unzip the downloaded Patch Set to `/tmp/patch11107`.

Login on the database host as user `oracle` and unzip the downloaded patch set file.

3. Edit the response file template.

To edit the response file template execute the following command:

```
vi /tmp/patch11107/response/patchset.rsp
```

Change the variables in the `patchset.rsp` response file to the following values:

```
UNIX_GROUP_NAME="oinstall"
FROM_LOCATION="/tmp/patch11107/stage/products.xml"
ORACLE_HOME="/opt/oracle/11g"
ORACLE_HOME_NAME="OraDbHome1"
RESTART_SYSTEM=false_HOSTNAME="w620113o.boeblingen.de.ibm.com"
METALINK_USERNAME=""
```

4. Run the Oracle Universal Installer using the patchset.rsp response file:

Execute the following command to start the oracle universal installer:

```
/tmp/patch11107/runInstaller -silent -responseFile
/tmp/patch11107/response/patchset.rsp
```

5. When the installation completes, log in as the root user and run the root.sh script:

```
su - root
password:
/opt/oracle/11g/root.sh
```

Chapter 5 Oracle database configuration

5.1 Creating the Oracle listener using a Response File

To configure and start an Oracle Net listener on the system, configure naming methods, and configure Oracle Net service names run Net Configuration Assistant in silent mode . To run Net Configuration Assistant in silent mode, edit a response file template. Oracle provides a response file template named netca.rsp in the response directory /tmp/db11106/response.

To run Net Configuration Assistant using a response file:

1. Open the response file in a text editor:

```
vi /tmp/db11106/response/netca.rsp
```

Change the variables in the netca.rsp response file to the following values:

```
SHOW_GUI=false
```

2. Log in as the Oracle software owner user (in this case oracle).

3. Run the network configuration assistent using the netca.rsp response file:

Execute the following command to start the oracle universal installer:

```
/opt/oracle/11g/bin/netca /silent /responsefile /tmp/db11106/response/netca.rsp
```

Wait until the installation is finished.
If any errors occur, refer to
Oracle® Database Installation Guide 11g Release 1 (11.1) for Linux
Part Number B32002-06:

http://download.oracle.com/docs/cd/B28359_01/install.111/b32002/toc.htm

5.1.1 Verifying the created listener

To verify the created listener:

1. log in with a separate session as user oracle.
2. and execute the following command:

```
lsnrctl status
```

The output will be like:

```
LSNRCTL for Linux: Version 11.1.0.7.0 - Production on 30-APR-2009 15:45:36

Copyright (c) 1991, 2008, Oracle.  All rights reserved.

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=EXTPROC1523)))
STATUS of the LISTENER
------------------------
Alias                     LISTENER
Version                   TNSLSNR for Linux: Version 11.1.0.7.0 -
Production
Start Date                30-APR-2009 15:45:30
Uptime                    0 days 0 hr. 0 min. 5 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Listener Parameter File   /opt/oracle/11g/network/admin/listener.ora
Listener Log File
/opt/oracle/diag/tnslsnr/W6201L3O/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)
(HOST=W6201L3O.boeblingen.de.ibm.com)(PORT=1521)))
The listener supports no services
The command completed successfully
```

5.2 Creating the WebSphere Process Server database

Start the database configuration assistant by executing the `/opt/oracle/11g/bin/dbca` command as user **oracle** to create the WebSphere Process Server database.

> **Note:** When X11 content is forwarded to a Windows system using an outdated X-Windows terminal emulators (or X *Server)*), the Oracle installation wizard may be displayed incorrectly. It is highly recommended to use one of the latest X Servers, for example the open source.Xming 6.9.0.23

For more information about Xming refer to

`http://sourceforge.net/projects/xming.`

The "Operations" panel is displayed:

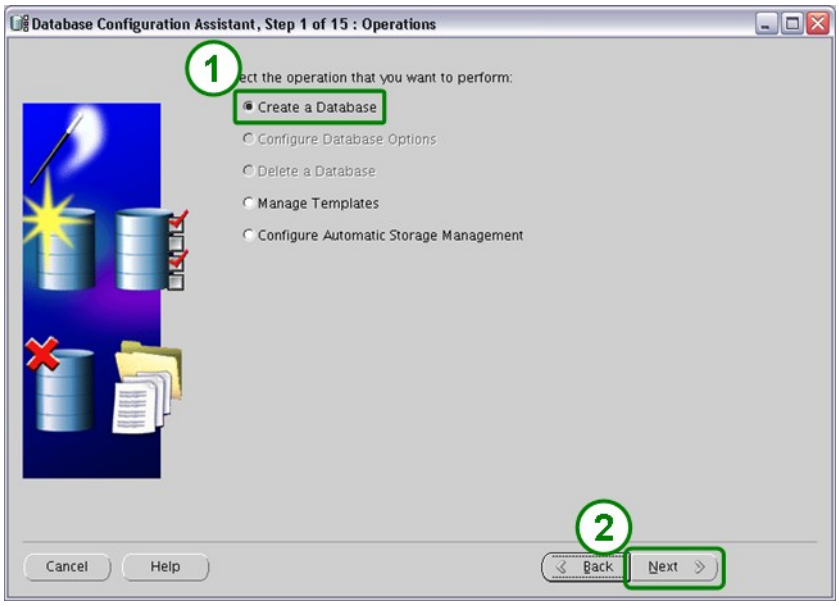

1. Select "Create aDatabase"

2. Press the **Next** button.

The “Database Templates” panel is displayed:

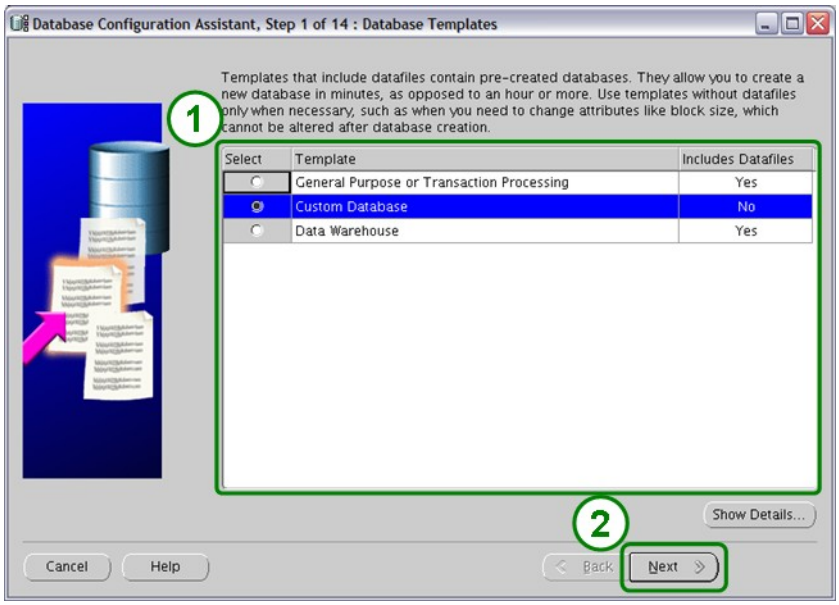

1. By default there are three database templates. Select the “Custom Database”.

2. Press the **Next** button.

The “Database Identification” panel is displayed:

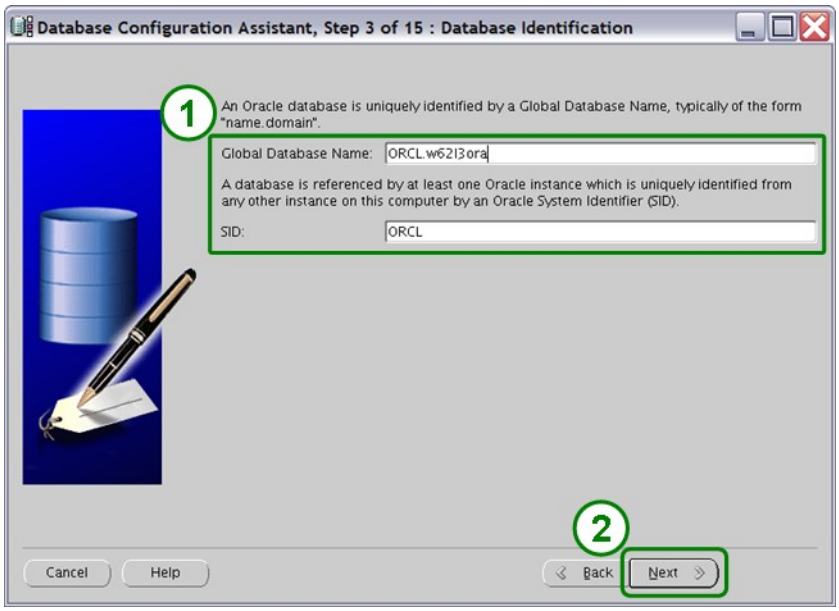

1. Specify the “Global Database Name and the System ID (SID).

The “Global Database Name” usually consists of the SID followed by the database domain.

The hostname of the database will be used as the database domain.

For example the database with the SID “ORCL” on host “hostname” would get the Global Database Name “ORCL.hostname”.

2. Press the **Next** button.

The “Management Options” panel is displayed:

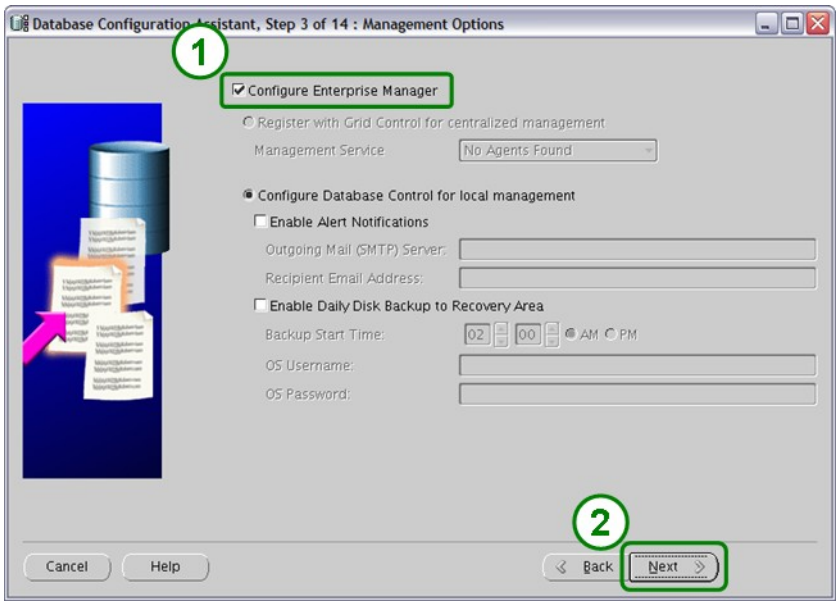

1. Select “Configure Enterprise Manager” to add the Oracle Enterprise Repository to the database.

It is recommended to choose this option unless Oracle Grid Control which is the Oracle Enterprise Infrastructure for managing multiple databases is used.

This will make database monitoring and tuning in the future much more easier and comfortable.

2. Press the **Next** button.

The “Database Credentials” panel is displayed:

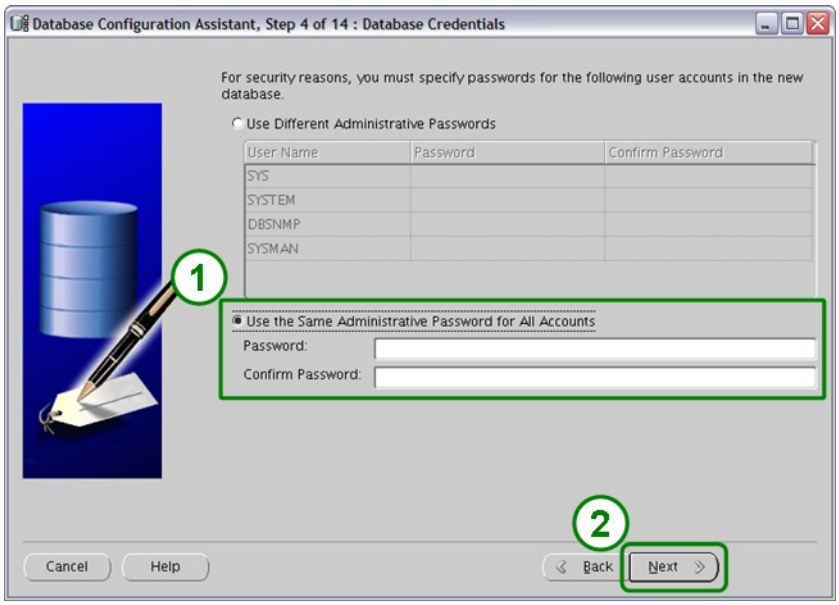

1. Select “Use the Same Administrative Password for All Account” and type in a preferred password.

2. Press the **Next** button.

The “Storage Options” panel is displayed:

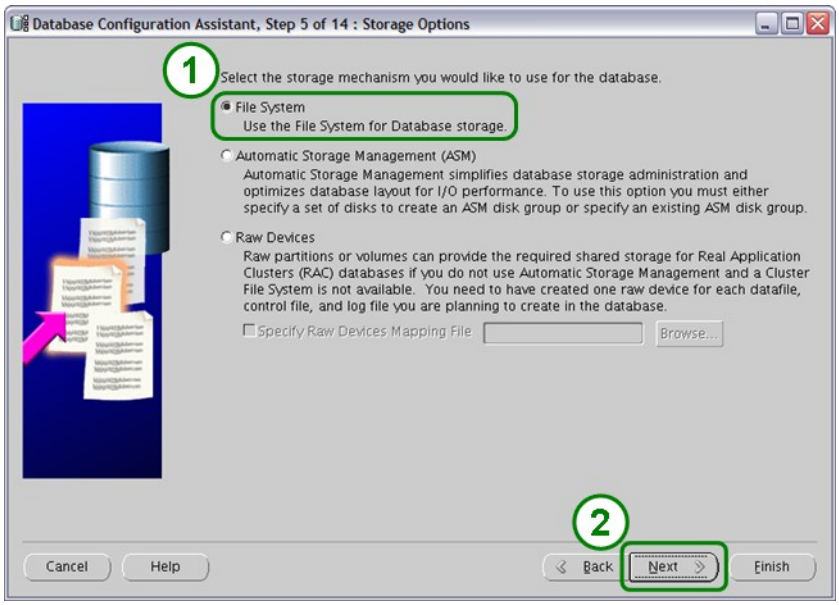

The Oracle storage options:

File System is the simplest way and only a little slower as the other mechanisms.

Automatic Storage Management (ASM) is new since Oracle 10g. Using ASM oracle accesses directly to the devices. In contrast to “Raw Devices” the disk space is managed by oracle.

Raw Devices have the advantage, that Oracle accesses directly to the raw partitions or volumes without having the operating system as indirection. The administration is very extensive, because a raw device needs to be created each datafile, control file, and log file that is planed to be created.

1. Select “File System” as the storage option.

2. Press the **Next** button.

The "Database File Locations" panel is displayed:

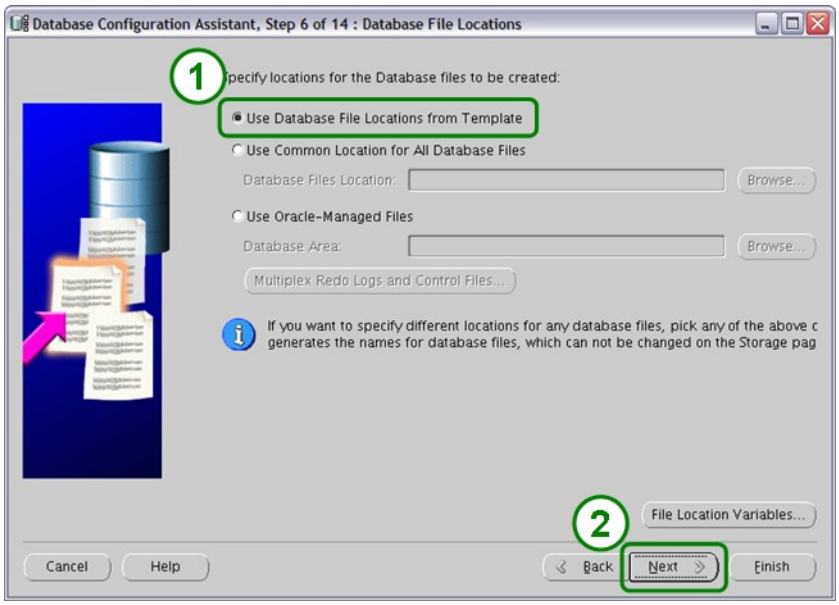

Oracle Database File Location options:

"**Database File Locations from Template**" gives the most control over the location where the database files will be created.

Choosing "**Common Location for All Database Files**", will create all database files in the same directory.

The mechanism of "**Oracle-Managed Files**" means, that Oracle is creating all database files in the database area. The database area is a simple directory. All database files will be created in that directory. In contrast to the "**Common Location**" Oracle assigns the name of the database files. This could be an advantage if a database file is often added to a tablespace. Simply the tablespace and the size have to be specified. Selecting "**Multiplex redo logs and Control Files**" stores a copy of all redo log files and a control file in the Flash Recovery Area. This area was introduced with Oracle 10g.

1. Select "Use Database File Locations from Template".

2. Press the **Next** button.

The "Recovery Configuration" panel is displayed:

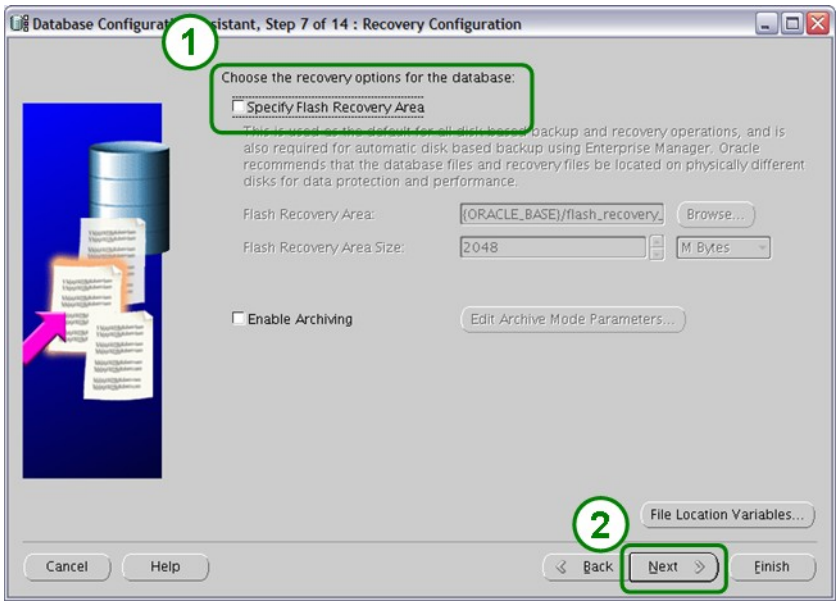

Flash Recovery Area is not needed if archiving is disabled and/or redo logs and control files are not multiplexed.

1. Deselect “Specify Flash Recovery Area”

2. Press the **Next** button.

The "Database Content" panel is displayed:

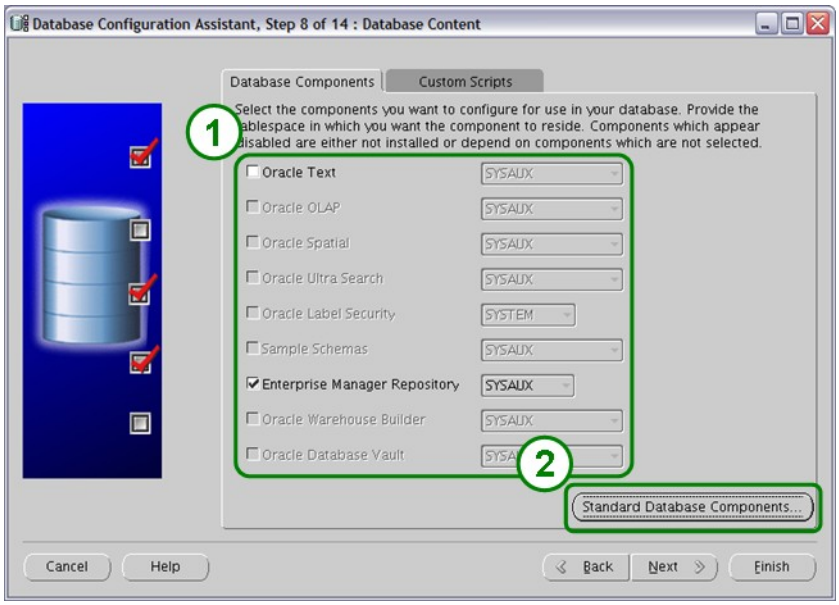

1. Deselect all database components except of "Enterprise Manager Repository"

2. Press the **Standard Database Components** button.

The "Standard Database Components" panel pops up:

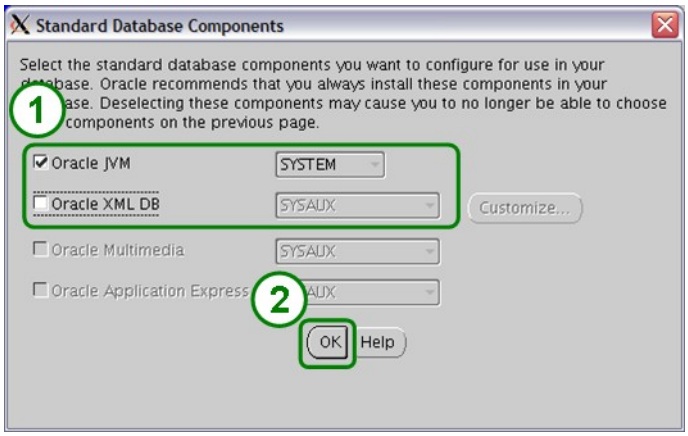

1. Deselect all standard database components except of "Oracle JVM".

2. Press the **OK** button.

The "Standard Database Components" is closed and the "Database Content" panel gets the focus back:

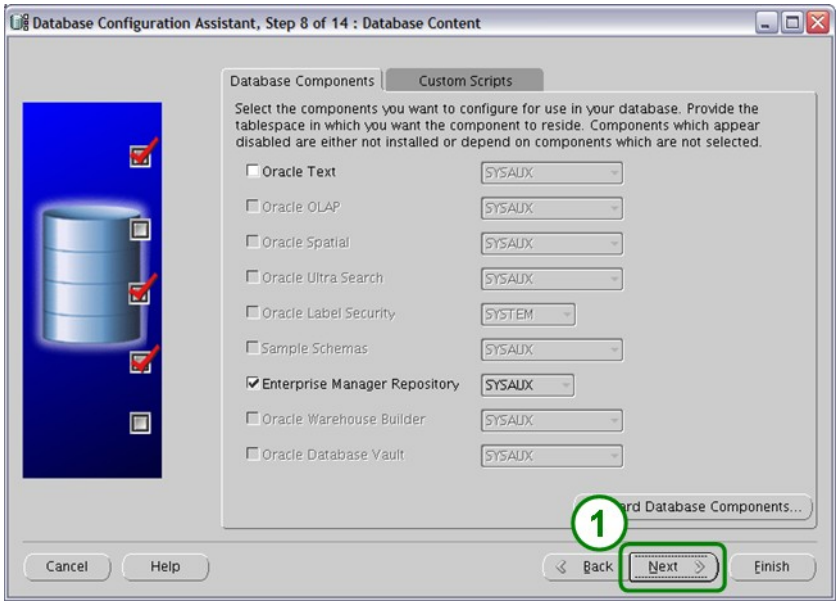

1. Press the **Next** button.

The "Initialization Parameters" panel with open "Memory" tab is displayed:

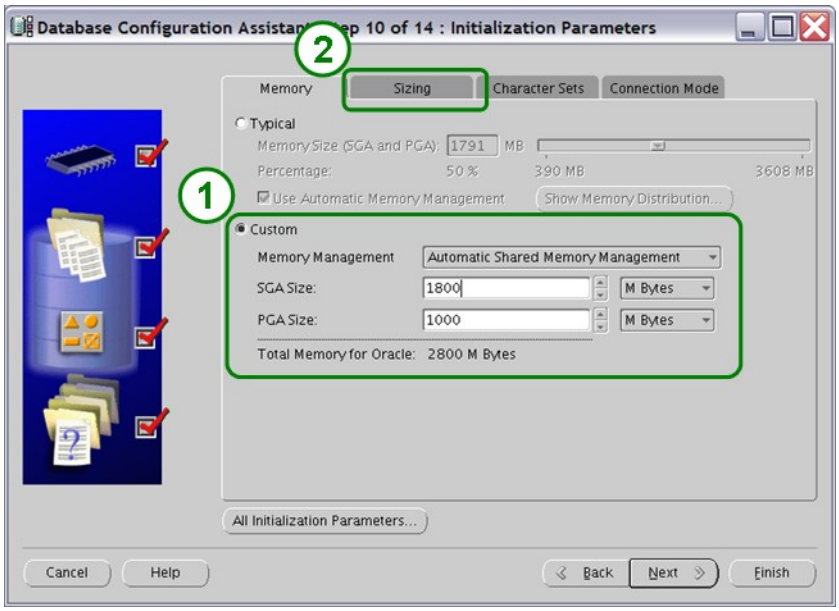

The memory management parameters options are:

"Typical" defines one memory size for both the System Global Area (SGA) and the Program Global Area (PGA) in one value and splits it automatically into SGA and PGA.

"Custom" defines the SGA and PGA memory size separately. The SGA size can not be set higher than defined in the kernel parameter kernel.shmmax.

1. Select "Custom" and type in the chosen SGA and PGA memory size.

2. Select the **Sizing** tab.

Note: When choosing the SGA and PGA size make sure, that the database host has enough memory and will not start swapping. Swapping will extremely slow the database.

Please also refer to the Oracle Installation Guide for Linux.

The "Initialization Parameters" panel with open "Sizing" tab is displayed:

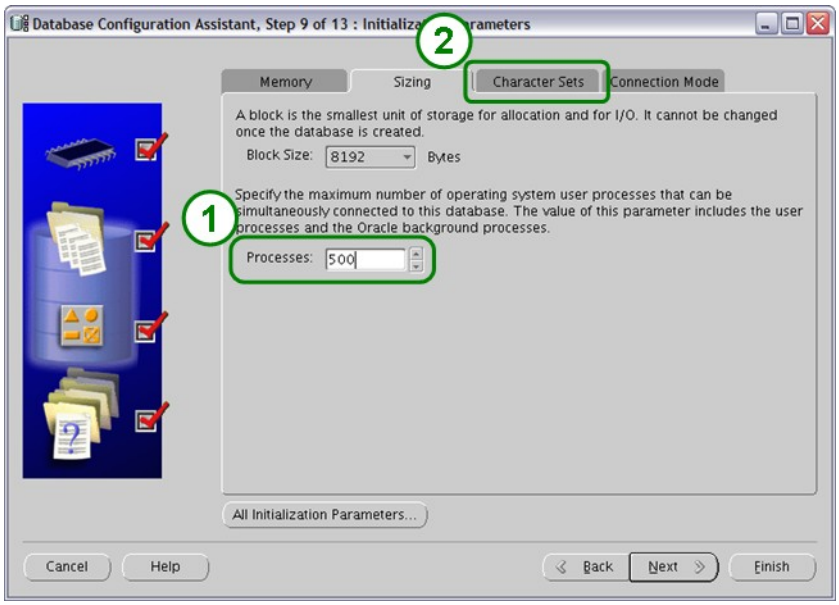

On the "Sizing" tab increase the number of processes. This parameter limits the number of parallel database connections if "Dedicated Server" is used as connection mode. It depends on the BPEL application and the number of process instances that are running in parallel.

1. Change the processes from 150 to 500 as first initial tuning action.

2. Select the **Character Sets** tab.

The "Initialization Parameters" panel with open "Character Sets" tab is displayed:

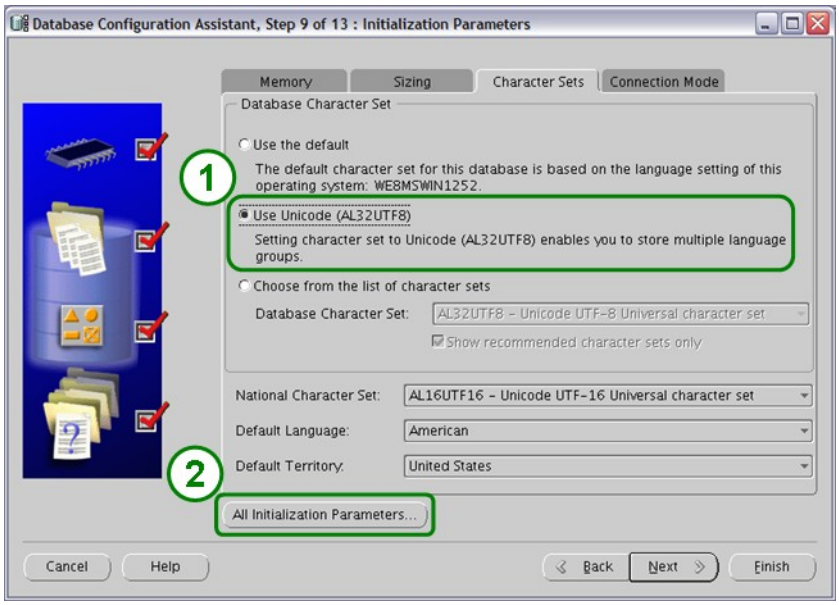

WebSphere Process Server needs a UTF8 database as a data store.

1. Select "Use Unicode (AL32UTF8)

2. Press the **All Initialization Parameters** button.

The "All Initialization Parameters" panel pops up:

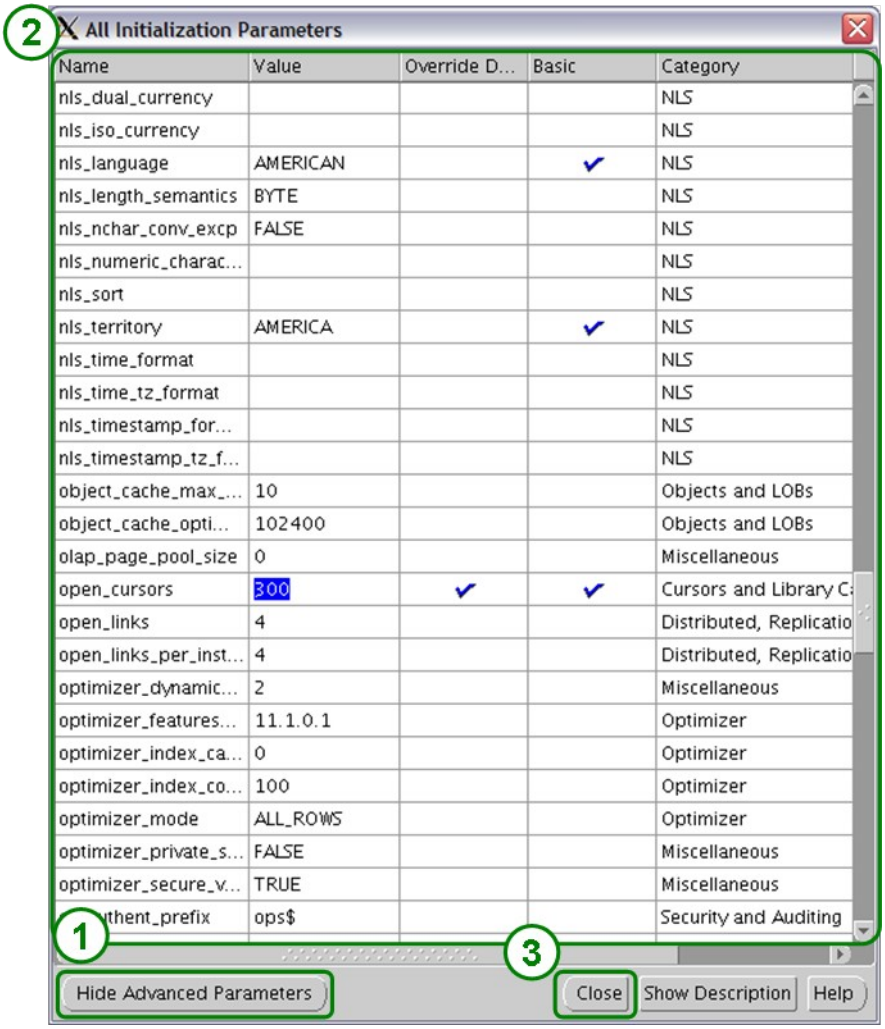

1. Press the **Show Advanced Parameters** button.
2. Change the Parameters as shown in the Table on the next side.
3. Press the **Close** button.

Parameter	*default*	*new value*	Remark
OPEN_CURSORS	300	1000	Specifies the maximum number of open cursors (context areas) a session can have at once, and constrains the PL/SQL cursor cache size which PL/SQL uses to avoid reparsing statements re-executed by a user. Set this value high enough to prevent your applications from running out of open cursors.
SESSION_CACHED_CURSORS	50	1000	Specifies the number of session cursors to cache. When the same SQL statement is parses several times, it's session cursor is moved into the session cursor cache.
FAST_START_MTTR_TARGET		900	This parameter specifies the mean time in seconds that the database should be able to recover if a crash occurs. Possible values are from 0 to 3600.

The "All Initialization Parameters" is closed and the "Initialization Parameters" panel gets the focus back:

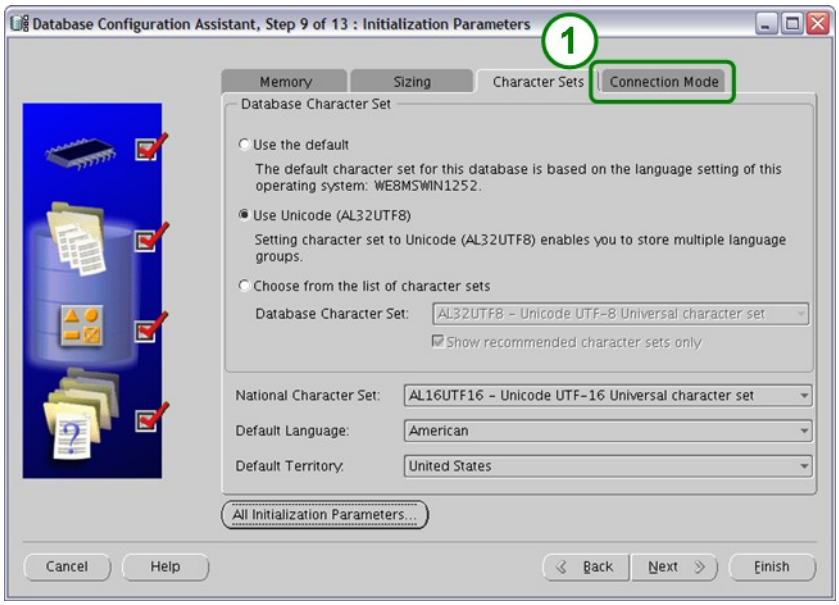

1. Select the **Connection Mode** tab.

The "Initialization Parameters" panel with open "Connection Mode" tab is displayed:

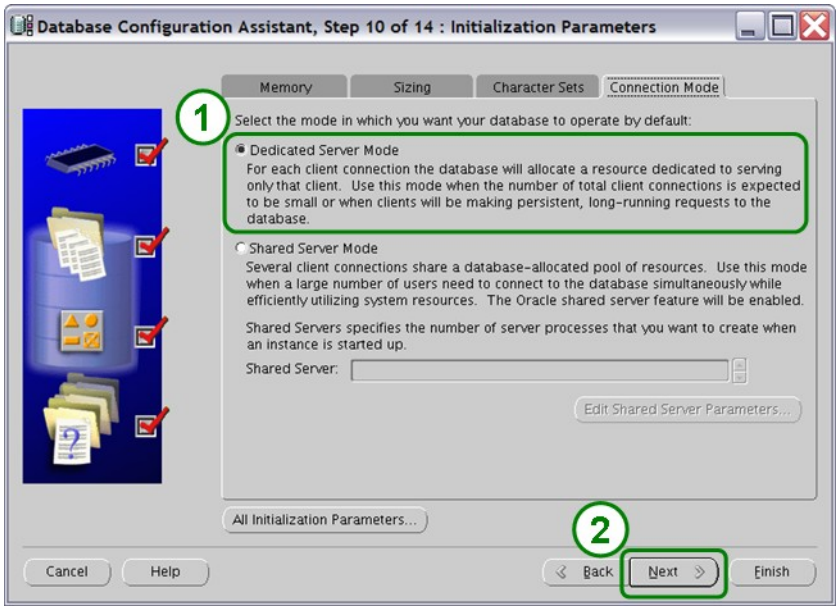

1. Verify that "Dedicated Server Mode" is selected.

2. Press the **Next** button.

The "Security Settings" panel is displayed:

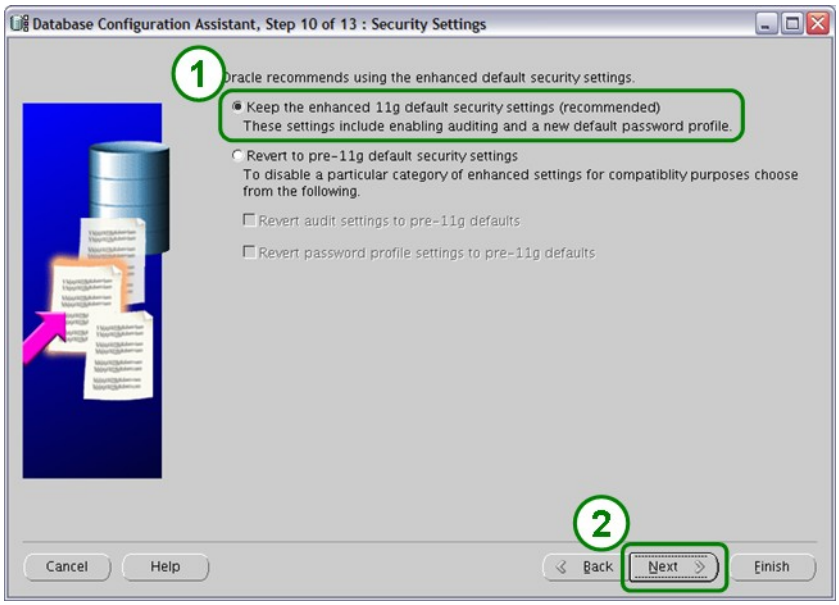

1. Select “Keep the enhanced 11g default security settings”.

2. Press the **Next** button.

The "Automatic Maintenance Tasks" panel is displayed:

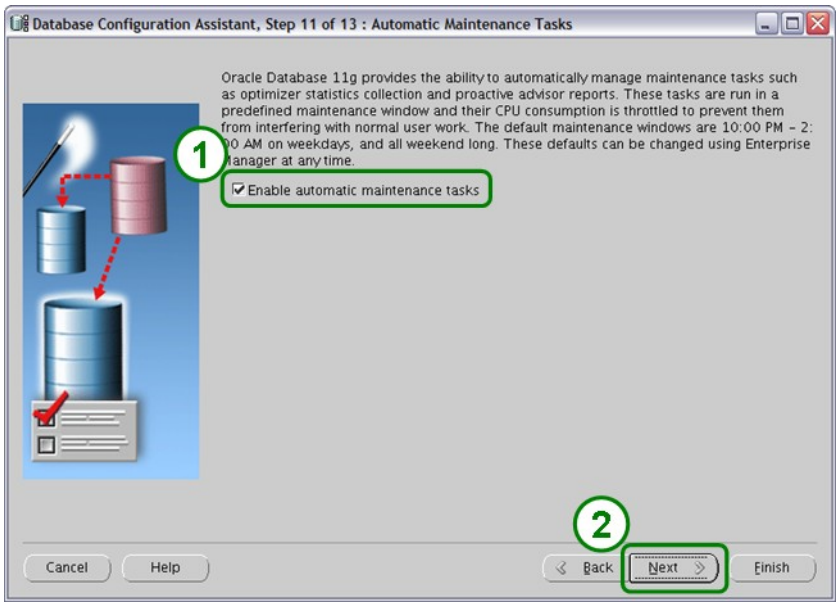

Oracle 11g has the possibility to run scheduled maintenance tasks like compute statistics and table space advisor.

1. Select “Enable automatic maintenance tasks”.

2. Press the **Next** button.

The "Database Storage" panel is displayed:

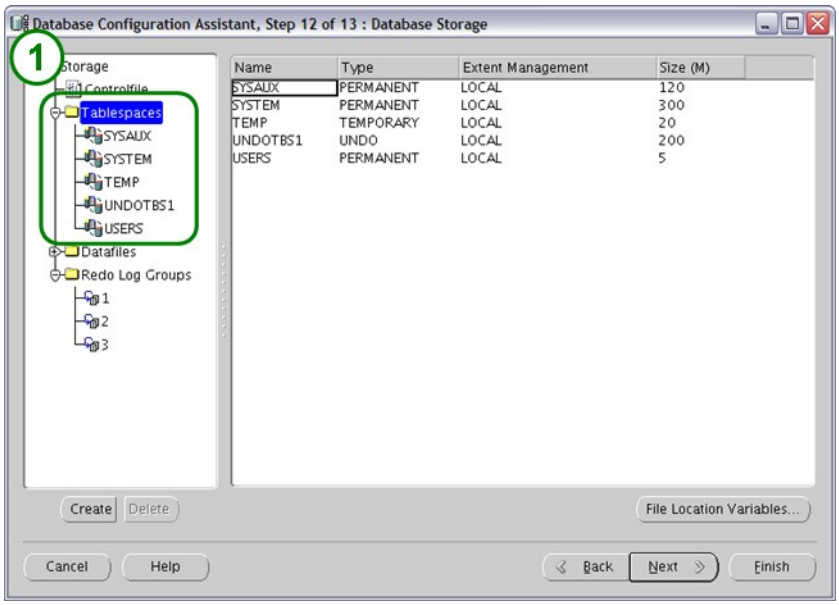

The default size of the database files of the tablespaces is too small; increase the sizes to the following values:

SYSAUX	500 MB
SYSTEM	700 MB
TEMP	200 MB
UNDOTBS1	1500 MB
USERS	150 MB

1. Select the tablespace SYSAUX.

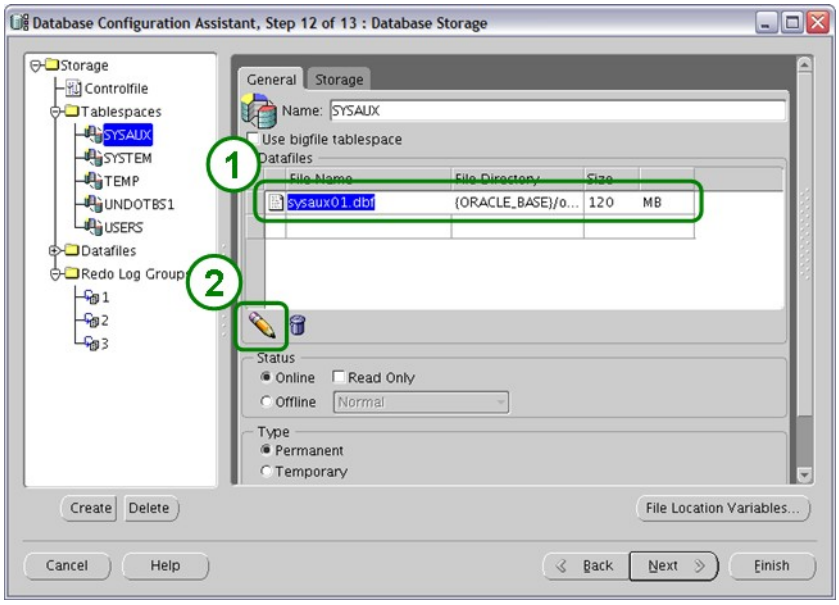

1. Select the datafile.

2. Press on the Pencil icon to edit the datafile.

The "Edit Datafile" panel pops up:

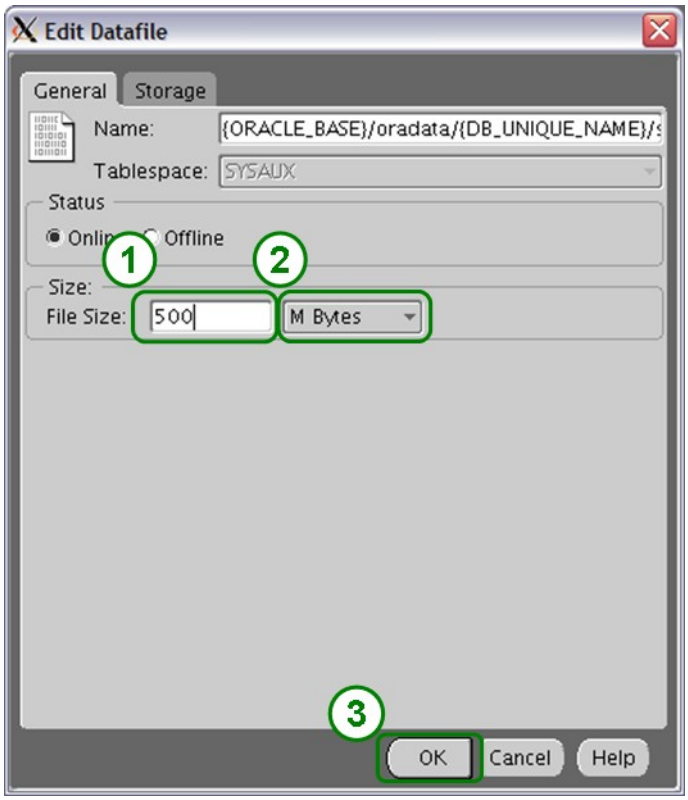

1. Change the file size to the value above.

2. Make sure that “M Bytes” is selected.

3. Press the **OK** button to confirm and save the changes.

The "Edit Datafile" is closed and the "Database Storage" panel gets the focus back:

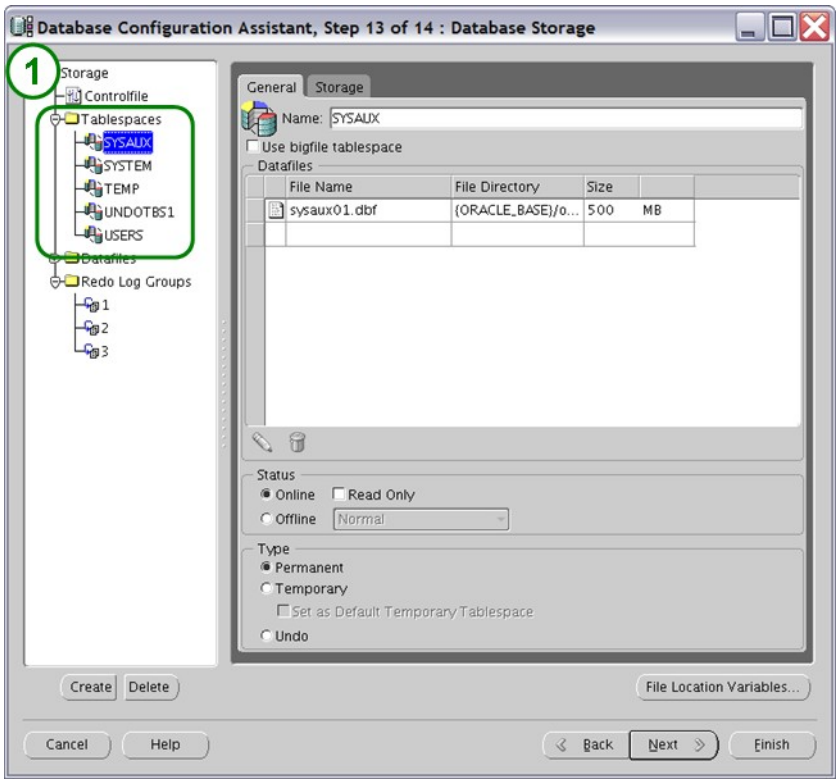

1. Repeat the last steps for the SYSTEM, TEMP, UNDOTBS1 and USERS table spaces to increase the size as listed above.

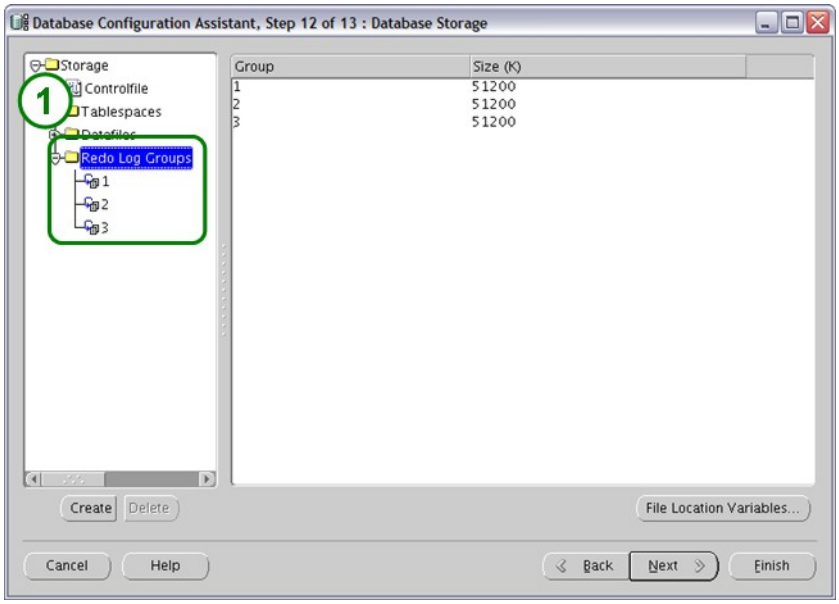

Specify the size and the location of the redo log files. The optimal size depends on the database load. To increase the performance of the database increase the size at least to **1500 MByte**. All redo log files should have the same size.

1. Select the first Redo Log Group.

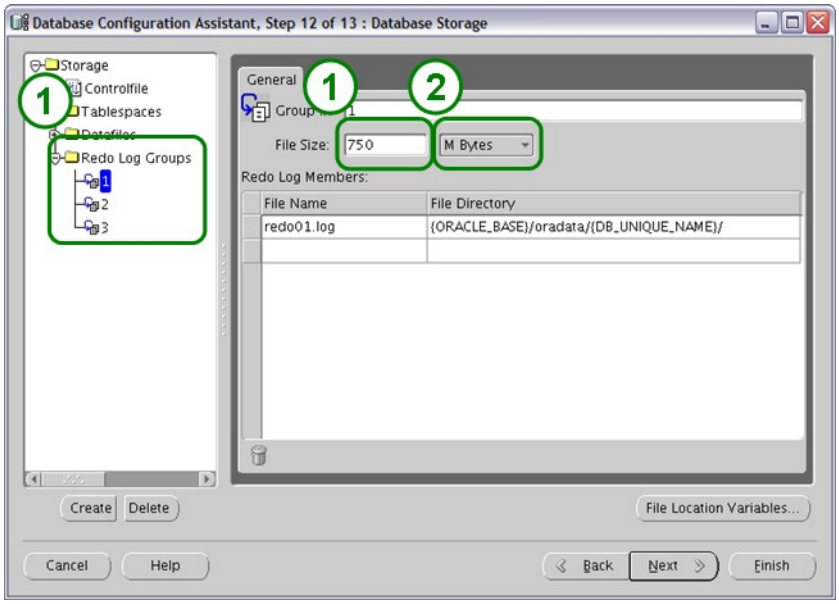

1. Insert the value 1500 in the file size field.

2. Make sure that "M Bytes" is selected.

3. Select the two other Redo Log Groups one after the other and repeat the steps.

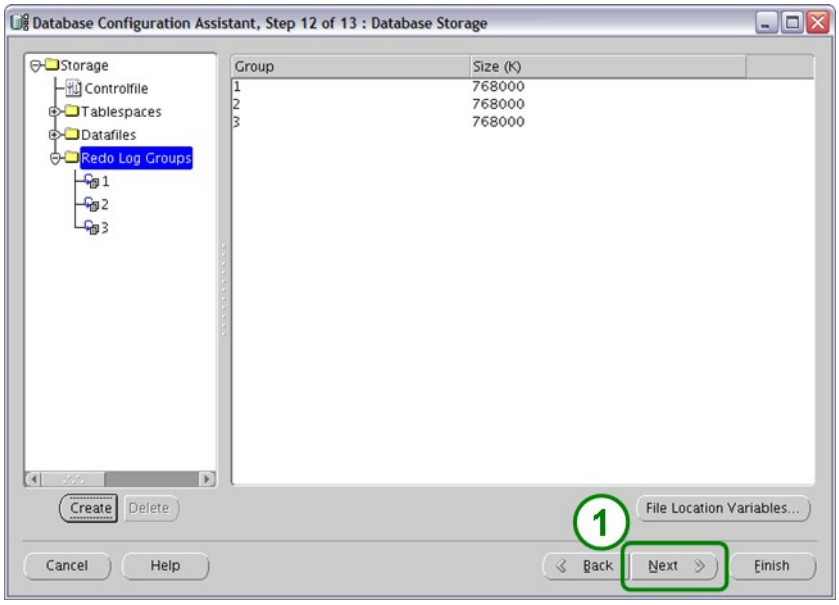

1. Press the **Next** button.

The "Creation Options" panel is displayed:

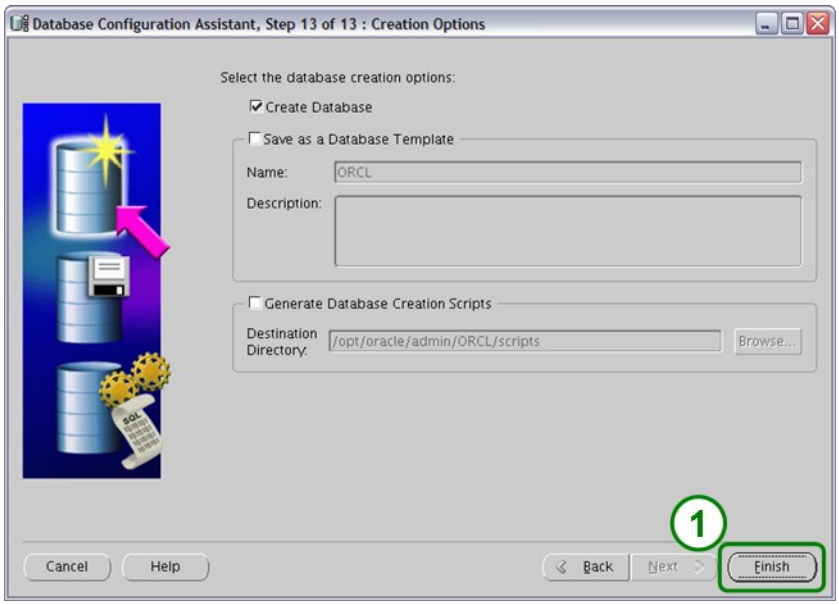

1. Press the **Finish** button.

The "Confirmation" panel pops up:

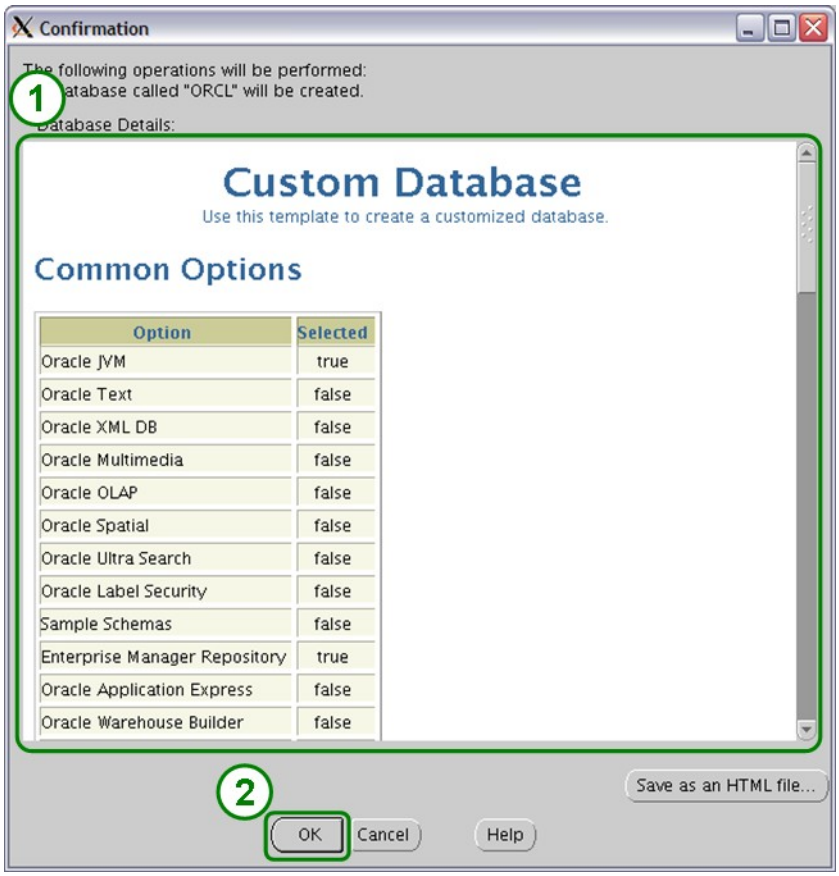

1. Verify the configured Oracle database options.

2. Press the **OK** button.

The "Database Configuration Assistant" starts:

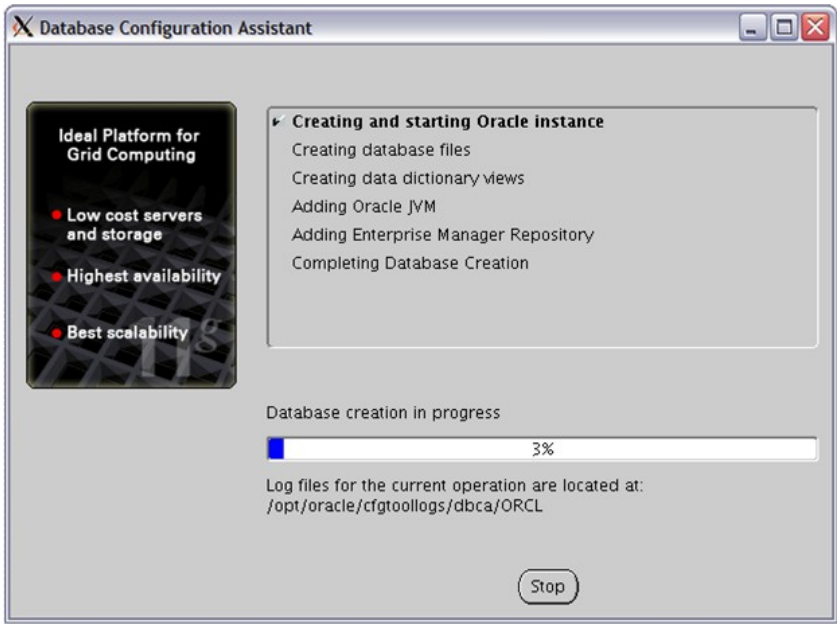

Wait until the database creation is finished.

1. Press the **Exit** button.

5.3 Database listener configuration

Login as user oracle. Goto the directory /opt/oracle/11g/network/admin and verify the settings in the file listener.ora:

```
# listener.ora Network Configuration File:
/opt/oracle/11g/network/admin/listener.ora
# Generated by Oracle configuration tools.

LISTENER =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS_LIST =
        (ADDRESS = (PROTOCOL = TCP)(HOST = <your_hostname>)(PORT = 1521))
      )
    )
  )
```

Make the necessary adjustments for the installation path und hostname.

The Oracle 11g listener needs the entry of an SID_LIST_LISTENER in the listerner.ora file. Add the following lines to the listener.ora file:

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ORCL)
      (ORACLE_HOME = /opt/oracle/11g)
    )
  )
```

The Oracle listener will discover the database instance ORCL automatically after a time period of about 1 minute.

The file tnsnames.ora located in the same directory has also to be checked. Add the following lines to the tnsnames.ora file if a similar entry is missing:

```
ORCL =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = <your_hostname>)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = ORCL)
    )
  )
```

Change <your_hostname> to the name of your Oracle database host.

Restart the listener after changing the files.

> Note: The commands to start and stop the listener are listed in the Oracle Appendix in this document.

5.4 Configuring database initialization parameters

During the creation of the database the right initialization parameters were already set.

> **Note:** Skip this chapter if a new database was created during the Oracle installation process and continue with chapter 5.5.Oracle JDBC driver.

Modify the initialization parameters for a existing database to match the "new value" size as in the table below:

Parameter	***default***	***new value***	**Remark**
OPEN_CURSORS	300	1000	Specifies the maximum number of open cursors (context areas) a session can have at once, and constrains the PL/SQL cursor cache size which PL/SQL uses to avoid reparsing statements re-executed by a user. Set this value high enough to prevent your applications from running out of open cursors.
SESSION_CACHED_CURSORS	50	1000	Specifies the number of session cursors to cache. When the same SQL statement is parses several times, it's session cursor is moved into the session cursor cache.
FAST_START_MTTR_TARGET		900	This parameter specifies the mean time in seconds that the database should be able to recover if a crash occurs. Possible values are from 0 to 3600.
PROCESSES	150	500	This parameter limits the number of parallel database connections if you are using "Dedicated Server" as connection mode. It depends on the BPEL application and the number of process instances that are running in parallel.

To change the parameters connect to the database ORCL as SYSDBA and use the ALTER SYSTEM command. The following script can be used to change the database initialization parameters.

Create a file `setWPSOraInitParameter.sql` with an editor and paste in the lines from below.

```
REM
*************************************************************************
REM File: setWPSOraInitParameter.sql
REM Date: 2009-01-15
REM
REM Desc: Set the initial Oracle database parameter for WPS.
REM
REM Usage:          Execute the sql script as user oracle on the database
host.
REM
*************************************************************************

ALTER SYSTEM SET OPEN_CURSORS=1000 SCOPE = spfile;
ALTER SYSTEM SET SESSION_CACHED_CURSORS=1000 SCOPE = spfile;
ALTER SYSTEM SET FAST_START_MTTR_TARGET=900 SCOPE = spfile;
ALTER SYSTEM SET PROCESSES=500 SCOPE = spfile;

REM Commit work
COMMIT;
EXIT
```

Execute the `setWPSOraInitParameter.sql` by typing the following command as user `oracle`:

```
sqlplus sys/<yourPassword>@ORCL AS SYSDBA @setWPSOraInitParameter.sql
```

Stop and restart the database after changing the initialization parameters.

> **Note:** How to start and stop the database please refer to the Oracle Appendix.

5.5 Oracle JDBC driver

This chapter describes where to find the needed Oracle JDBC drivers.

1. Create the directory /opt/oracle/driver on the Deployment manager host and on each WPS and WBM host.

2. Locate the jdbc driver files in the directory "$ORACLE_HOME/jdbc/lib" on the host where Oracle was installed.

3. Copy all files from the lib directory into the directory /opt/oracle/driver on the Deployment manager host and on each WPS and WBM host.

4. Go to the following Oracle website:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html

5. Download the latest ojdbc14.jar and the latest ojdbc14_g.jar from the Oracle 10g JDBC Drivers.

Note: At publishing time the latest driver release was version 10.2.0.4.0.

6. Copy the downloaded JDBC drivers in the directory /opt/oracle/driver on the Deployment manager host and on each WPS and WBM host.

Note: If a non-root WPS installation is done, make sure the driver directory is readable by the used WPS os user.